Last updated: December 24, 2010

COMPUTER NETWORKS

Performance and Quality of Service

IVAN MARSIC

Department of Electrical and Computer Engineering

Rutgers University

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Book website: http://www.ece.rutgers.edu/~marsic/books/QoS/

Preface

This book reviews modern computer networks with a particular focus on performance and quality of service. There is a need to look towards future, where wired and wireless/mobile networks will be mixed and where multimedia applications will play greater role. In reviewing these technologies, I put emphasis on underlying principles and core concepts, rather than meticulousness or completeness.

Audience

This book is designed for upper-division undergraduate and graduate courses in computer networking. It is intended primarily for learning, rather than reference. I also believe that the book's focus on basic concepts should be appealing to practitioners interested in the "whys" behind the commonly encountered networking technologies. I assume that the readers will have basic knowledge of probability and statistics, which are reviewed in the Appendix. Most concepts do not require mathematical sophistication beyond a first undergraduate course.

Most of us have a deep desire to understand logical cause-effect relationships in our world. However, some topics are either inherently difficult or poorly explained and they turn us off. I tried to write a computer networking book for the rest of us, one that has a light touch but is still substantial. I tried to present a serious material in a fun way so the reader may have fun and learn something nontrivial. I do not promise that it will be easy, but I hope it will be worth your effort.

Approach and Organization

In structuring the text, I faced the choice between logically grouping the topics vs. gradually identifying and addressing issues. The former creates a neater structure and the latter is more suitable for teaching new material. I compromised by opportunistically adopting both approaches. I tried to make every chapter self-contained, so that entire chapters can be skipped if necessary.

Chapter 1 reviews essential networking technologies. It is condensed but more technical than many current networking books. I tried to give an engineering overview that is sufficiently detailed but not too long. This chapter serves as the basis for the rest of the book.

Chapter 2 reviews the mechanisms for congestion control and avoidance in data networks. Most of these mechanisms are implemented in different variants of Transmission Control Protocol (TCP), which is the most popular Internet protocol.

Chapter 3 reviews requirements and solutions for multimedia networking.

Chapter 4 describes how network routers forward data packets. It also describes simple techniques for modeling queuing delays.

Chapter 5 describes router techniques for reducing or redistributing queuing delays across data packets. These include scheduling and policing the network traffic.

Chapter 6 describes wireless networks, focusing on the network and link layers, rather than on physical layer issues.

Chapter 7 describes network measurement techniques.

Chapter 8 describes major protocols used in the Internet that I are either not essential or are specific implementations of generic protocols presented in earlier chapters. The most essential Internet protocols, such as TCP and IP are presented in earlier chapters.

The Appendix provides a brief review of probability and statistics.

Solved Problems

This book puts great emphasis on problems for two reasons. First, I believe that specific problems are the best way to explain difficult concepts. Second, I wanted to keep the main text relatively short and focused on the main concepts; therefore, I use problems to illustrate less important or advanced topics. Every chapter (except for Chapter 9) is accompanied with a set of problems. Solutions for most of the problems can be found at the back of the text, starting on page 401.

Additional information about team projects and online links to related topics can be found at the book website: http://www.ece.rutgers.edu/~marsic/books/QoS/.

Contents at a Glance

Table of Contents

Chapter 1
Introduction to Computer Networks

Contents

1.1 Introduction

A network is a set of devices (often referred to as nodes) connected by communication links that are built using different physical media. A node can be a computer, telephone, or any other device capable of sending and receiving messages. The communication medium is the physical path by which message travels from sender to receiver. Example media include fiber-optic cable, copper wire, or air carrying radio waves.

1.1.1 The Networking Problem

Networking is about transmitting messages from senders to receivers (over a "communication channel"). Key issues we encounter include:

- "Noise" damages (corrupts) the messages; we would like to be able to *communicate reliably* in the presence of noise
- Establishing and maintaining physical communication lines is costly; we would like to be able to *connect arbitrary senders and receivers* while keeping the economic cost of network resources to a minimum
- Time is always an issue in information systems as is generally in life; we would like to be able to provide *expedited delivery* particularly for messages that have short deadlines

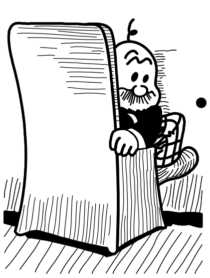

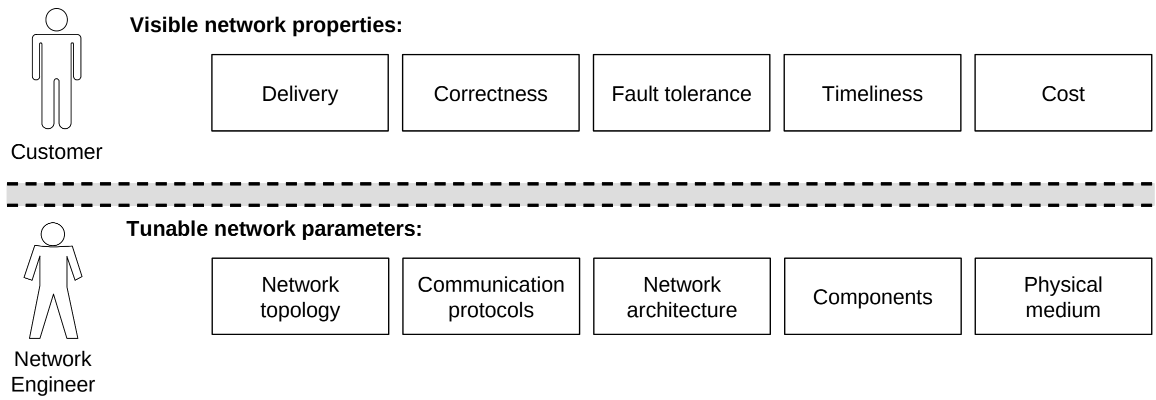

Figure 1-1: The customer cares about the visible network properties that can be controlled by the adjusting the network parameters.

Figure 1-1 illustrates what the customer usually cares about and what the network engineer can do about it. The visible network variables ("symptoms"), easily understood by a non-technical person include:

Delivery: The network must deliver data to the correct destination(s). Data must be received only by the intended recipients and not by others.

Correctness: Data must be delivered accurately, because distorted data is generally unusable.

Timeliness: Data must be delivered before they need to be put to use; else, they would be useless.

Fault tolerance and cost effectiveness are important characteristics of networks. For some of these parameters, the acceptable value is a matter of degree, judged subjectively. Our focus will be on network performance (objectively measurable characteristics) and quality of service (psychological determinants).

Limited resources can become overbooked, resulting in message loss. A network should be able to deliver messages even if some links experience outages.

The tunable parameters (or "knobs") for a network include: network topology, communication protocols, architecture, components, and the physical medium (connection lines) over which the signal is transmitted.

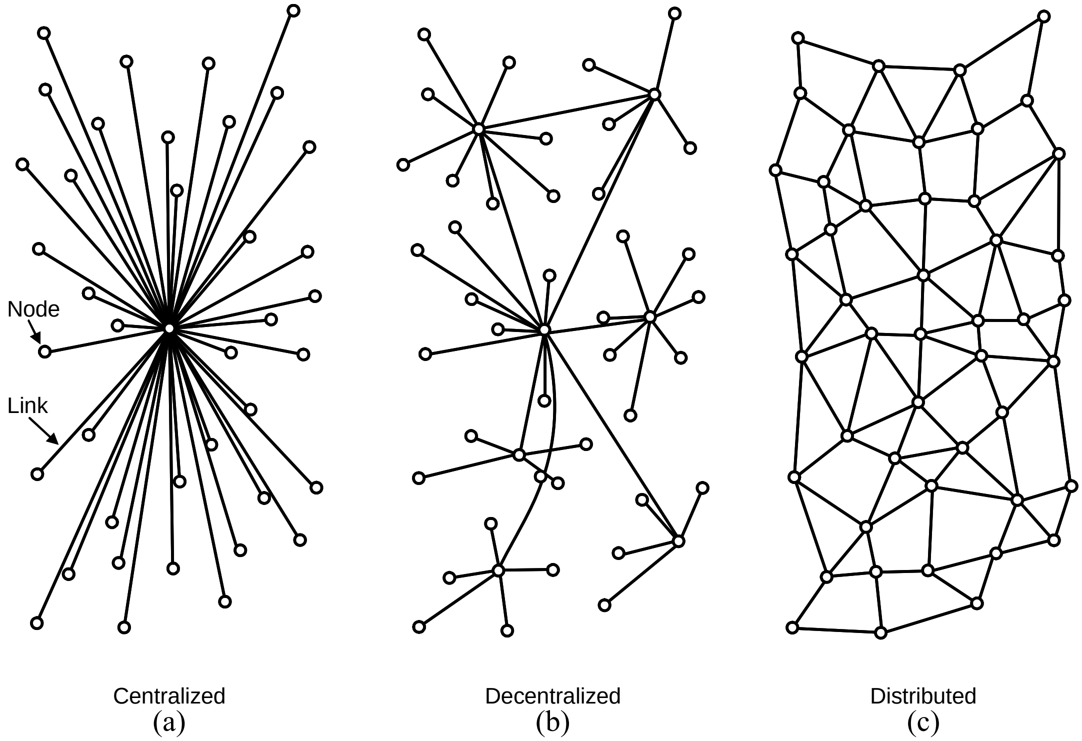

Figure 1-2: Different network topologies have different robustness characteristics relative to the failures of network elements.

- Connection topology: completely connected graph compared to link sharing with multiplexing and demultiplexing. Paul Baran considered in 1964 theoretically best architecture for survivability of data networks (Figure 1-2). He considered only network graph topology and assigned no qualities to its nodes and links[1]. He found that the distributed-topology network which resembles a fisherman's net, Figure 1-2(c), has the greatest resilience to element (node or link) failures. Figure 1-3 shows the actual topology of the entire Internet (in 1999). This topology evolved over several decades by incremental contributions from many independent organizations, without a "grand plan" to guide the overall design. In a sense, one could say that the Internet topology evolved in a "self-organizing" manner. Interestingly, it resembles more the decentralized-topology network with many hubs (Figure 1-2(b)), and to a lesser extent the distributed topology (Figure 1-2(c)).

[1] When discussing computer networks, the term "host" is usually reserved for communication endpoints and "node" is used for intermediary computing nodes that relay messages on their way to the destination.

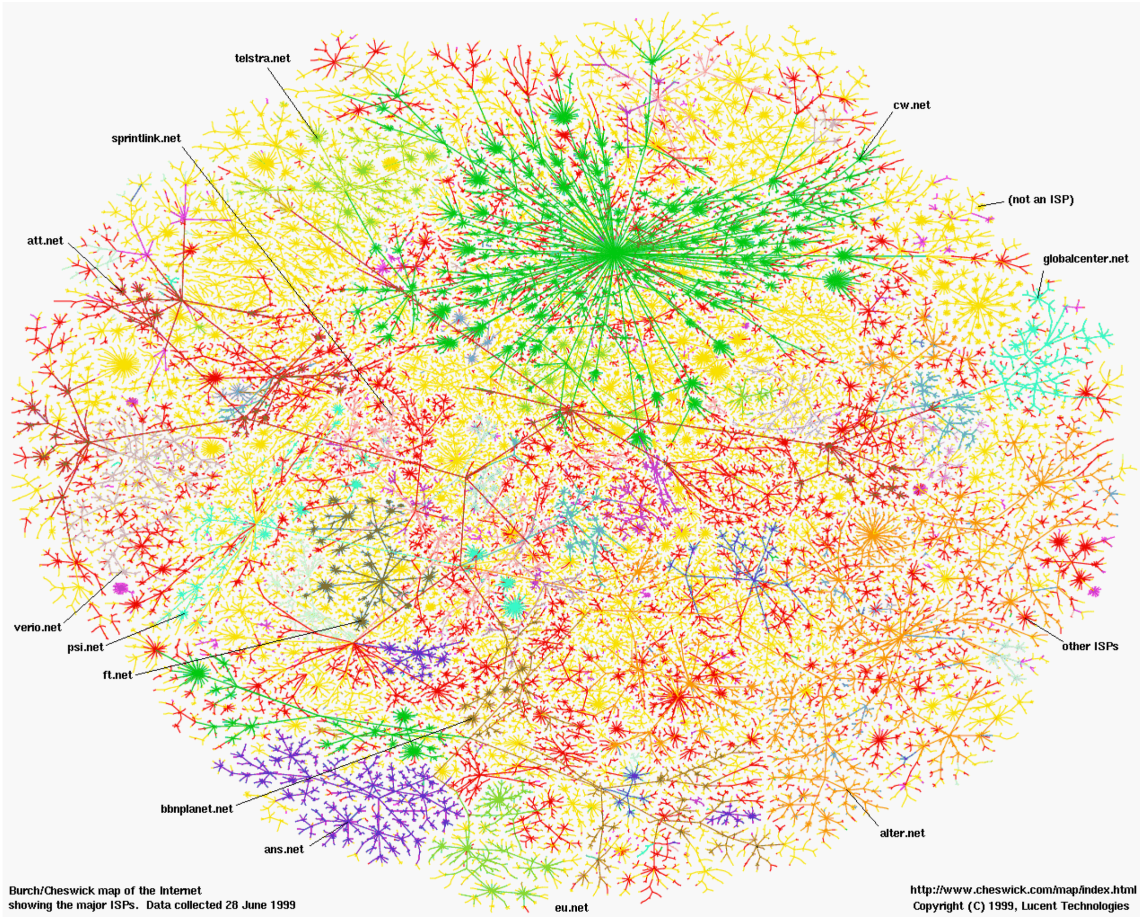

Figure 1-3. The map of the connections between the major Internet Service Providers (ISPs). [From the Internet mapping project: http://www.cheswick.com/]

- Network architecture: what part of the network is a fixed infrastructure as opposed to being ad hoc built for a temporary need

- Component characteristics: reliability and performance of individual hardware components (nodes and links). Faster and more reliable components are also more costly. When a network node (called switch or router) relays messages from a faster to a slower link, a congestion and a waiting-queue buildup may occur under a heavy traffic. In practice, all queues have limited capacity of their "waiting room," so loss occurs when messages arrive at a full queue.

- Performance metrics: success rate of transmitted packets (or, packet loss rate), average delay of packet delivery, and delay variability (also known as jitter)

- Different applications (data/voice/multimedia) have different requirements: sensitive to loss vs. sensitive to delay/jitter

There are some major problems faced by network engineers when building a large-scale network, such as the Internet that is now available worldwide. Some of these problems are non-technical:

- *Heterogeneity*: Diverse software and hardware of network components need to coexist and interoperate. The diversity results from different user needs and their economic capabilities, as

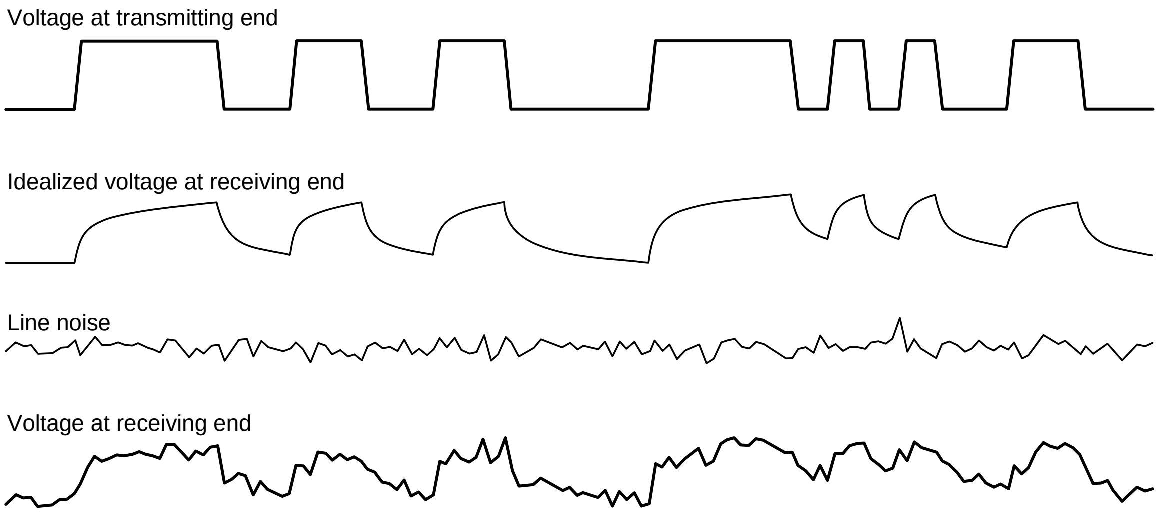

Figure 1-4: Digital signal distortion in transmission due to noise and time constants associated with physical lines.

well as because installed infrastructure tends to live long enough to become mixed with several new generations of technologies.

- *Autonomy*: Different parts of the Internet are controlled by independent organizations. Even a sub-network controlled by the same multinational organization, such as IBM or Coca Cola, may cross many state borders. These independent organizations are generally in competition with each other and do not necessarily provide one another the most accurate information about their own networks. The implication is that the network engineer can effectively control only a small part of the global network. As for the rest, the engineer will be able to receive only limited information about the characteristics of others' autonomous sub-networks. Any local solutions must be developed based on that limited information about the rest of the global network.

- *Scalability*: Although a global network like the Internet consists of many autonomous domains, there is a need for standards that prevent the network from becoming fragmented into many non-interoperable pieces ("islands"). Solutions are needed that will ensure smooth growth of the network as many new devices and autonomous domains are added. Again, information about available network resources is either impossible to obtain in real time, or may be proprietary to the domain operator.

1.1.2 Communication Links

There are many phenomena that affect the transmitted signal, some of which are illustrated in Figure 1-4. Although the effects of time constants and noise are exaggerated, they illustrate an important point. The input pulses must be well separated because too short pulses will be "smeared" together. This can be observed for the short-duration pulses at the right-hand side of the pulse train. Obviously, the receiver of the signal in the bottom row of Figure 1-4 will have great difficulty figuring out whether or not there were pulses in the transmitted signal. You can also see that longer pulses are better separated and easier to recognize in the distorted signal. The minimum tolerable separation depends on the physical characteristics of a transmission line (e.g., copper vs. optical fiber). If each pulse corresponds to a

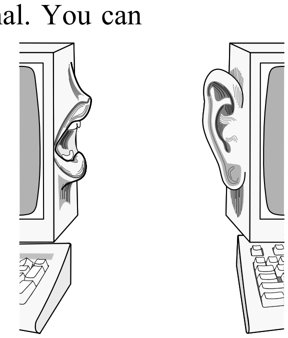

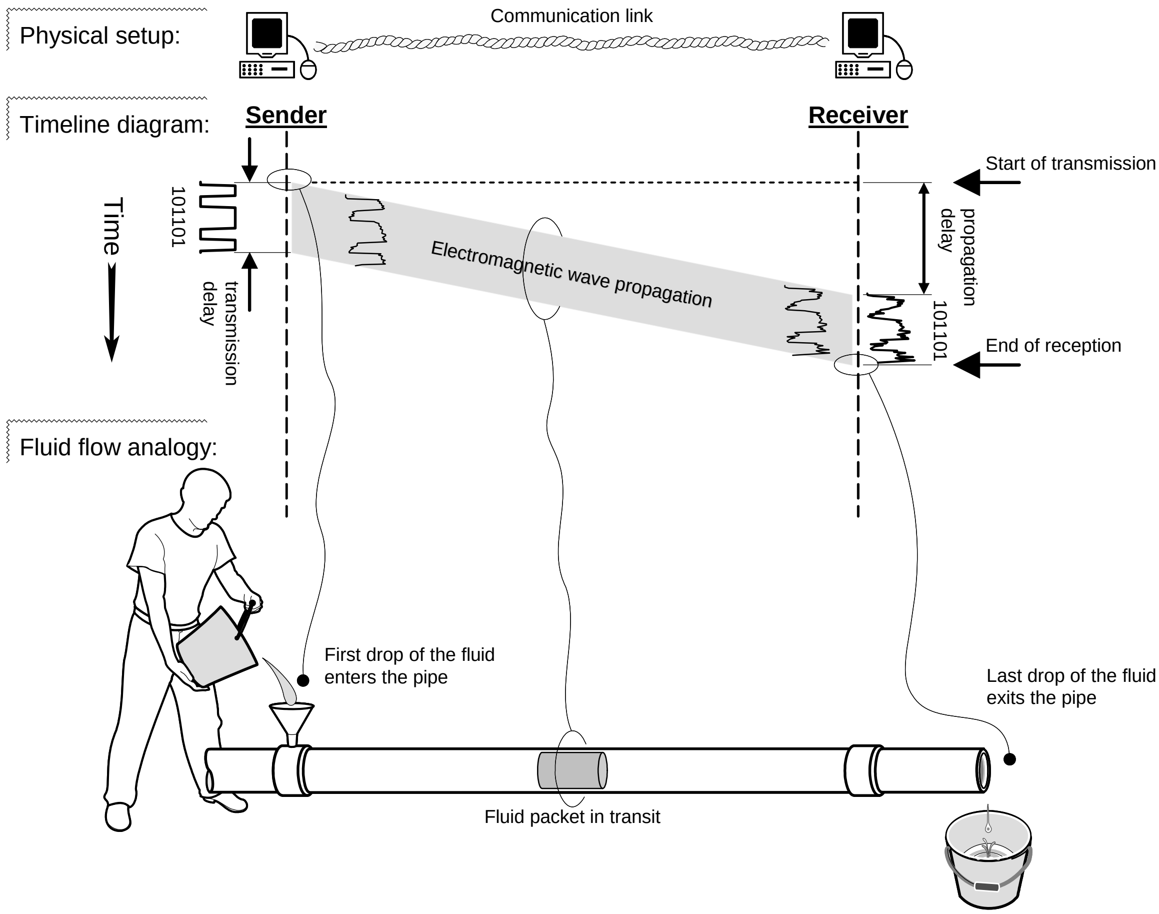

Figure 1-5: Timeline diagram for data transmission from sender to receiver.

single bit of information, then the minimum tolerable separation of pulses determines the maximum number of bits that can be transmitted over a particular transmission line.

It is common to represent data transmissions on a timeline diagram as shown in Figure 1-5. This figure also illustrates delays associated with data transmission. Although information bits are not necessarily transmitted as rectangular pulses of voltage, all transmission lines are conceptually equivalent, as represented in Figure 1-6, because the transmission capacity for every line is expressed in bits/sec or bps. The time required to transmit a single bit on a given link is known as **bit time**. In this text, we will always visualize transmitted data as a train of digital pulses. The reader interested in physical methods of signal transmission should consult a communications-engineering textbook, such as [Haykin, 2006].

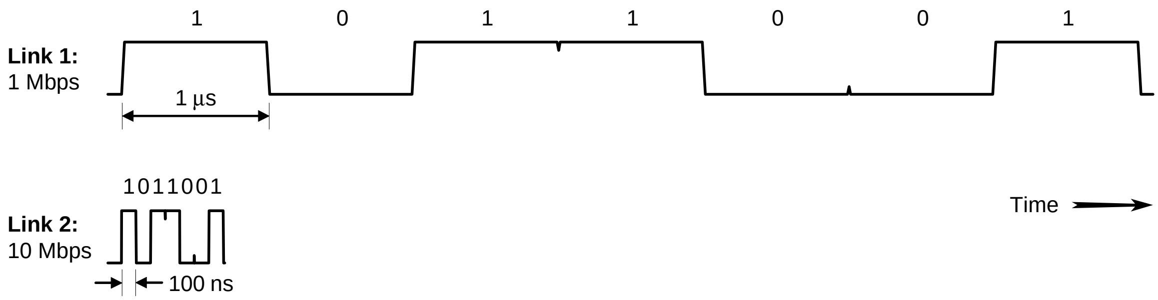

Figure 1-6: Transmission link capacity determines the speed at which the link can transmit data. In this example, each bit on Link 1 is 1 μs wide, while on Link 2 each bit is 100 ns wide. Hence, Link 2 can transmit ten times more data than Link 1 in the same time interval.

A common characterization of noise on transmission lines is **bit error rate** (BER): the fraction of bits received in error relative to the total number of bits received in transmission. Given a packet n bits long and assuming that bit errors occur independently of each other, a simple approximation for the packet error rate is

$$PER = 1 - (1 - BER)^n \approx 1 - e^{-n \cdot BER} \tag{1.1}$$

An important attribute of a communication link is how many bitstreams can be transmitted on it at the same time. If a link allows transmitting only a single bitstream at a time, then the nodes connected to the link must coordinate their transmissions to avoid different bitstreams corrupting each other (known as data collision). Such links are known as *broadcast links* or *multiple-access links*. Point-to-point links often support data transmissions in both directions simultaneously. This kind of a link is said to be **full duplex**. A point-to-point link that supports data flowing in only one direction at a time is called **half duplex**. In other words, the nodes on each end of this kind of a link can both transmit and receive, but not at the same time—they only can do it by taking turns. It is like a one-lane road with bidirectional traffic. We will assume that all point-to-point links are full duplex, unless stated otherwise. A full-duplex link can be implemented in two ways: either the link must contain two physically separate transmission paths, one for sending and one for receiving, or the capacity of the communication channel is divided between signals traveling in opposite directions. The latter is usually achieved by *time division multiplexing* (TDM) or *frequency division multiplexing* (FDM).

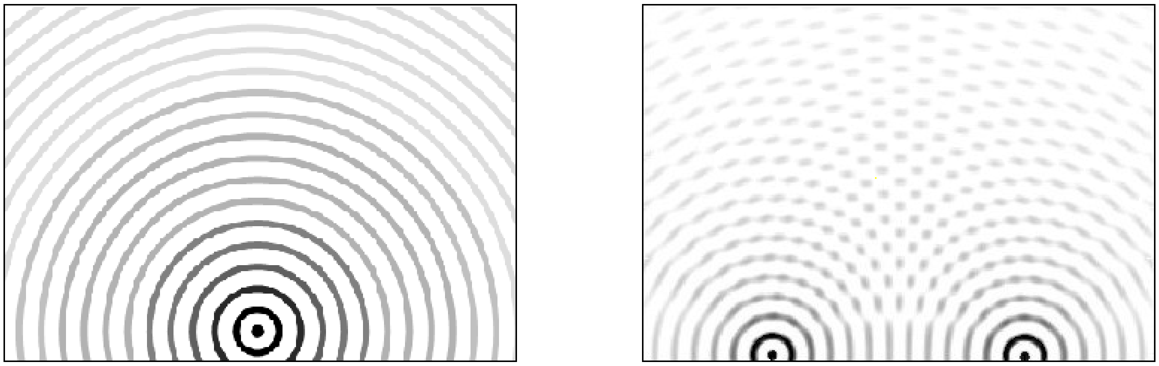

Figure 1-7: Wireless transmission. Left: single point source. Right: interference of two point sources.

Wireless Link

Consider a simple case of a point source radiating electromagnetic waves in all directions (Figure 1-7, left). The received signal strength decreases exponentially with the sender-receiver distance. As with any other communication medium, the wireless channel is subject to thermal noise, which distorts the signal randomly according to a Gaussian distribution of noise amplitudes. As the distance between a transmitter and receiver increases, the received signal strength decreases to levels close to the background noise floor. At a certain distance from the sender, the signal strengths will become so weak that the receiver will not be able to discern reliably signal from noise. This distance, known as **transmission range**, is decided arbitrarily, depending on what is considered acceptable bit error rate. For example, we can define the transmission range as the sender-receiver distance for which the packet error rate is less than 10 %.

In addition to thermal noise, the received signal may be distorted by parallel transmissions from other sources (Figure 1-7, right). This phenomenon is known as **interference**. Because this normally happens only when both sources are trying to transmit data (unknowingly of each other's parallel transmissions), this scenario is called **packet collision**. A key observation is that *collisions occur at the receiver*—the sender is not disturbed by concurrent transmissions, but receiver cannot correctly decode sender's message if it is combined with an interfering signal. If the source and receiver nodes are far away from the interfering source, the interference effect at the receiver will be a slight increase in the error rate. If the increased error rate is negligible, the source and receiver will be able to carry out their communication despite the interference. Notice, however, that the interference of simultaneously transmitting sources never disappears—it only is reduced exponentially with an increasing mutual distance (Figure 1-8). The minimum distance (relative to the receiver) at which interferer's effect can be considered negligible is called **interference range**. In Figure 1-8, node *D* is within the interference range of receiver *B*. Nodes *C* and *E* are outside the interference range. However, although outside the interference range defined for a single interferer, if nodes *C* and *E* are transmitting simultaneously their combined interference at *B* may be sufficiently high to cause as great or greater number of errors as a single interferer within the interference range.

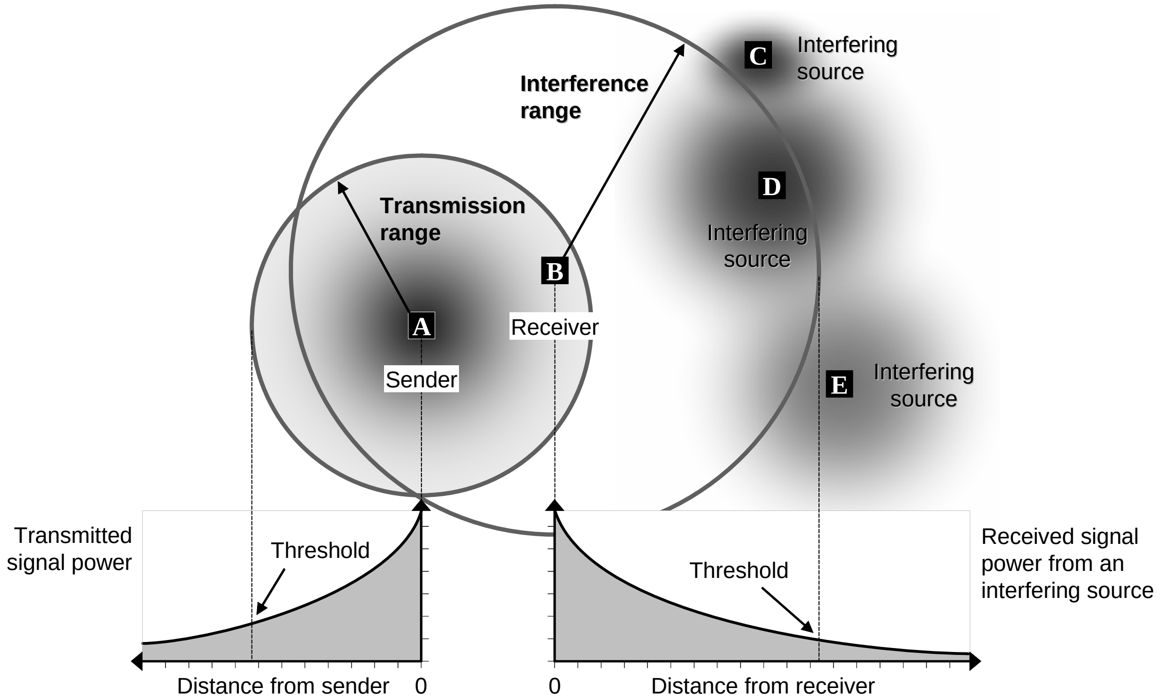

Figure 1-8: Transmission range and interference range for wireless links.

1.1.3 Packets and Statistical Multiplexing

The communication channel essentially provides an abstraction of a continuous stream of symbols transmitted that are subject to a certain error probability. When interacting with another person, whether face-to-face or over the telephone, we think of units of communication in terms of conversational turns: first one person takes a turn and delivers their message, then the other person takes a turn, and so on. *Messages* could be thought of as units of communication exchanged by two (or more) interacting persons. We notice that there are benefits of slicing a long oration into a sequence of smaller units of discourse. This slicing into messages gives the other person chance to clarify a misunderstanding or give a targeted response to a specific item.

In computer communication networks, messages are represented as strings of binary symbols (0 or 1), known as bits. Generally, messages are of variable length and some of them may still be considered too long for practical network transmission. There are several reasons for imposing a limit on message length. One is that longer messages stand a higher chance of being corrupted by an error (see equation (1.1)). Another reason is to avoid the situation where a sending application seizes the link for itself by sending very long messages while other applications must wait for a long time. Therefore, messages are broken into shorter bit strings known as **packets**. These packets are then transmitted independently and reassembled into messages at the destination. This allows individual packets to opportunistically take *alternate routes* to the destination and *interleave* the network usage by multiple sources, thus avoiding inordinate waiting periods for some sources to transmit their information.

Different network technologies impose different limits on the size of data blocks they can handle, which is known as the **maximum transmission unit** (MTU). For example, a regular Ethernet

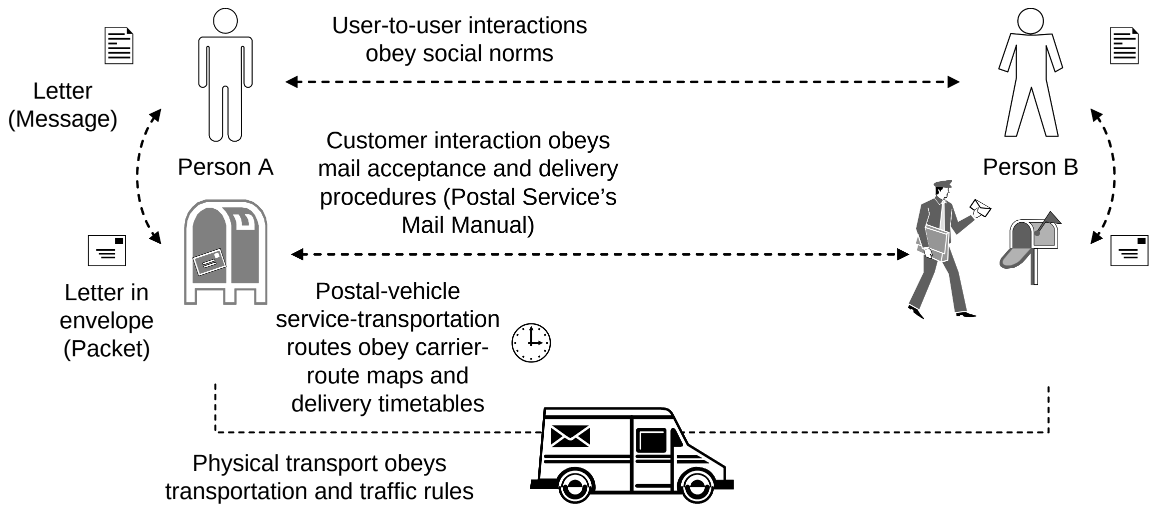

Figure 1-9: Protocol layers for conventional mail transport.

frame uses a frame format that limits the size of the payload it sends to 1,500 bytes. Notice that the MTU value specifies the maximum payload size and does not include the header size of the header that is prepended to the payload of a packet.

Statistical Multiplexing

Link sharing using packet multiplexers

Real-world systems are designed with sub-peak capacity for economic reasons. As a result, they experience congestion and delays during peak usage periods. Highways experience traffic slowdown during rush hours; restaurants or theaters experience crowds during weekend evenings; etc. Designing these systems to support peak usage without delays would not be economically feasible—most of the time they would be underutilized. Figure 1-10

1.1.4 Communication Protocols

A **protocol** is a *set of rules* agreed-upon by interacting entities, e.g., computing devices, that govern their communicative exchanges. It is hard to imagine accomplishing any task involving multiple entities without relying on a protocol. For example, one could codify the rules for how a customer (C) purchases goods from a merchant (M) as follows:

1. C→M Request catalog of products

2. C←M Respond catalog

3. C→M Make selections

4. C←M Deliver selections

5. C→M Confirm delivery

6. C←M Issue bill

7. C→M Make payment

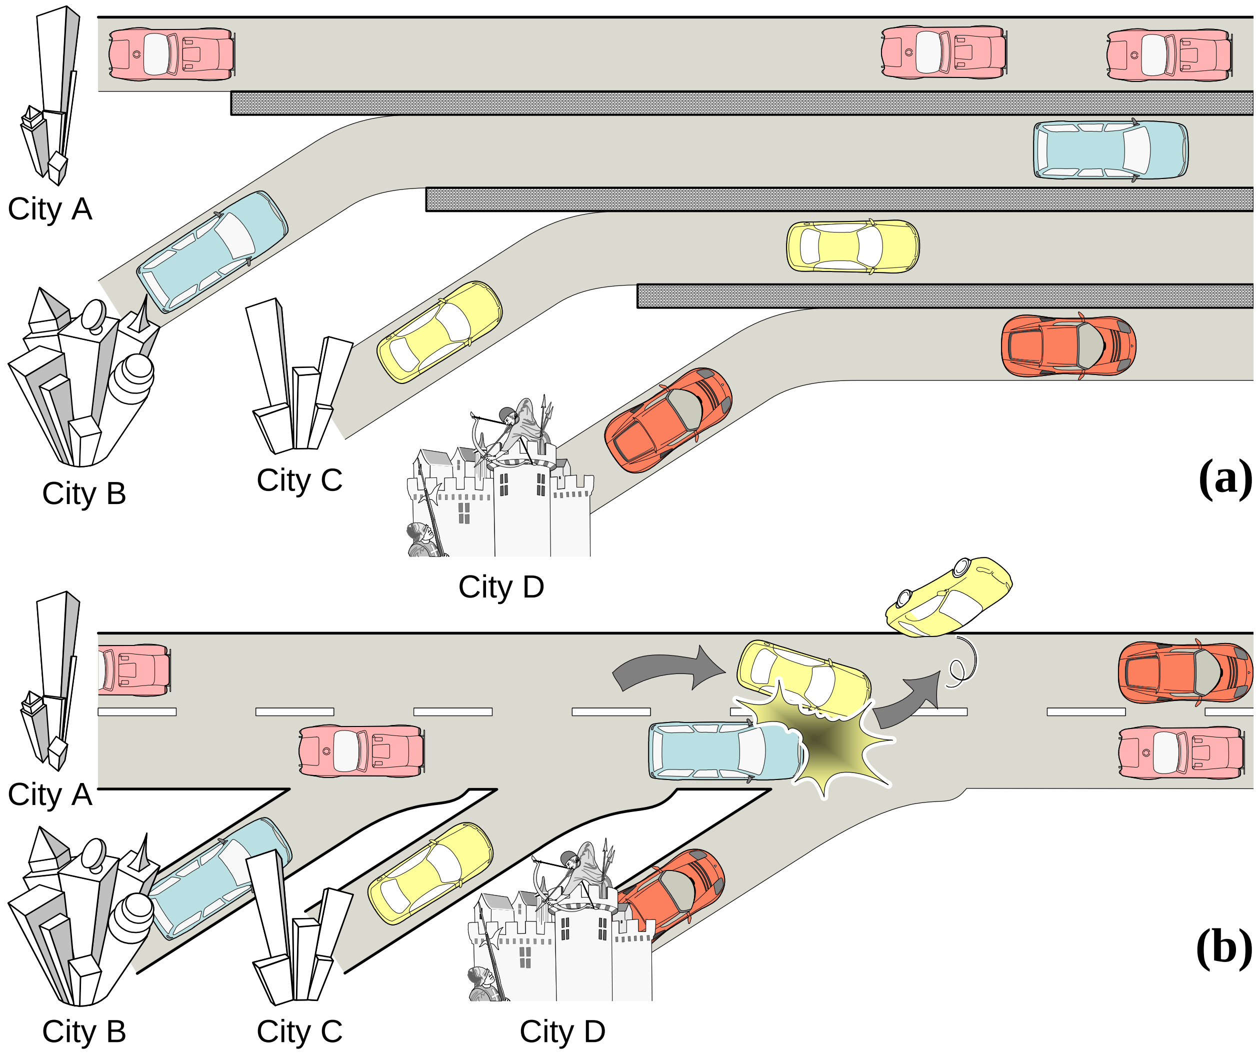

Figure 1-10: An analogy illustrating dedicated lines (a) compared to statistical maultiplexing (b).

8. C←M Issue confirmation

The customer and merchant may be located remote from each other and using other entities to help accomplish the purchasing task, such as a bank for credit-card transactions, or a postal service for parcel delivery.

An important characteristic of protocols is that the units of communication are **data packets**. Each data packet consists of a **header** that contains the packet guidance information to help guide the packet from its source to its destination, and the **payload**, which is the user information to be delivered to the destination address. In packet-switched networks, packets are transmitted at random times, and the receiver at the end of a communication link must have a means to distinguish an arriving packet from random noise on the line. For this purpose, each packet is preceded with a special sequence of bits that mark the start of the packet. This special bit pattern is usually called the **preamble**. Each receiver is continuously hunting for the preamble to catch the arriving packet. If the preamble is corrupted by random noise, the packet will be lost (i.e., unnoticed by the receiver).

Communication in computer networks is very complex. One effective means of dealing with complexity is known as *modular design* with *separation of concerns*. In this approach, the system is split into modules and each module is assigned separate tasks to do ("concerns"). Network designers usually adopt a restricted version of modular design, know as **layered design**. Each **layer** defines a collection of conceptually similar functions (or, services) distinct from those of the other layers. The restriction in layered design is that a module in a given layer provides services to the layer just above it and receives services from the layer just below it. The layering approach forbids the modules from using services from (or providing to) non-adjacent layers.

Each layer of the layered architecture contains one or more software *modules* that offer services characteristic for this layer. Each module is called **protocol**. A protocol defines two application-programming interfaces (APIs):

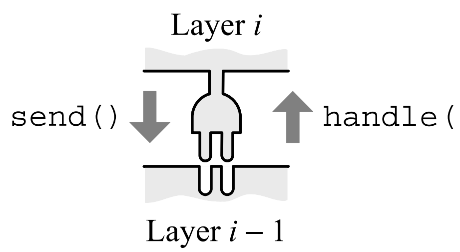

1. *Service interface* to the protocols in the layer above this layer. The upper-layer protocols use this interface to "plug into" this layer and hand it data packets to `send()`. Each layer also defines a `handle()` callback operation through which the lower layer calls this layer to handle an incoming packet.
2. *Peer interface* to the counterpart protocol on a remote machine. This interface defines the format and meaning of data packets exchanged between peer protocols to support communication.

There are many advantages of layered design, primarily because it decomposes the problem of building a network into more manageable components. Each component can be developed independently and used interchangeably with any other component that complies with its service interface. However, there are some disadvantages, as well. For example, when a layer needs to make a decision about how to handle a data packet, it would be helpful to know what kind of information is inside the packet. Because of strict separation of concerns, particularly between the non-adjacent layers, this information is not available, so a more intelligent decision cannot be made. This is the reason why recently **cross-layered designs** are being adopted, particularly for wireless networks (see Chapter 6).

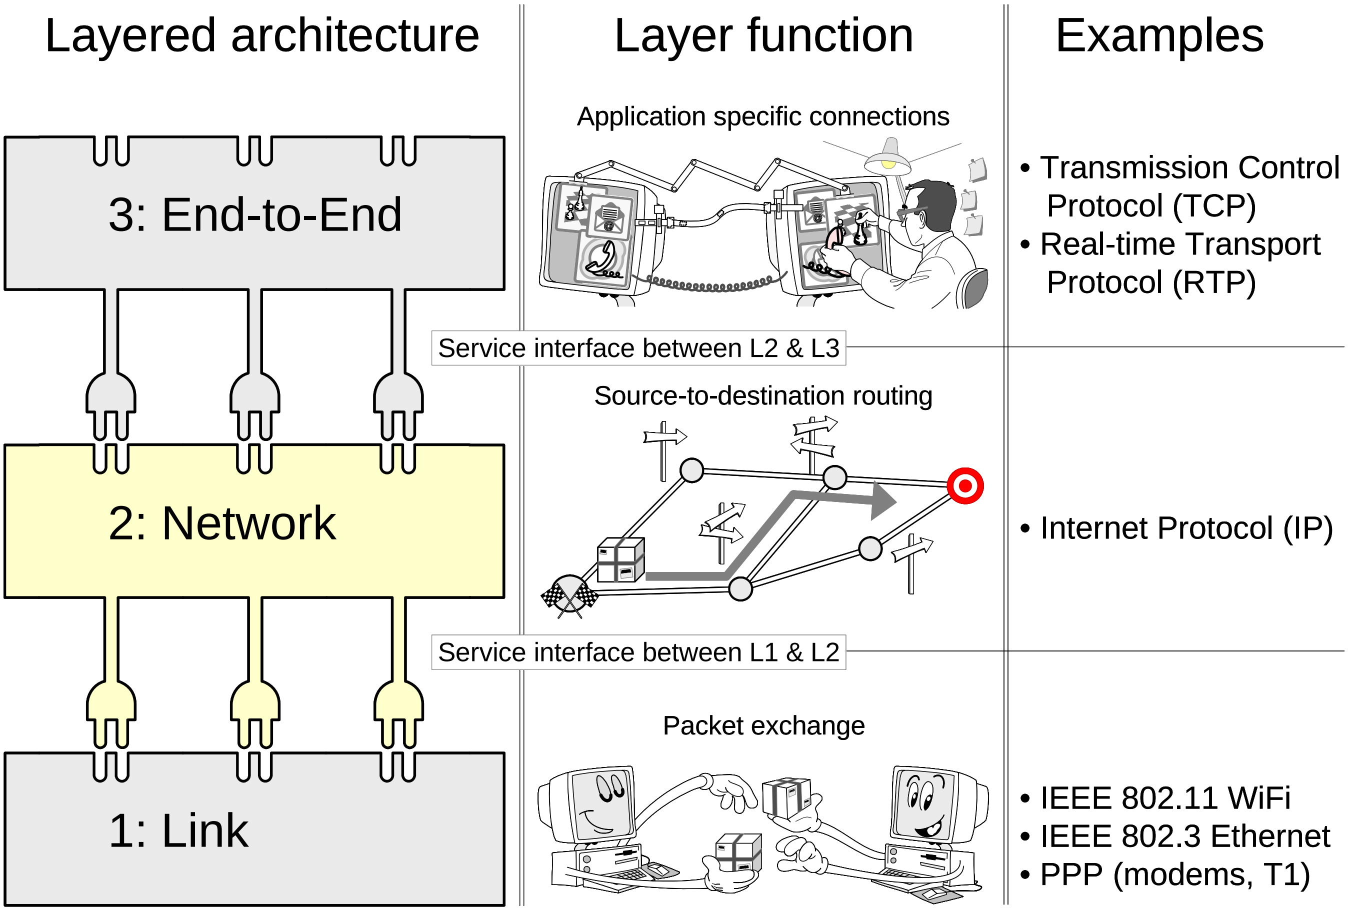

Figure 1-11: Three-layer reference architecture for communication protocol layering.

Three-Layer Model

In recent years, the three-layer model (Figure 1-11) has emerged as reference architecture of computer networking protocols.

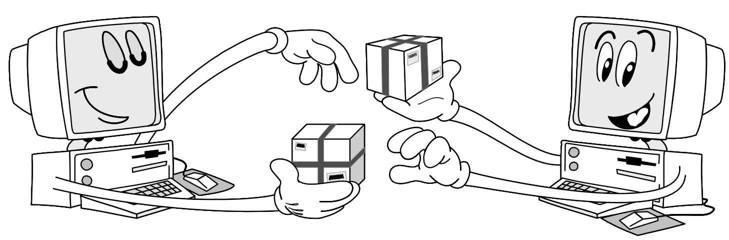

LAYER-1 – Link layer: is at the bottom of the protocol stack and implements a packet delivery service between nodes that are attached to the same physical link (or, physical medium). The physical link may be point-to-point from one transmitter to a receiver, or it may be shared by a number of transmitters and receivers (known as "broadcast link," Section 1.3.3).

There is the "physical layer," which implements a digital transmission system that delivers bits. But, you would not know it because it is usually tightly coupled with the link layer by the link technology standard. Link and physical layers are usually standardized together and technology vendors package them together, as will be seen later in Section 1.5.

In wireless networks, physical communication is much more complex than in wired networks. Therefore, it may be justifiable to distinguish the physical and link layers, to keep both manageable. Because this book is mainly about protocols and not about physical communication, I will consider both together as a single, Link layer.

The link layer is not concerned with bridging end hosts across (many) intermediate links; this is why we need the network layer.

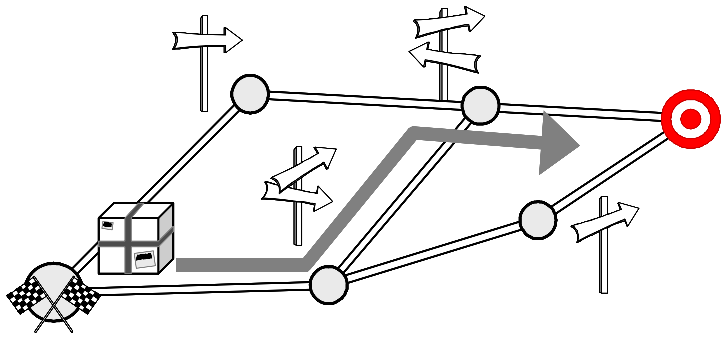

LAYER-2 – Network layer: provides concatenation of links to connect arbitrary end hosts. It will be elaborated in Section 1.4 where we describe the most popular network layer protocol: the Internet Protocol (IP). However, host computers are not the endpoints of communication—application programs running on these hosts are the actual endpoints of communication! The network layer is not concerned with application requirements. It may provide a range of choices for an upper layer to select from. For example, the network layer may support "quality of service" through "service level agreements," "resource reservation," but it does not know which one of these choices is the best for a particular application program; this is why we need the end-to-end layer.

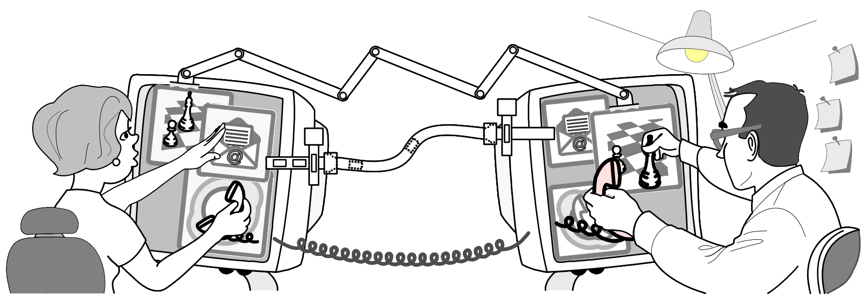

LAYER-3 – End-to-end layer: this layer brings together (encapsulates) all communication-specific features of an application program. Here is the first time that we are concerned with application requirements.

The figure on the right is meant to illustrate that different applications need different type of connection for optimal performance. For example, manipulation-based applications (such as video games) require an equivalent of mechanical links to convey user's action. Telephony applications need an equivalent of a telephone wire to carry user's voice, etc. A most prominent example of an end-to-end protocol is TCP, described in Chapter 2.

A fundamental design principle of network protocols and distributed systems in general is the **end-to-end principle**. The principle states that, whenever possible, communications protocol operations should occur at the end-points of a communications system, or as close as possible to the resource being controlled. According to the end-to-end principle, protocol features are only justified in the lower layers of a system if they are a performance optimization.

Figure 1-12 shows the layers involved when a message is sent from an application running on one host to another running on a different host. The application on host *A* hands the message to the end-to-end layer, which passes it down the protocol stack on the same host machine. Every layer accepts the payload handed to it and processes it to add its characteristic information in the form of an additional header (Figure 1-13). The link layer transmits the message over the physical

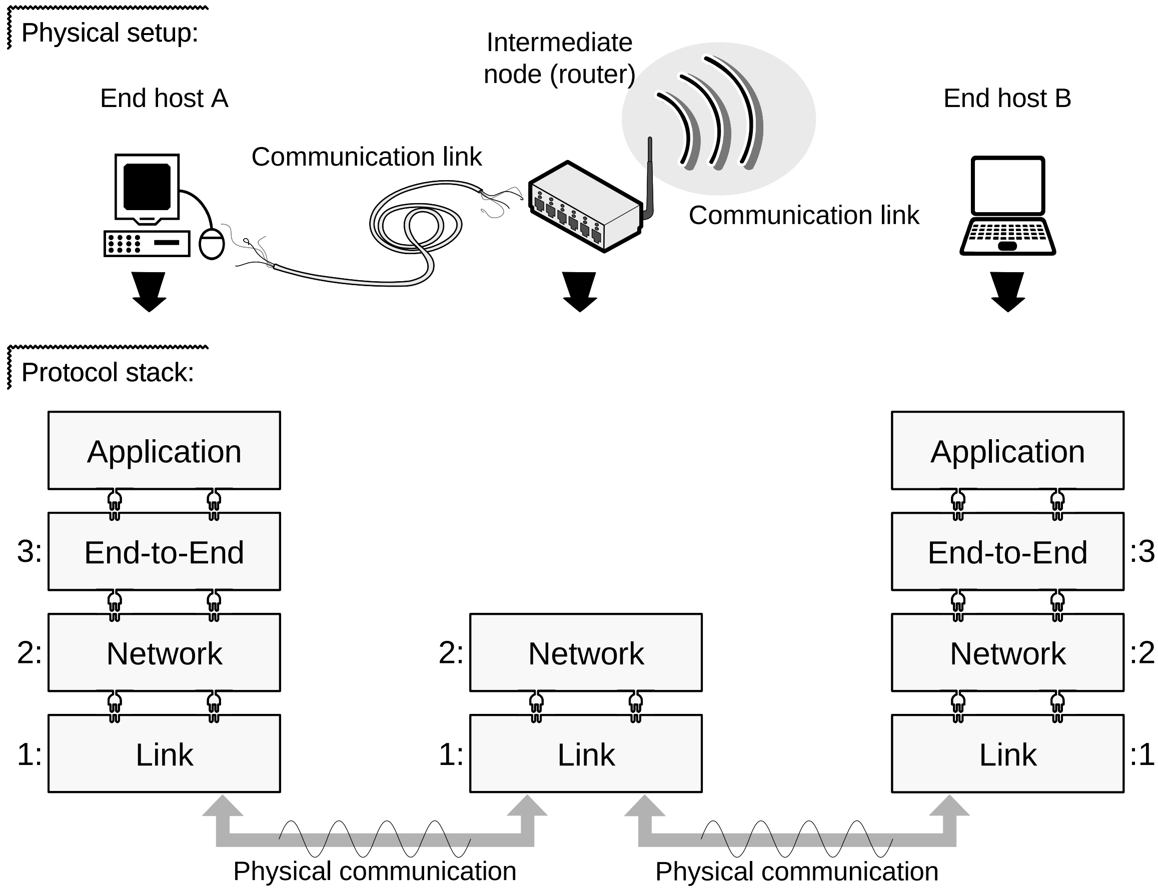

Figure 1-12: Protocol layering in end hosts and intermediate nodes (switches or routers).

medium. As the message travels from *A* to *B*, it may pass through many intermediate nodes, known as *switches* or *routers*. In every receiving node (including the intermediate ones), the message is received by the bottommost (or, link) layer and passed up through their protocol stack. Because intermediate nodes are not the final destination (or, end point), they do not have the complete protocol stack, but rather only the two bottommost layers: link and network layers (see Figure 1-12).

Pseudo code of a generic protocol module in layer *i* is given in Listing 1-1.

Listing 1-1: Pseudo code of a protocol module in layer *i*.

```
// Definition of packet format for layer i.
// Implementing the java.io.Externalizable interface makes possible to serialize
// a Packet object to a byte stream (which becomes the payload for the lower-layer protocol).
public class PacketLayer_i implements java.io.Externalizable {
    // packet header
    private String sourceAddress;
    private String receiverAddress;
    private String packetID;   // this packet's identifier
    private String receivingProtocol; // upper layer protocol at receiver

    // packet payload
```

```
    private byte[] payload;

    // constructor
    public PacketLayer_i(
        byte[] data, String recvAddr, String upperProtocol
    ) {
        payload = data;
        sourceAddress = address of my host computer;
        receiverAddress = recvAddr;
        packetID = generate unique identifier for this packet;
        receivingProtocol = upperProtocol;
    }

    public void writeExternal(ObjectOutput out) {
        // Packet implements java.io.Externalizable instead of java.io.Serializable
        // to be able to control how the serialization stream is written, because
        // it must follow the standard packet format for the given protocol.
    }
    public void readExternal(ObjectOutput out) {
        // reconstruct a Packet from the received bytestream
    }
}
```

```
// Definition of a generic protocol module in layer i.

public class ProtocolLayer_i {
    // maximum number of outstanding packets at sender (zero, if NOT a persistent sender)
    public static final int N; // (N is also called the sliding window size

    // lower layer protocol that provides services to this protocol
    private ProtocolLayer_iDOWN lowerLayerProtocol;

    // look-up table of upper layer protocols that use services of this protocol
    private HashMap upperLayerProtocols;

    // look-up table of next-receiver node addresses based on final destination addresses
    //      (this object is shared with the routing protocol, shown in Listing 1-2)
    private HashMap forwardingTable;

    // list of unacknowledged packets that may need to be retransmitted
    //     (maintained only for persistent senders that provide reliable transmission)
    private ArrayList unacknowledgedPackets = new ArrayList();

    // constructor
    public ProtocolLayer_i(
        ProtocolLayer_iDOWN lowerLayerProtocol
    ) {
        this.lowerLayerProtocol = lowerLayerProtocol;
    }

    // sending service offered to the upper layer protocols, called in a top-layer thread
    public void send(
        byte[] data, String destinationAddr,
        ProtocolLayer_iUP upperProtocol
    ) throws Exception {
        // if persistent sender and window of unacknowledged packets full, then do nothing
```

```
        if ((N > 0) && (N - unacknowledgedPackets.size() <= 0)) ) {
           throw exception: admission refused by overbooked sender;
        }

        // create the packet to send
        PacketLayer_i outgoingPacket =
           new PacketLayer_i(data, destinationAddr, upperProtocol);

        // serialize the packet object into a byte-stream (payload for lower-layer protocol)
        java.io.ByteArrayOutputStream bout =
            new ByteArrayOutputStream();
        java.io.ObjectOutputStream outstr =
            new ObjectOutputStream(bout);
        outstr.writeObject(outgoingPacket);

        // look-up the receiving node of this packet based on the destination address
        //       (requires synchronized access because the forwarding table is shared
        //       with the routing protocol)
        synchronized (forwardingTable) { // critical region
            String recvAddr = forwardingTable.get(destinationAddr);
        } // end of the critical region

        // hand the packet as a byte-array down the protocol stack for transmission
        lowerLayerProtocol.send(
             bout.toByteArray(), recvAddr, this
        );
    }

    // upcall method, called from the layer below this one, when data arrives
    //    from a remote peer (executes in a bottom-layer thread!)
    public void handle(byte[] data) {

        // reconstruct a Packet object from the received data byte-stream
        ObjectInputStream instr = new ObjectInputStream(
            new ByteArrayInputStream(data)
        );
        PacketLayer_i receivedFrame =
            (PacketLayer_i) instr.readObject();

        // if this packet is addressed to me ... (on a broadcast medium)
        if (receivedFrame.getReceiverAddress() == my address) {
            // ...determine which upper layer protocol should handle this packet's payload
            synchronized (upperLayerProtocols) { // critical region
              ProtocolLayer_iUP upperProtocol = (ProtocolLayer_iUP)
                  upperLayerProtocols.get(
                      receivedFrame.getReceivingProtocol()
                  );
            } // end of the critical region

            // remove this protocol's header and
            //      hand the payload over to the upper layer protocol
            upperProtocol.handle(receivedFrame.getPayload());
        }
    }
```

```
51      public void setHandler(
51a         String receivingProtocol, ProtocolLayer_iUP upperProtocol
51b     ) {
52          // add a <key, value> entry into the routing look-up table
53          upperLayerProtocols.put(receivingProtocol, upperProtocol);
54      }

55      // Method called by the routing protocol (running in a different thread or process)
56      public void setReceiver(
56a         String destinationAddr, String receiverAddr
56b     ) {
57          // add a <key, value> entry into the forwarding look-up table
58          synchronized (forwardingTable) { // critical region
59              forwardingTable.put(destinationAddr, receiverAddr);
60          } // end of the critical region
61      }
62 }
```

Here I provide only a brief description of the above pseudocode. We will encounter and explain the details later in this chapter, as new concepts are introduced. The attribute `upperProtocol` is used to decide to which upper-layer protocol to deliver a received packet's payload. This process is called *protocol demultiplexing* and allows the protocol at a lower-layer layer to serve different upper-layer protocols.

To keep the pseudo code manageable, some functionality is not shown in Listing 1-1. For example, in case of a persistent sender in the method `send()` we should set the retransmission timer for the sent packet and store the packet in the `unacknowledgedPackets` list. Similarly, in the method `handle()` we should check what packet is acknowledged and remove the acknowledged packet(s) from the `unacknowledgedPackets` list. Also, the method `send()` is shown to check only the `forwardingTable` to determine the intermediary receiver (`recvAddr`) based on the final destination address. In addition, we will see that different protocol layers use different addresses for the same network node (Section 8.3.1). For this reason, it is necessary to perform address translation from the current-layer address (`recvAddr`) to the address of the lower layer before passing it as an argument in the `send()` call, in Line 33. (See Section 8.3 for more about address translation.)

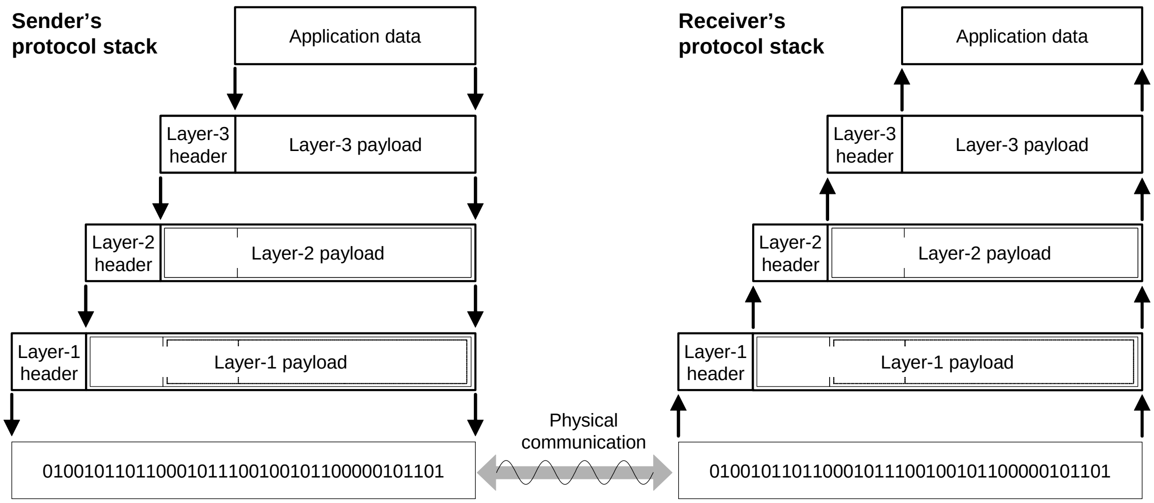

Figure 1-13: Packet nesting across the protocol stack: an entire packet of an upper layer becomes the data payload of a lower layer.

The reader who carefully examined Listing 1-1 will have noticed that packets from higher layers become nested inside the packets of lower layers as they are passed down the protocol stack (Figure 1-13). The protocol at a lower layer is *not* aware of any structure in the data passed down from the upper layer, i.e., it does not know if the data can be partitioned to header and payload or where their boundary is—it simply considers the whole thing as an unstructured data payload.

The generic protocol implementation in Listing 1-1 works for all protocols in the layer stack. However, each layer will require some layer-specific modifications. The protocol in Listing 1-1 best represents a Network layer protocol for source-to-destination packet delivery.

The Link layer is special because it is at the bottom of the protocol stack and cannot use services of any other layer. It runs the `receiveBits()` method in a continuous loop (perhaps in a separate thread or process) to hunt for arriving packets. This method in turn calls Link layer's `handle()`, which in turn calls the upper-layer (Network) method `handle()`. Link layer's `send()` method, instead of using a lower-layer service, itself does the sending by calling this layer's own method `sendBits()`.

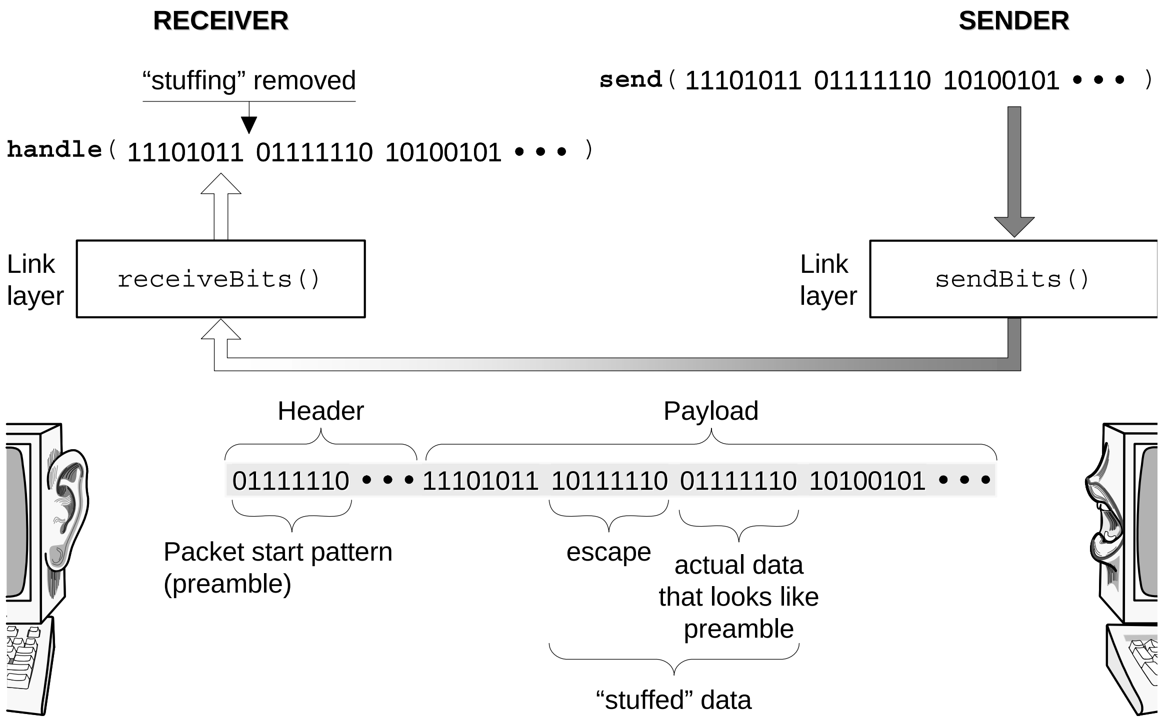

Figure 1-14: Bit stuffing to escape special control patterns in the frame data.

An important feature of a link-layer protocol is **data transparency**, which means that it must carry any bit pattern in the data payload. An example of a special bit pattern is the packet preamble that helps the receiver to recognize the start of an arriving packet (mentioned at the start of this section). Data transparency means that the link layer must not forbid the upper-layer protocol from sending data containing special bit patterns. To implement data transparency, link-layer protocol uses a technique known as **bit stuffing** (or, **byte stuffing**, depending on what the smallest units used to measure the payload is). Bit stuffing defines a special control escape bit pattern, call it *ESC* (Figure 1-14). The method `sendBits()` examines the payload received from the upper-layer protocol. If it encounters a special control sequence, say preamble (call it *PRE*), then it "stuffs" (adds) a control escape sequence *ESC* into the transmitted data stream, before *PRE* (resulting in *ESC PRE*), to indicate that the following *PRE* is *not* a preamble but is, in fact, actual data. Similarly, if the control escape pattern *ESC* itself appears as actual data, it too must be preceeded by an *ESC*. The method `receiveBits()` removes any control escape patterns that it finds in the received packet before delivering it to the upper-layer method `handle()` (Figure 1-14).

The pseudo code in Listing 1-1 is only meant to illustrate how one would write a protocol module. It is extremely simplified and certainly not optimized for performance. My main goal is to give the reader an idea about the issues involved in protocol design. We will customize the pseudo code from Listing 1-1 for different protocols, such as routing protocols in Listing 1-2, Section 1.4, and TCP sender in Listing 2-1, Section 2.1.1.

When Is a "Little in the Middle" OK? The Internet's End-to-End Principle Faces More Debate; by Gregory Goth -- http://ieeexplore.ieee.org/iel5/8968/28687/01285878.pdf?isnumber

Why it's time to let the OSI model die; by Steve Taylor and Jim Metzler, Network World, 09/23/2008 -- http://www.networkworld.com/newsletters/frame/2008/092208wan1.html

Open Systems Interconnection (OSI) Reference Model

The OSI model has seven layers (Figure 1-15). The layer functionality is as follows:

Layer 7 – Application: Its function is to provide application-specific services. Examples include call establishment and management for a telephony application (SIP protocol, Section 8.6.2), mail services for e-mail forwarding and storage (SMTP protocol), and directory services for looking up global information about various network objects and services (LDAP protocol). Notice that this layer is distinct from the application itself, which provides business logic and user interface.

Layer 6 – Presentation: Its function is to "dress" the messages in a "standard" manner. It is sometimes called the *syntax layer* because it deals with the syntax and semantics of the information exchanged between the network nodes. This layer performs *translation* of data representations and formats to support interoperability between different encoding systems (ASCII vs. Unicode) or hardware architectures. It also performs *encryption* and *decryption* of sensitive information. Lastly, this layer also performs data *compression* to reduce the number of bits to be transmitted, which is particularly important for multimedia data (audio and video).

Layer 5 – Session: Its function is to maintain a "conversation" across multiple related exchanges between two hosts (called *session*), to keep track of the progress of their communication. This layer establishes, manages, and terminates sessions. Example services include keeping track of whose turn it is to transmit (dialog control) and checkpointing long conversations to allow them to resume after a crash.

Layer 4 – Transport: Its function is to provide reliable or expedient delivery of messages, or error recovery.

Layer 3 – Network: Its function is to move packets from source to destination in an efficient manner (called *routing*), and to provide internetworking of different network types (a key service is address resolution across different networks or network layers).

Layer 2 – Link: Its function is to organize bits into packets or frames, and to provide packet exchange between adjacent nodes.

Layer 1 – Physical: Its function is to transmit bits over a physical medium, such as copper wire or air, and to provide mechanical and electrical specifications.

Visit http://en.wikipedia.org/wiki/OSI_model for more details on the OSI Reference Architecture

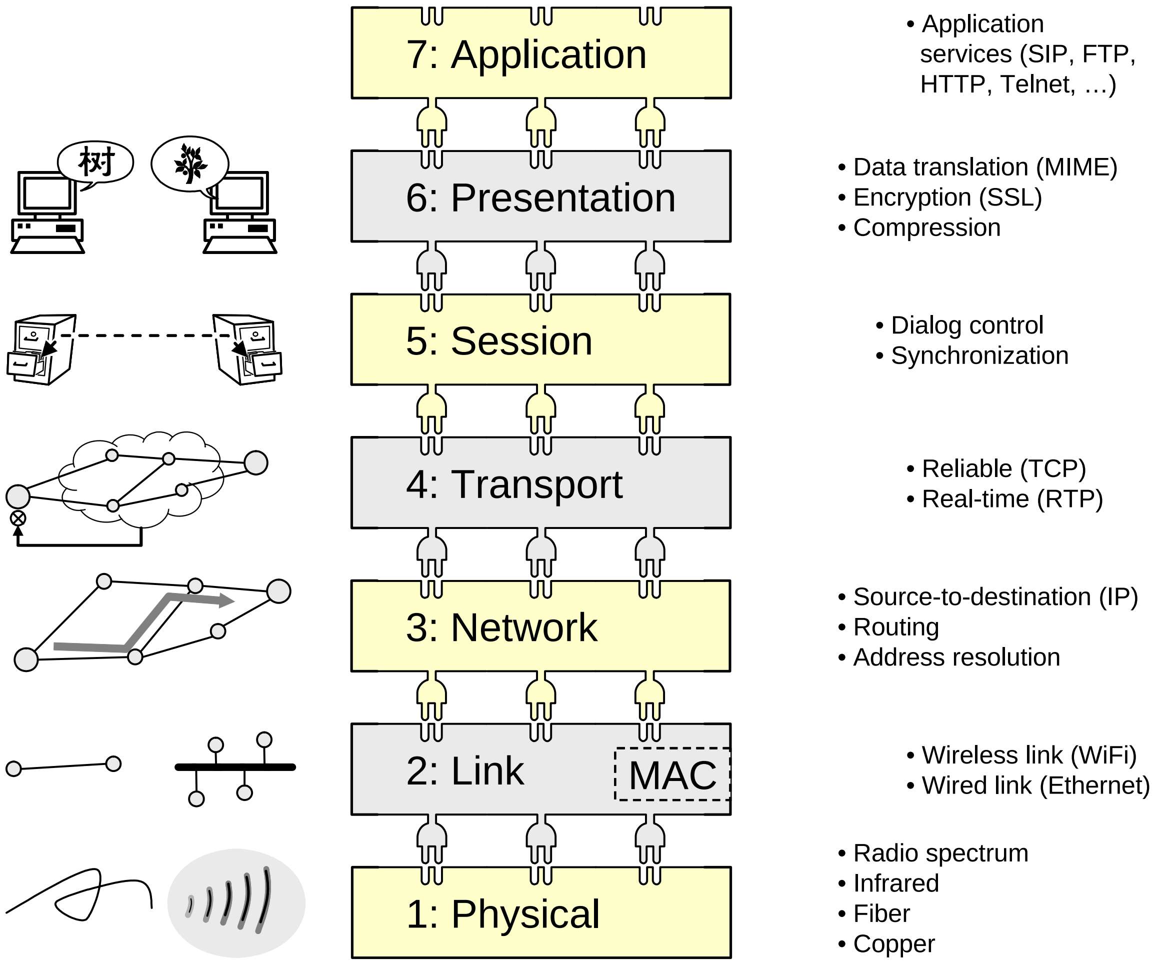

Figure 1-15: OSI reference architecture for communication protocol layering.

The seven layers can be conceptually organized to three subgroups. First, layers 1, 2, and 3—physical, link, and network—are the *network support layers*. They deal with the physical aspects of moving data from one device to another, such as electrical specifications, physical connections, physical addressing, etc. Second, layers 5, 6, and 7—session, presentation, and application—can be thought of as *user support layers*. They allow interoperability among unrelated software systems. Third, layer 4—the transport layer—ensures end-to-end reliable data transmission, while layer 2 may ensure reliable data transmission on a single link.

When compared to the three-layer model (Figure 1-11), OSI layers 1, 2 correspond to layer 1, the Link layer, in the three-layer model. OSI layer 3 corresponds to layer 2, the Network layer, in the three-layer model. Finally, OSI layers 4, 5, 6, and 7 correspond to layer 3, the End-to-end layer, in the three-layer model.

The OSI model serves mainly as a reference for thinking about protocol architecture issues. There are no actual protocol implementations that follow the OSI model. Because it is dated, I will mainly use the three-layer model in the rest of this text.

1.2 Reliable Transmission via Redundancy

To counter the line noise, a common technique is to add redundancy or context to the message. For example, assume that the transmitted word is "information" and the received word is "inrtormation." A human receiver would quickly figure out that the original message is "information," because this is the closest meaningful word to the one that is received. Similarly, assume you are tossing a coin and want to transmit the sequence of outcomes (head/tail) to your friend. Instead of transmitting H or T, for every H you can transmit HHH and for every T you can transmit TTT. The advantage of sending two redundant letters is that if one of the original letters flip, say TTT is sent and TTH is received, the receiver can easily determine that the original message is TTT, which corresponds to "tail." Of course, if two letters become flipped, catastrophically, so TTT turns to THH, then the receiver would erroneously infer that the original is "head." We can make messages more robust to noise by adding greater redundancy. Therefore, instead of two redundant letters, we can have ten: for every H you could transmit HHHHHHHHHH and for every T you could transmit TTTTTTTTTT. The probability that the message will be corrupted by noise catastrophically becomes progressively lower with more redundancy. However, there is an associated penalty: the economic cost of transmitting the longer message grows higher because every communication line can transmit only a limited number of bits per unit of time. Finding the right *tradeoff* between robustness and cost requires the knowledge of the physical characteristics of the transmission line as well as the knowledge about the importance of the message to the receiver (and the sender).

Example of adding redundancy to make messages more robust will be seen in Internet telephony (VoIP), where forward error correction (FEC) is used to counter the noise effects.

If damage/loss can be detected, then an option is to request retransmission but, request + retransmission takes time ⇒ large response latency. FEC is better but incurs overhead.

1.2.1 Error Detection and Correction by Channel Coding

To bring the message home, here is a very simplified example for the above discussion. Notice that this oversimplifies many aspects of error coding to get down to the essence. Assume that you need to transmit 5 different messages, each message containing a single integer number between 1 – 5. You are allowed to "encode" the messages by mapping each message to a number between 1 – 100. Assume that the noise amplitude is distributed according to the normal distribution, as shown in [Figure X]. What are the best choices for the codebook?

Note: this really represents a *continuous* case, not digital, because numbers are not binary and errors are not binary. But just for the sake of simplicity…

1.2.2 Interleaving

Redundancy and error-correcting codes are useful when errors are randomly distributed. If errors are clustered, they are not effective. Consider the following example. Say you want to send the following message to a friend: "*All science is either physics or stamp collecting*."[2] A random noise in the communication channel may result in the following distorted message received by your friend: "*All scidnce is eitjer physocs or statp colletting*." By simply using a spelling checker, your friend may easily recover the original message. One the other hand, if the errors were clustered, the received message may appear as: "*All science is either checker or stamp collecting*." Obviously, it is impossible to guess the original message unless you already know what Rutherford said.

This kind of clustered error is usually caused by a *jamming source*. It may not necessarily be a hostile adversary trying to prevent the communication, but it could be a passive narrow-band jamming source, such as microwave owen, which operates in the same frequency range as Wi-Fi wireless networking technology.

To recover from such errors, one can use **interleaving**. Let us assume that instead of sending the original message as-is, you first scramble the letters and obtain the following message:

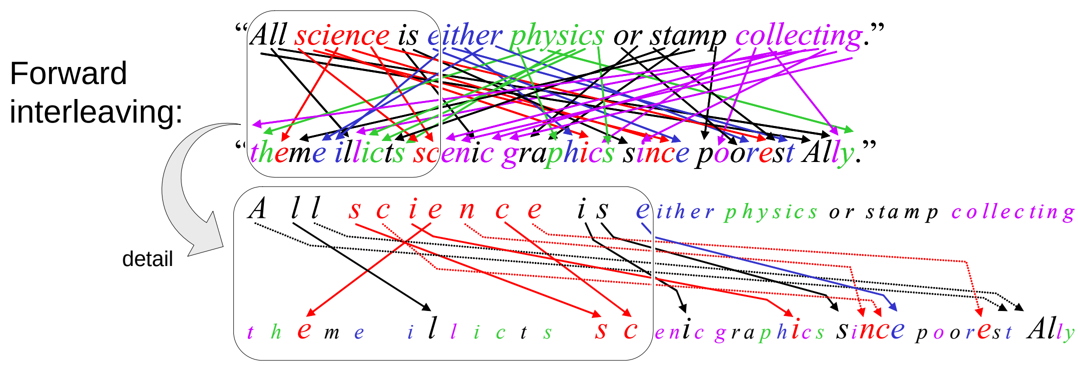

Now you transmit the message "*theme illicts scenic graphics since poorest Ally*." Again, the jamming source inflicts a cluster of errors, so the word "graphics" turns into "strictly," and your friend receives the following message: "*theme illicts scenic strictly since poorest Ally*." Your friend must know how to unscramble the message by applying an inverse mapping to obtain:

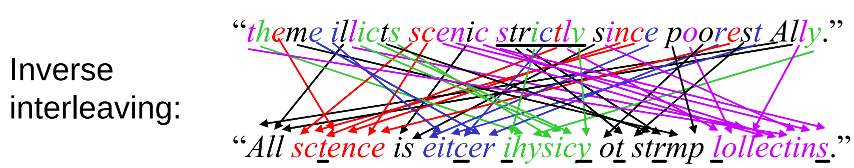

Therefore, with interleaving, the receiver will obtain a message with errors randomly distributed, rather than missing a complete word. By applying a spelling checker, your friend will recover the original message.

---

[2] Ernest Rutherford, in J. B. Birks, "Rutherford at Manchester," 1962.

1.3 Reliable Transmission by Retransmission

We introduced channel encoding as a method for dealing with errors (Section 1.2). But, encoding provides only probabilistic guarantees about the error rates—it can *reduce* the number errors to an arbitrarily small amount, but it cannot *eliminate* them. When error is detected that cannot be corrected, it may be remedied by repeated transmission. This is the task for Automatic Repeat Request (ARQ) protocols. In case retransmission fails, the sender should *persist* with repeated retransmissions until it succeeds or decides to give up. Of course, even ARQ retransmission is a probabilistic way of ensuring reliability and the sender should not persist infinitely with retransmissions. After all, the link to the receiver may be broken, or the receiver may be dead. There is no absolutely certain way to guarantee reliable transmission.

Failed transmissions manifest themselves in two ways:

- Packet error: Receiver receives the packet and discovers error via error control
- Packet loss: Receiver never receives the packet (or fails to recognize it as such)

If the former, the receiver can request retransmission. If the latter, the sender must detect the loss by the lack of response from the receiver within a given amount of time.

Common requirements for a reliable protocol are that: (1) it delivers at most one copy of a given packet to the receiver; and, (2) all packets are delivered in the same order they are presented to the sender. "Good" protocol:

- Delivers a single copy of every packet to the receiver application
- Delivers the packets in the order they were presented to the sender

A lost or damaged packet should be retransmitted. A **persistent sender** is a protocol participant that tries to ensure that at least one copy of each packet is delivered, by sending repeatedly until it receives an acknowledgment. To make retransmission possible, a copy is kept in the transmit buffer (temporary local storage) until it is successfully received by the receiver and the sender received the acknowledgement. Buffering generally uses the fastest memory chips and circuits and, therefore, the most expensive memory, which means that the buffering space is scarce. Disk storage is cheap but not practical for packet buffering because it provides relatively slow data access.

During network transit, different packets can take different routes to the destination, and thus arrive in a different order than sent. The receiver may temporarily store (buffer) the out-of-order packets until the missing packets arrive. Different ARQ protocols are designed by making different choices for the following issues:

- Where to buffer: at sender only, or both sender and receiver?
- What is the maximum allowed number of outstanding packets, waiting to be acknowledged?

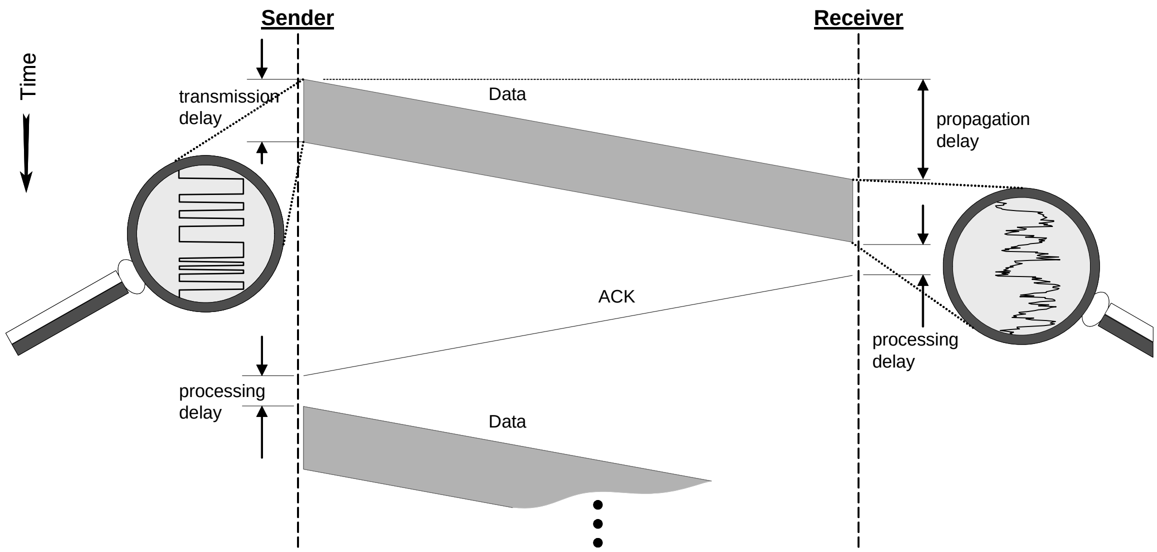

Figure 1-16: Timeline diagram for reliable data transmission with acknowledgements.

- How is a packet loss detected: a timer expires, or the receiver explicitly sends a "negative acknowledgement" (NAK)? (Assuming that the receiver is able to detect a damaged packet.)

The ***sender utilization*** of an ARQ connection is defined as the fraction of time that the sender is busy sending data.

The ***throughput*** of an ARQ connection is defined as the average rate of successful message delivery.

The ***goodput*** of an ARQ connection is defined as the rate at which data are sent uniquely, i.e., this rate does not include error-free data that reach the receiver as duplicates. In other words, the goodput is the fraction of time that the receiver is receiving data that it has not received before.

The transmissions of packets between a sender and a receiver are usually illustrated on a timeline as in Figure 1-16. There are several types of delay associated with packet transmissions. To illustrate, here is an analogy: you are in your office, plan to go home, and on your way home you will stop at the bank to deposit your paycheck. From the moment you start, you will get down to the garage ("transmission delay"), drive to the bank ("propagation delay"), wait in the line ("queuing delay"), get served at teller's window ("processing delay" or "service delay"), and drive to home (additional "propagation delay").

The first delay type is **transmission delay**, which is the time that takes the sender to place the data bits of a packet onto the transmission medium. In other words, transmission delay is measured from when the first bit of a packet enters a link until the last bit of that same packet enters the link. This delay depends on the transmission rate R offered by the medium (in bits per second or bps), which determines how many bits (or pulses) can be generated per unit of time at the transmitter. It also depends on the length L of the packet (in bits). Hence, the transmission delay is:

$$t_x = \frac{\text{packet length}}{\text{bandwidth}} = \frac{L \text{ (bits)}}{R \text{ (bits per second)}} \tag{1.2}$$

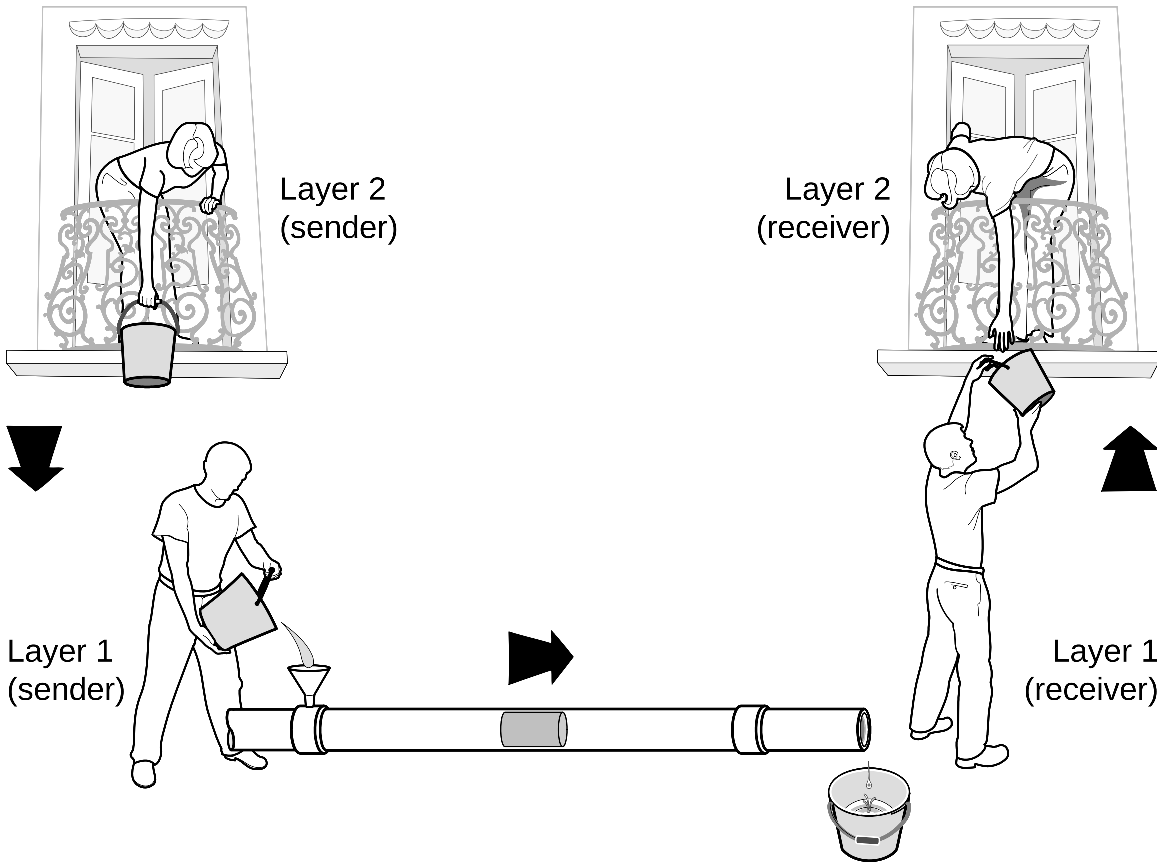

Figure 1-17: Fluid flow analogy for delays in packet delivery between the protocol layers.

Propagation delay is defined as the time elapsed between when a bit is sent at the sender and when it is received at the receiver. This delay depends on the distance *d* between the sender and the receiver and the velocity *v* of electromagnetic waves in the transmission medium, which is proportional to the speed of light in vacuum ($c \approx 3\times10^8$ m/s), $v = c/n$, where *n* is the index of refraction of the medium. Both in copper wire and glass fiber or optical fiber $n \approx 3/2$, so $v \approx 2 \times 10^8$ m/s. The index of refraction for dry air is approximately equal to 1. The propagation delay is:

$$t_p = \frac{\text{distance}}{\text{velocity}} = \frac{d\ (\text{m})}{v\ (\text{m/s})} \tag{1.3}$$

Processing delay is the time needed for processing a received packet. At the sender side, the packet may be received from an upper-layer protocol or from the application. At the receiver side, the packet is received from the network or from a lower-layer protocol. Examples of processing include conversion of a stream of bytes to frames or packets (known as *framing* or *packetization*), data compression, encryption, relaying at routers, etc. Processing delays usually can be ignored when looking from an end-host's viewpoint. However, processing delay is very critical for routers in the network core that need to relay a huge number of packets per unit of time, as will be seen later in Section 1.4.4.

Another important parameter is the **round-trip time** (or **RTT**), which is the time a bit of information takes from departing until arriving back at the sender if it is immediately bounced back at the receiver. This time on a single transmission link is often assumed to equal RTT =

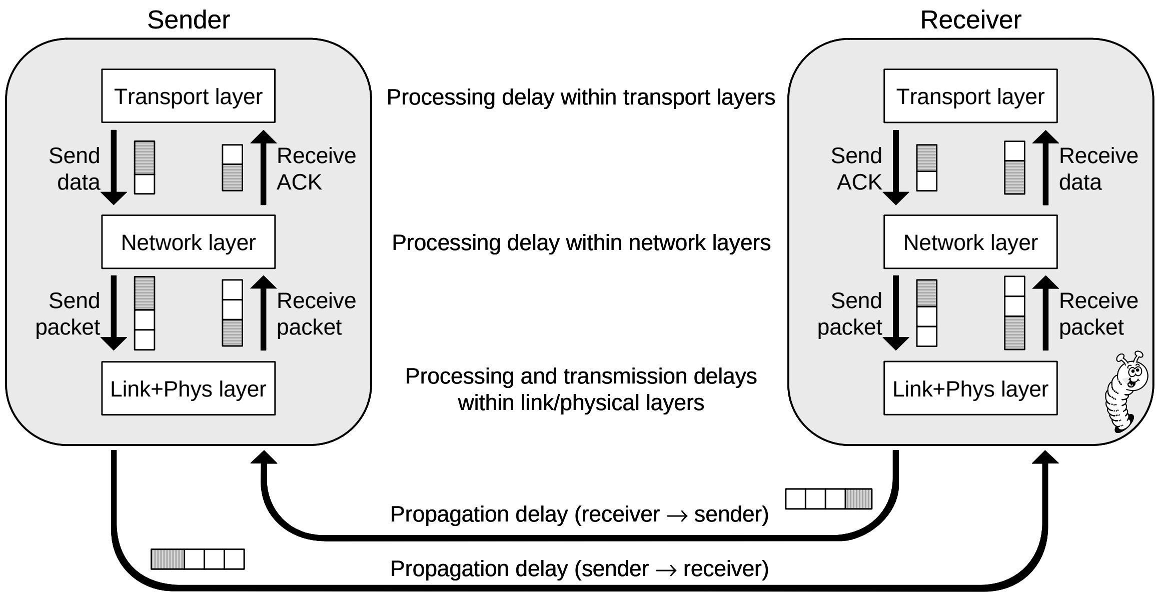

Figure 1-18: Delay components that contribute to round-trip time (RTT).

$2 \times t_p$. Determining the RTT is much more complex if the sender and receiver are connected over a network where multiple alternative paths exist, as will be seen later in Section 2.1.2. However, even on a single link, the notion of RTT is much more complex than just double the propagation delay. To better understand RTT, we need to consider what it is used for and how it is measured. RTT is most often used by sender to set up its retransmission timer in case a packet is lost. Obviously, network nodes do not send individual bits; they send packets. RTT is measured by recording the time when a packet is sent, reading out the time when the acknowledgement is received, and subtracting these two values:

$$\text{RTT} = (\text{time when the acknowledgement is received}) - (\text{time when the packet is sent}) \quad (1.4)$$

To understand what contributes to RTT, we need to look at how packets travel through the network. First, acknowledgements may be *piggybacked* on data packets coming back for the receiver. Therefore, even if the transmission delay is not included at the sender side, receiver's transmission delay does contribute to the RTT. (However, when an acknowledgement is piggybacked on a regular data packet from receiver to sender, the transmission time of this packet must be taken into account.)

Second, we have to remember that network nodes use layered protocols (Section 1.1.4). Continuing with the fluid flow analogy from Figure 1-5, we illustrate in Figure 1-17 how delays are introduced between the protocol layers. The physical-layer (layer 1) receiver waits until the bucket is full (i.e., the whole packet is received) before it delivers it to the upper layer (layer 2).

The delay components for a single link and a three-layer protocol are illustrated in Figure 1-18. Sender's transmission delay will not be included in the measured RTT only if the sender operates at the link/physical layer. A sender operating at any higher layer (e.g., network or transport layers), cannot avoid having the transmission delay included in the measured RTT, because it cannot know when the packet transmission on the physical medium will actually start or end.

Third, lower layers of sender's protocol stack may incur significant *processing delays*. Suppose that the sender is at the transport layer and it measures the RTT to receive the acknowledgement from the receiver, which is also at the transport layer. When a lower layer receives a packet from a higher layer, the lower layer may not forward the packet immediately, because it may be busy with sending some other packets. Also, if the lower layer uses error control, it will incur processing delay while calculating the checksum or some other type of error-control code. Later we will learn about other types of processing delays, such as time spent looking up forwarding tables in routers (Section 1.4), time spent dividing a long message into fragments and later reassembling it (Section 1.4.1), time spent compressing data, time spent encrypting and decrypting message contents, etc.

Fourth, lower layers may implement their own reliable transmission service, which is transparent to the higher layer. An example are broadcast links (Section 1.3.3), which keep retransmitting lost packets until a retry-limit is reached. The question, then, is: what counts as the transmission delay for a packet sent by a higher layer and transmitted by a lower layer, which included several retransmissions? Should we count only the successful transmission (the last one), or the preceding unsuccessful transmissions, as well?

In summary, the reader should be aware that RTT estimation is a complex issue even for a scenario of a single communication link connecting the sender and receiver. Although RTT is often approximated as double the propagation delay, this may be grossly inaccurate and the reader should examine the feasibility of this approximation individually for each scenario.

Mechanisms needed for reliable transmission by retransmission:

- Error detection for received packets, e.g., by checksum
- Receiver feedback to the sender, via acknowledgement or negative acknowledgement
- Retransmission of a failed packet, which requires storing the packet at the sender until the sender obtains a positive acknowledgement that the packet reached the receiver error-free
- Sequence numbers, so the receiver can distinguish duplicate packets
- Retransmission timer, if packet loss on the channel is possible (not only error corruption), so that the sender can detect the loss

Several popular ARQ protocols are described next.

1.3.1 Stop-and-Wait

Problems related to this section: Problem 1.2 → Problem 1.4; also see Problem 1.12

The simplest retransmission strategy is *stop-and-wait*. This protocol buffers only a single packet at the sender and does not deal with the next packet before ensuring that the current packet is correctly received (Figure 1-16). A packet loss is detected by the expiration of a timer, which is set when the packet is transmitted.

When the sender receives a corrupted ACK/NAK, it could send back to the receiver a NAK (negative acknowledgement). For pragmatic reasons (to keep the sender software simple), receiver does nothing and the sender just re-sends the packet when its retransmission timer expires.

Assuming error-free communication, the utilization of a Stop-and-wait sender is determined as follows. The entire cycle to transport a single packet takes a total of $(t_x + 2 \times t_p)$ time. (We assume that the acknowledgement packets are tiny, so their transmission time is negligible.) Of this time, the sender is busy t_x time. Therefore

$$U_{sender}^{S\&W} = \frac{t_x}{t_x + 2 \cdot t_p} \tag{1.5}$$

Given a probability of packet transmission error p_e, which can be computed using Eq. (1.1), we can determine *how many times*, on average, a packet will be (re-)transmitted until successfully received and acknowledged. This is known as the **expected number of transmissions**. Our simplifying assumption is that error occurrences in successively transmitted packets are independent events[3]. A successful transmission in one round requires error-free transmission of two packets: forward data and feedback acknowledgement. We again assume that these are independent events, so the joint probability of success is

$$p_{\text{succ}} = \left(1 - p_e^{\text{DATA}}\right) \cdot \left(1 - p_e^{\text{ACK}}\right) \tag{1.6}$$

The probability of a failed transmission in one round is $p_{\text{fail}} = 1 - p_{\text{succ}}$. Then, the number of attempts K needed to transmit successfully a packet is a *geometric random variable*. The probability that the first k attempts will fail and the $(k+1)^{\text{st}}$ attempt will succeed equals:

$$Q(0, k+1) = \binom{k}{0} \cdot \left(1 - p_{\text{succ}}\right)^k \cdot p_{\text{succ}}^1 = p_{\text{fail}}^k \cdot p_{\text{succ}} \tag{1.7}$$

where $k = 1, 2, 3, \ldots$. The round in which a packet is successfully transmitted is a random variable N, with the probability distribution function given by (1.7). Its expected value is

$$E\{N\} = \sum_{k=0}^{\infty} (k+1) \cdot Q(0, k+1) = \sum_{k=0}^{\infty} (k+1) \cdot p_{\text{fail}}^k \cdot p_{\text{succ}} = p_{\text{succ}} \cdot \left(\sum_{k=0}^{\infty} p_{\text{fail}}^k + \sum_{k-0}^{\infty} k \cdot p_{\text{fail}}^k \right)$$

Recall that the well-known summation formula for the geometric series is

$$\sum_{k=0}^{\infty} x^k = \frac{1}{1-x}, \qquad \sum_{k=0}^{\infty} k \cdot x^k = \frac{x}{(1-x)^2}$$

Therefore we obtain (recall that $p_{\text{fail}} = 1 - p_{\text{succ}}$):

$$E\{N\} = p_{\text{succ}} \cdot \left(\frac{1}{1 - p_{\text{fail}}} + \frac{p_{\text{fail}}}{(1 - p_{\text{fail}})^2} \right) = \frac{1}{p_{\text{succ}}} \tag{1.8}$$

We can also determine the **average delay per packet** as follows. Successful transmission of one packet takes a total of $t_{\text{succ}} = t_x + 2 \times t_p$, assuming that transmission time for acknowledgement packets can be ignored. A single failed packet transmission takes a total of $t_{\text{fail}} = t_x + t_{\text{out}}$, where t_{out} is retransmission timer's countdown time. If a packet is successfully transmitted after k failed

[3] This is valid only if we assume that thermal noise alone affects packet errors. However, the independence assumption will not be valid for temporary interference in the environment, such as a microwave oven interference on a wireless channel.

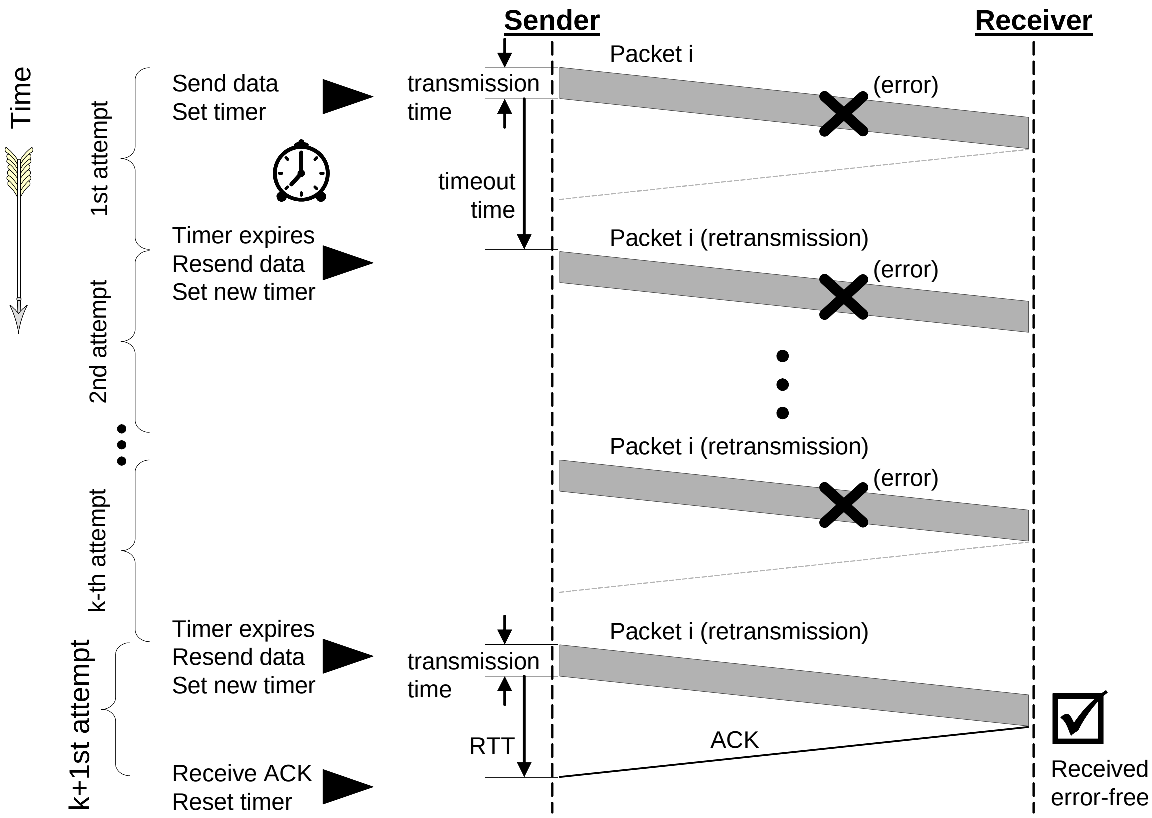

Figure 1-19: Stop-and-Wait with errors. The transmission succeeds after *k* failed attempts.

attempts, then its total transmission time equals: $T_{k+1}^{\text{total}} = k \cdot t_{\text{fail}} + t_{\text{succ}}$, where k = 0, 1, 2, … (see Figure 1-19). The total transmission time for a packet is a random variable T_{k+1}^{total}, with the probability distribution function given by (1.7). Its expected value is

$$E\{T^{\text{total}}\} = \sum_{k=0}^{\infty} \left(k \cdot t_{\text{fail}} + t_{\text{succ}}\right) \cdot p_{\text{fail}}^{k} \cdot p_{\text{succ}} = p_{\text{succ}} \cdot \left(t_{\text{succ}} \sum_{k=0}^{\infty} p_{\text{fail}}^{k} + t_{\text{fail}} \sum_{k=0}^{\infty} k \cdot p_{\text{fail}}^{k} \right)$$

Following a derivation similar as for Eq. (1.8), we obtain

$$E\{T^{\text{total}}\} = p_{\text{succ}} \cdot \left(\frac{t_{\text{succ}}}{1 - p_{\text{fail}}} + \frac{p_{\text{fail}} \cdot t_{\text{fail}}}{\left(1 - p_{\text{fail}}\right)^2} \right) = t_{\text{succ}} + \frac{p_{\text{fail}}}{p_{\text{succ}}} \cdot t_{\text{fail}} \tag{1.9}$$

The expected sender utilization in case of a noisy link is

$$E\left\{U_{sender}^{S\&W}\right\} = \frac{t_x \cdot E\{N\}}{E\{T^{\text{total}}\}} = \frac{t_x}{p_{\text{succ}} \cdot t_{\text{succ}} + p_{\text{fail}} \cdot t_{\text{fail}}} \tag{1.10}$$

Here, we are considering the expected fraction of time the sender will be busy of the total expected time to transmit a packet successfully. That is, ($t_x \cdot E\{N\}$) includes both unsuccessful and successful (the last one) transmissions.

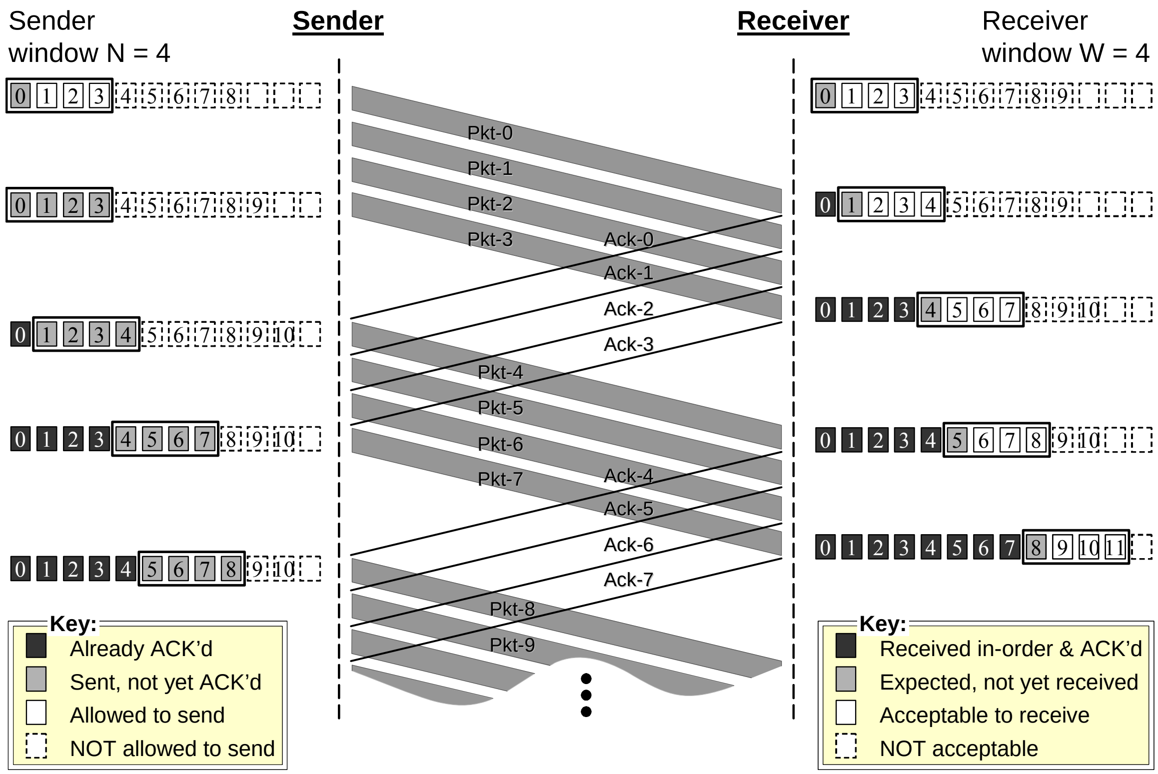

Figure 1-20: Sliding window protocol in operation over an error-free link.

1.3.2 Sliding-Window Protocols

Problems related to this section: Problem 1.5 → Problem 1.12

Stop-and-wait is very simple but also very inefficient, because the sender spends most of the time idle waiting for the acknowledgement. We would like the sender to send as much as possible, short of causing path congestion or running out of the memory space for buffering copies of the outstanding packets. One type of ARQ protocols that offer higher efficiency than Stop-and-wait is the *sliding window protocol*.

The **sender window size** N is a measure of the maximum number of outstanding (i.e., unacknowledged) packets in the network. Figure 1-20 shows the operation of sliding window protocols in case of no errors in communication. The **receiver window size** W gives the upper bound on the number of out-of-order packets that the receiver is willing to accept. In the case shown in Figure 1-20, both sender and receiver have the same window size. In general case it is required that $N \leq W$.

The sliding window sender should store in local memory (buffer) all outstanding packets for which the acknowledgement has not yet been received. Therefore, the send buffer size should be N packets large. The sent-but-unacknowledged packets are called "in-flight" packets or "in-transit" packets. The sender must ensure that the number of in-flight packets is always $\leq N$.

The sliding window protocol is actually a family of protocols that have some characteristics in common and others different. Next, we review two popular types of sliding window protocols:

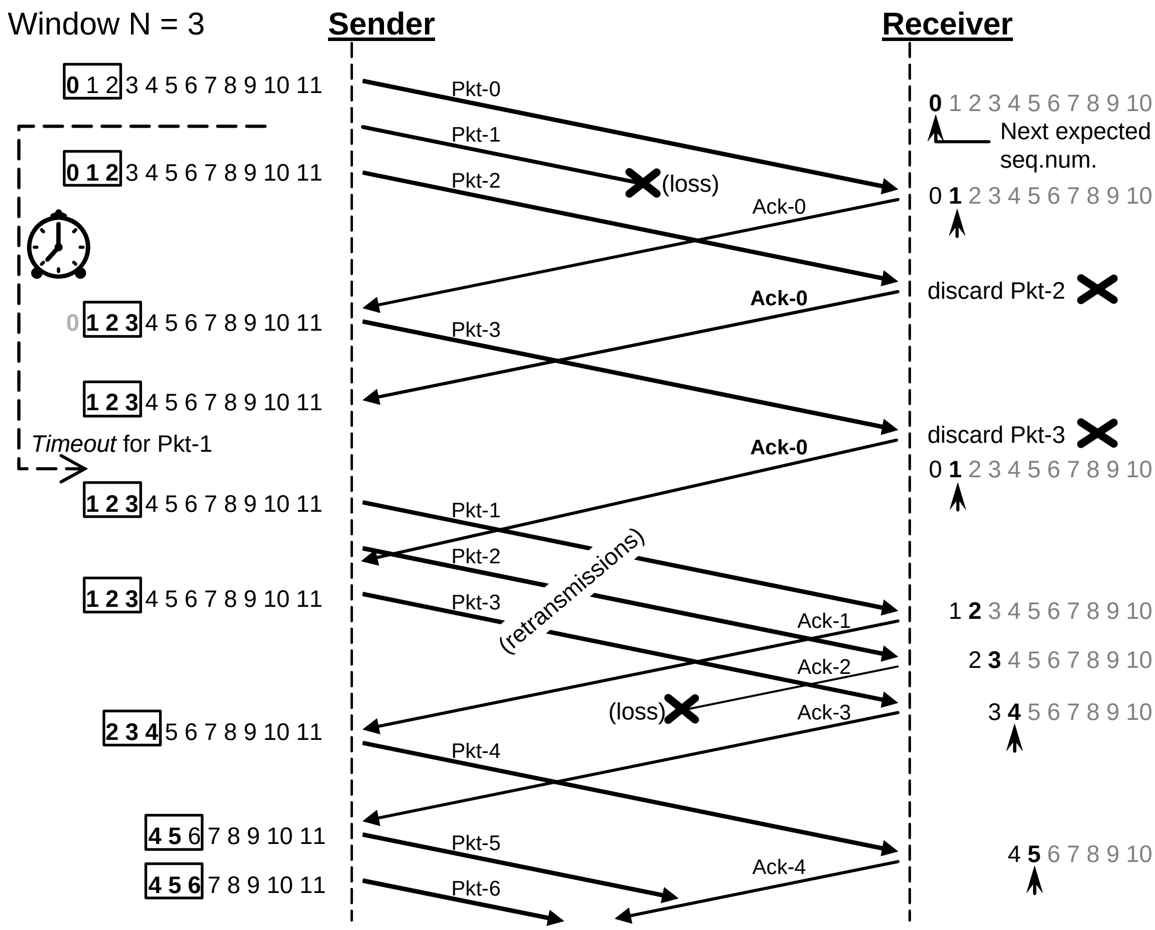

Figure 1-21: Go-back-N protocol in operation under communication errors.

Go-back-*N* (GBN) and Selective Repeat (SR). The TCP protocol described in Chapter 2 is another example of a sliding window protocol. The key difference between GBN and SR is in the way they deal with communication errors.

Go-back-*N*

The key idea of the Go-back-*N* protocol is to have the receiver as simple as possible. This means that the receiver accepts only the next expected packet and immediately discards any packets received out-of-order. Hence, the receiver needs only to memorize what is the next expected packet (single variable), and does not need any memory to buffer the out-of-order packets.

As with other sliding-window protocols, the Go-back-*N* sender should be able to buffer up to *N* outstanding packets.

The operation of Go-back-*N* is illustrated in Figure 1-21. The sender sends sender-window-size (N = 3) packets and stops, waiting for acknowledgements to arrive. When Ack-0 arrives, acknowledging the first packet (Pkt-0), the sender slides its window by one and sends the next available packet (Pkt-3). The sender stops again and waits for the next acknowledgement.

Because Pkt-1 is lost, it will never be acknowledged and its retransmission timer will expire. When a timeout occurs, the Go-back-*N* sender resends *all* packets that have been previously sent

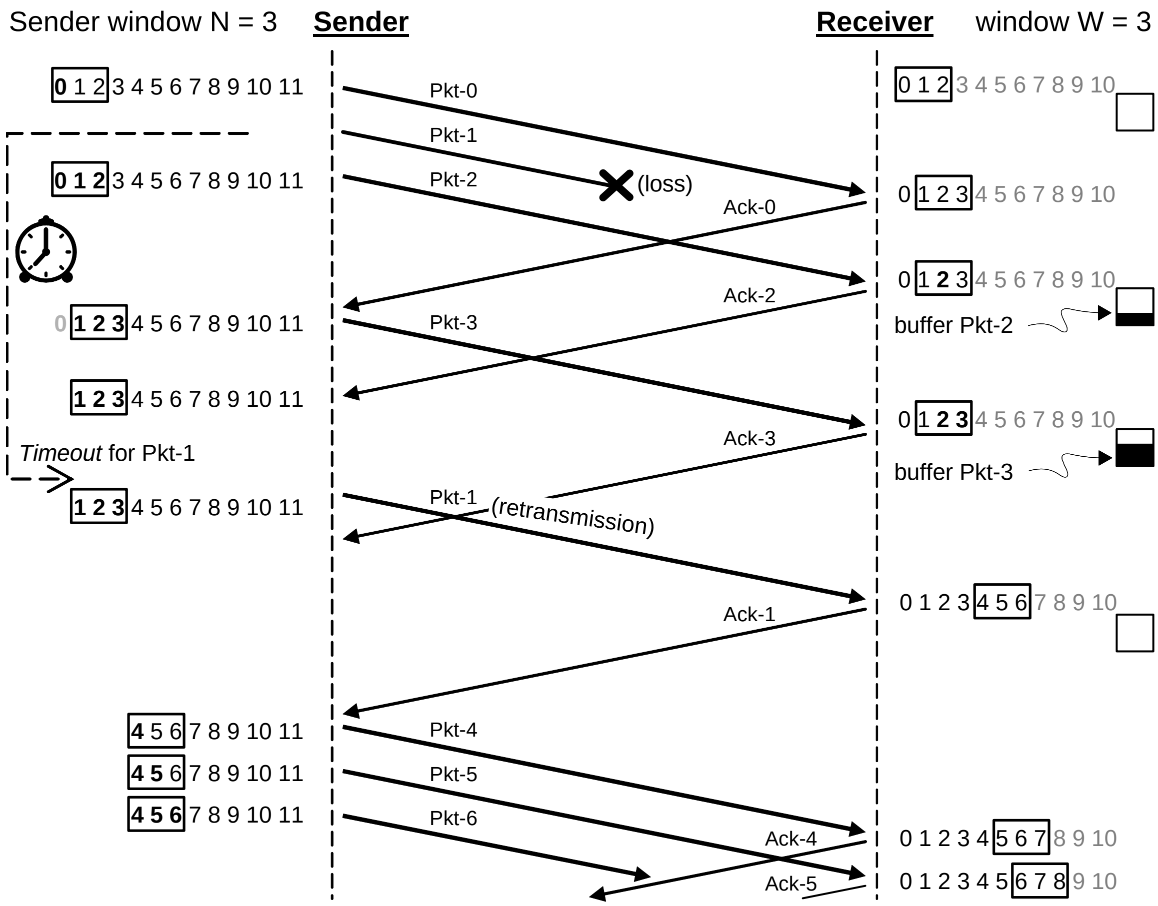

Figure 1-22: Selective Repeat protocol in operation under communication errors.

but have not yet been acknowledged, that is, all "in-flight" packets. This is where this protocol's name comes from. Because the sender will usually have N in-flight packets, a timeout will cause it to go back by N and resend the N outstanding packets. The rationale for this behavior is that if the oldest outstanding packet is lost, then all the subsequent packets are lost as well, because the Go-back-N receiver automatically discards out-of-order packets.

As mentioned, the receiver memorizes a single variable, which is the sequence number of the next expected packet. The Go-back-N receiver considers a packet *correctly received* if and only if

1. The received packet is error-free
2. The received packet arrived in-order, i.e., its sequence number equals next-expected-sequence-number.

In this example, Pkt-1 is lost, so Pkt-2 arrives out of order. Because the Go-back-N receiver discards any packets received out of order, Pkt-2 is automatically discarded. One salient feature of Go-back-N is that the receiver sends **cumulative acknowledgements**, where an acknowledgement with sequence number m indicates that all packets with a sequence number up to and including m have been correctly received at the receiver. The receiver sends acknowledgement even for incorrectly received packets, but in this case, the previously correctly received packet is being acknowledged. In Figure 1-21, the receipt of Pkt-2 generates a duplicate acknowledgement Ack-0. Notice also that when Ack-2 is lost, the sender takes Ack-3 to

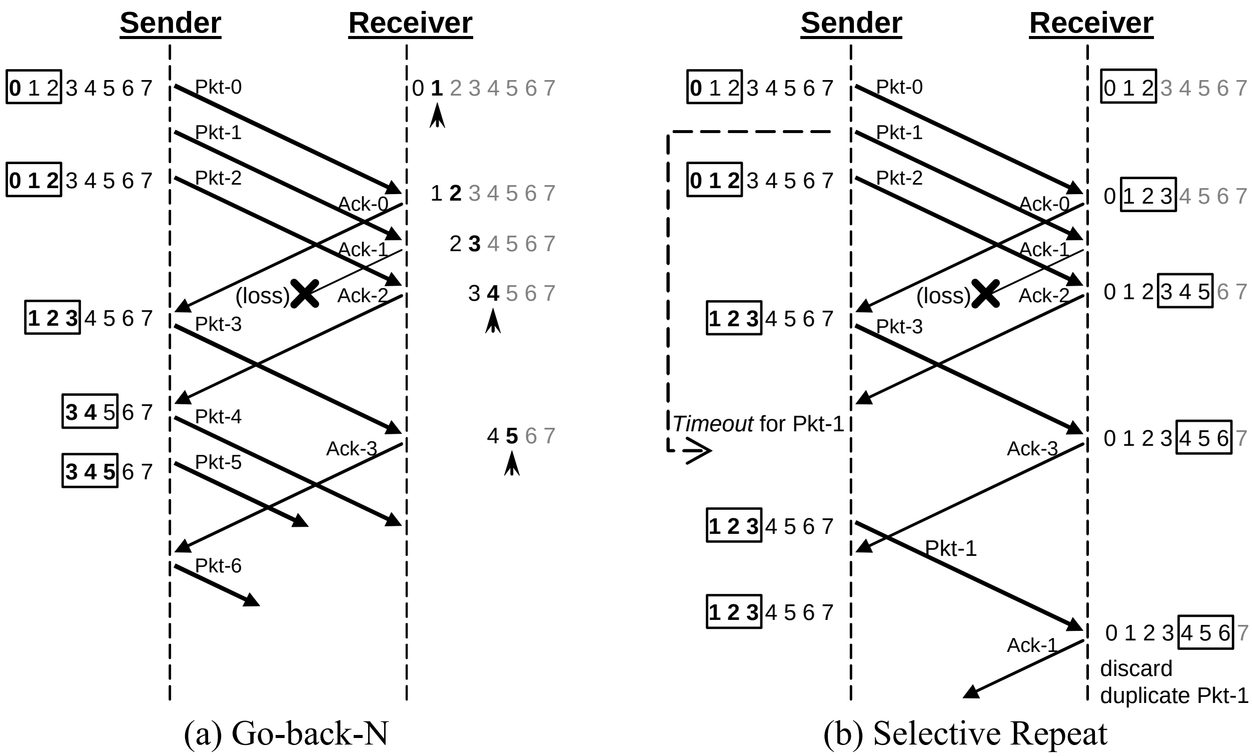

Figure 1-23: Comparison of Go-back-N and Selective-Repeat acknowledgments.

acknowledge all previous packets, including Pkt-2. Hence, a lost acknowledgement does not need to be retransmitted as long as the acknowledgement acknowledging the following packet arrives before the retransmission timer expires.

Selective Repeat (SR)

The key idea of the Selective Repeat protocol is to avoid discarding packets that are received error-free and, therefore, to avoid unnecessary retransmissions. Go-back-*N* suffers from performance problems, because a single packet error can cause Go-back-*N* sender to retransmit a large number of packets, many of them unnecessarily. Figure 1-22 illustrates the operation of the Selective Repeat protocol. Unlike a Go-back-*N* sender which retransmits all outstanding packets when a retransmission timer times out, a Selective Repeat sender retransmits only a single packet—the oldest outstanding one.

Unlike a Go-back-*N* receiver, a Selective Repeat receiver sends **individual acknowledgements**, where an acknowledgement with sequence number *m* indicates only that the packet with sequence number *m* has been correctly received. There is no requirement that packets are received in order. If a packet is received out of order but error-free, it will be buffered in the receiver's memory until the in-order missing packets are received.

Figure 1-23 illustrates the difference between the behaviors of GBN cumulative acknowledgements and SR individual acknowledgements. Notice that both protocols require the sender to acknowledge duplicate packets, which were received and acknowledged earlier. The reason for this requirement is that a duplicate packet usually indicates that the acknowledgement has been lost. Without an acknowledgement, the sender window would never move forward and the communication would come to a halt. (Notice also that a packet can be retransmitted when its acknowledgement is delayed, so the timeout occurs before the acknowledgement arrives. In this case, the sender window would simply move forward.) Again, SR acknowledges only the last received (duplicate) packet, whereas GBN *cumulatively* acknowledges all the packets received up to and including the last one. In Figure 1-23(a), Ack-1 was lost, but when Ack-2 arrives it acknowledges all packets up to and including Pkt-2. This acknowledgement shifts the sender's window forward by 2 and the sender advances uninterrupted. Unlike this, in Figure 1-23(b) the SR sender needs to retransmit Pkt-1 because it never receives Ack-1 before its timeout expired.

In practice, a combination of selective-ACK and Go-back-*N* is used, as will be seen with TCP in Chapter 2.

SIDEBAR 1.1: The Many Faces of Acknowledgements

♦ The attentive reader may have noticed that acknowledgements are used for multiple purposes. For example, earlier we saw that a received ACK informs the sender that: (a) the corresponding data packet arrived at the receiver; (b) the sender may stop the retransmission-timer countdown for the corresponding data packet; (c) the sender may discard the copy of the acknowledged packet and release the memory buffer space; and, (d) the sender may send another data packet. Notice that an acknowledgment usually only confirms that the data packet arrived error-free to the receiver, but it does not say anything about whether the receiver acted upon the received data and completed the required processing. This requires additional acknowledgement at the application level. Later, in Chapter 2, we will learn about some additional uses of acknowledgements in the TCP protocol.

1.3.3 Broadcast Links

Broadcast links allow connecting multiple network nodes via the same link. Hence, when one node transmits, all or most other nodes on the link can hear the transmission. If two or more nodes are transmitting simultaneously, their signals will interfere with each other (see Figure 1-7 for interference on a wireless link). A receiver that receives the interference signal will not be able to decode either of the original signals; this is known as a **collision**. Therefore, the nodes should take turns transmitting their packets. However, this is easier said than done: when a node has a packet ready for transmission, it does not know whether any other nodes are also about to transmit. A key technical problem for broadcast links is **coordination of transmissions**, to control collisions.

There are several techniques for transmission coordination on broadcast links. Collisions could be prevented by designing strictly timed transmissions; avoided by listening before speaking; or, detected after they happen and remedied by retransmission of corrupted information. An example of preventing collisions by design is TDMA (Time Division Multiple Access). It creates unique time slots and assigns a different time slot to each node. A node is allowed to transmit only within its assigned time slot. After all nodes are given opportunity to transmit, the cycle is repeated. The problem with this technique is that if some nodes do not have data ready for transmission, their

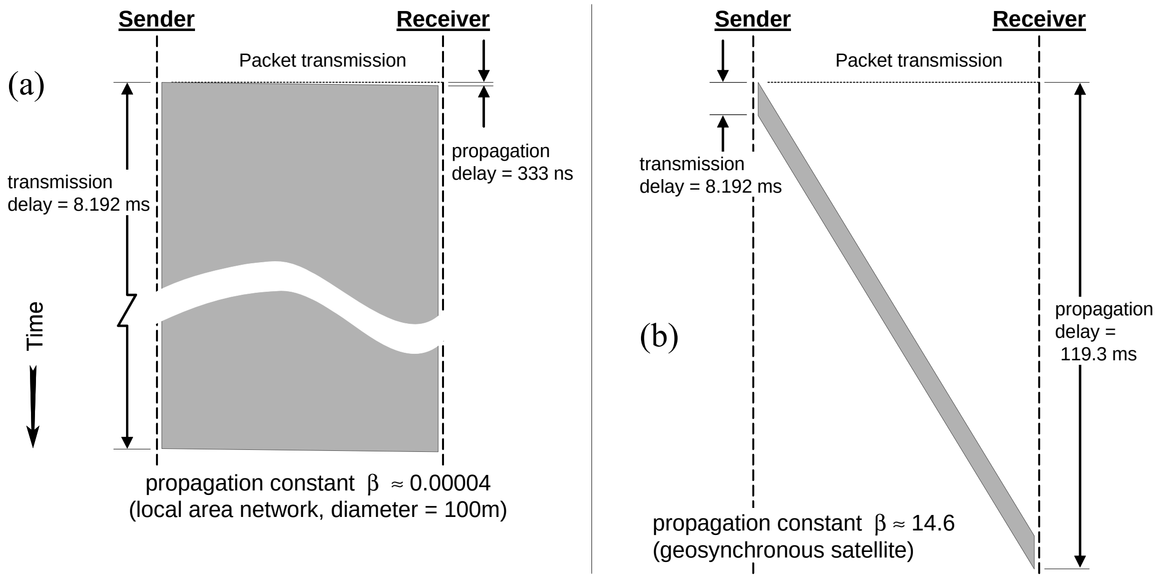

Figure 1-24: Propagation constant β for different wireless networks.

slot goes unused. Any other nodes that may wish to transmit are delayed and have to wait for their predetermined slot even though the link is currently idle.

A popular class of protocols for broadcast links is **random-access protocols**, which are based on the Stop-and-Wait ARQ, with addition of the backoff delay mechanism. **Backoff mechanism** is a key mechanism for coordination of multiple senders on a broadcast link. Stop-and-wait has no reason to backoff because it assumes that the sender is *not* contending for the link against other senders and any loss is due to a transmission error. Conversely, a random-access protocol assumes that any packet loss is due to a collision of concurrent senders and it tries to prevent further collisions by introducing a random amount of delay (backoff) before attempting a re-transmission. It is a way to provide stations with "polite behavior." This method is commonly used when multiple concurrent senders are competing for the same resource; another example will be seen for the TCP protocol (Section 2.1.2). The sender usually doubles the range of backoff delays for every failed transmission, which is why this method is also known as **binary exponential backoff**. Increasing the backoff range increases the number of choices for the random delay. This, in turn, makes it less likely that several stations will select the same delay value and, therefore, reduces the probability of repeated collisions. It is like first deciding how long to wait by tossing a coin (two choices: heads or tails); if both make the same choice and again experience a collision, then they try by rolling a dice (six choices), etc.

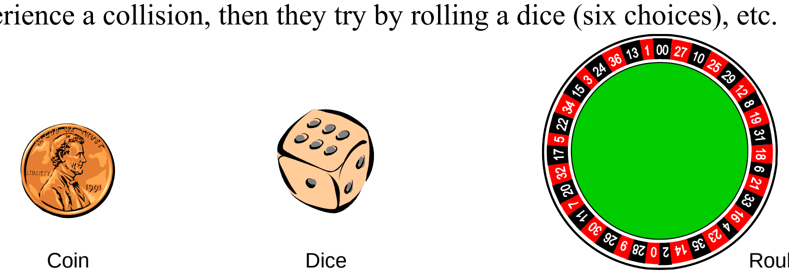

Coin
two choices
{0, 1}

Dice
six choices
{1, 2, 3, 4, 5, 6}

Roulette
38 choices (American)
{0, 00, 1, 2, 3, …, 36}

The reason for addressing reliability at the link layer is as follows. A wireless link is significantly more unreliable than a wired one. Noise, interference, and other propagation effects result in the loss of a significant number of frames. Even with error-correction codes, a number of MAC frames may not successfully be received. This situation can be dealt with by reliability mechanisms at a higher layer, such as transport-layer protocol. However, timers used for retransmission at higher layers (which control paths comprising many links) are typically on the order of seconds (see TCP timers in Section 2.1.2). It is therefore more efficient to deal with errors at the link level and retransmit the corrupted packets.

The time (in packet transmission units) required for all network nodes to detect a start of a new transmission or an idle channel after a transmission ends is an important parameter. Intuitively, the **parameter β** is the number of bits that a transmitting station can place on the medium before the station furthest away receives the first bit of the first packet.

Recall that signal propagation time is t_p = distance/velocity, as given earlier by Eq. (1.3). The transmission delay is t_x = packet-length/bandwidth, as given by Eq. (1.2). The parameter β is calculated as

$$\beta = \frac{t_p}{t_x} = \frac{d \cdot R}{v \cdot L} \tag{1.11}$$

The velocity of electromagnetic waves in dry air equals $v \approx 3 \times 10^8$ m/s, and in copper or optical fiber it equals $v \approx 2 \times 10^8$ m/s. Therefore, propagation time is between 3.33 and 5 nanoseconds per meter (ns/m). Given a wireless local area network (W-LAN) where all stations are located within a 100 m diameter, the (maximum) propagation delay is $t_p \approx 333$ ns. If the bandwidth (or, data rate) of the same W-LAN is 1 Mbps, the transmission delay for a 1 Kbytes packet equals t_x = 8.192 ms. The relationship is illustrated in Figure 1-24(a). Recall from Figure 1-6 that on a 1 Mbps link, 1 bit is 1 μs wide, so the leading edge of the first bit will reach the receiver long before the sender is done with the transmission of this bit. In other words, the propagation delay is practically negligible. On the other hand, the altitude of a geosynchronous satellite is 35,786 km above the Earth surface, so the propagation delay is $t_p \approx 119.3$ ms. As shown in Figure 1-24, the respective β parameters for these networks are $\beta_{\mathrm{LAN}} \approx 0.00004$ and $\beta_{\mathrm{GSS}} \approx 14.6$. The time taken by the electronics for detection should also be added to the propagation time when computing β, but it is usually ignored as negligible. We will see later how parameter β plays an important role in network design.

The ALOHA Protocol

Problems related to this section: Problem 1.13 → Problem 1.15

A simple protocol for broadcast media is called ALOHA. There are two versions of ALOHA: *pure* or *plain ALOHA* transmits packets as soon as they become ready, and *slotted ALOHA* which transmits packets only at regular intervals. The state diagram for the sender side of both variations of the protocol is shown in Figure 1-25. Plain ALOHA sends the packet immediately as it becomes ready, while slotted ALOHA has to wait for the start of the next time interval (or, **slot**). In other words, in slotted ALOHA, all transmissions are strictly clocked at regular time intervals. After transmission, the sender stops-and-waits for the acknowledgement. If the acknowledgement arrives, this is the end of the current cycle, and the sender expects the next packet to become available for sending. If the acknowledgement does not arrive, the sender assumes that this is because collision happened and it increases its backoff interval. The **backoff interval** is the

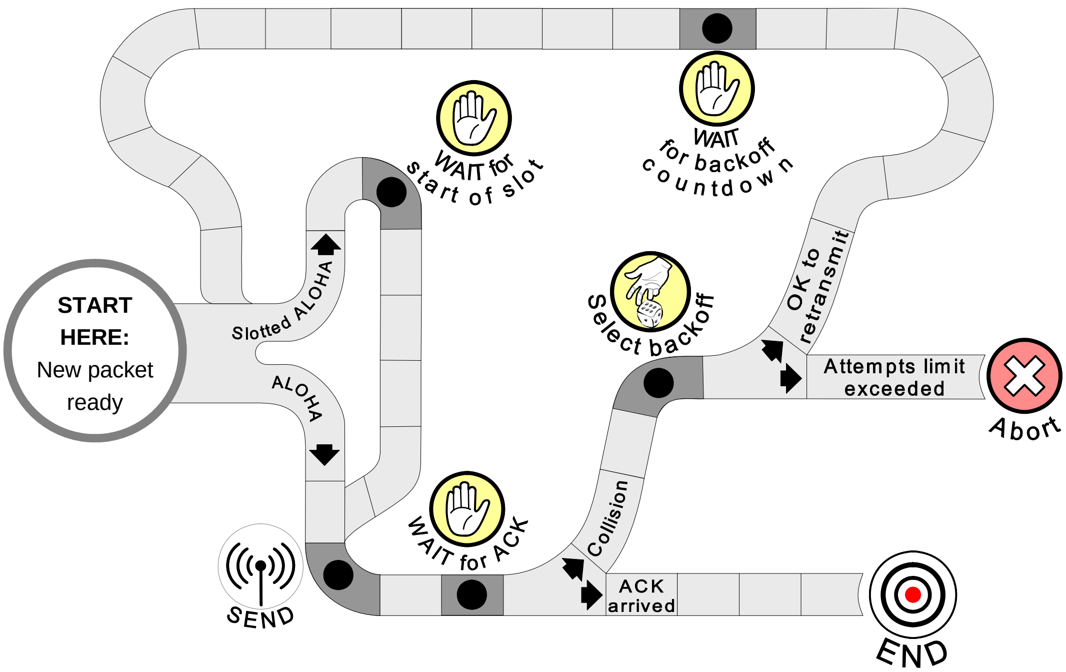

Figure 1-25: The sender's state diagram for ALOHA and Slotted ALOHA protocols.

amount of time the sender waits to reduce the probability of collision with another sender that also has a packet ready for transmission. After waiting for backoff countdown, the sender repeats the cycle and retransmits the packet. As we know from the earlier discussion, the sender does not persist forever in resending the packet, and if it exceeds a given threshold, the sender gives up and aborts the retransmissions of this packet.

ALOHA is a very simple protocol, almost identical to Stop-and-Wait ARQ, except for the backoff interval. ALOHA also does not initiate transmission of the next packet before ensuring that the current packet is correctly received. Let us first consider a pure ALOHA protocol. To derive its throughput, we make the following assumptions:

- There are a total of m wireless nodes and each node generates new packets for transmission according to a Poisson process (see Appendix) with rate λ/m.
- Each node can hold only a single packet at a time, and the packet is stored until the node receives a positive acknowledgement that the packet is successfully transmitted. While storing the packet, the node is said to be *backlogged*.
- When a new packet is generated, what happens to it depends on whether or not the node is already backlogged. If the node it is backlogged, the newly generated packet is discarded; if the node is *not* backlogged, the newly generated packet is immediately transmitted (in pure ALOHA) and stored until acknowledgement is received.
- A backlogged node can retransmit at any moment with a certain probability.
- All packets have the same length. The time needed for packet transmission (transmission delay) is called *slot* length, and it is normalized to equal 1.
- If only a single node transmits, the transmission is always successful (noiseless channel); if two or more nodes transmit simultaneously, there will be a collision and all collided

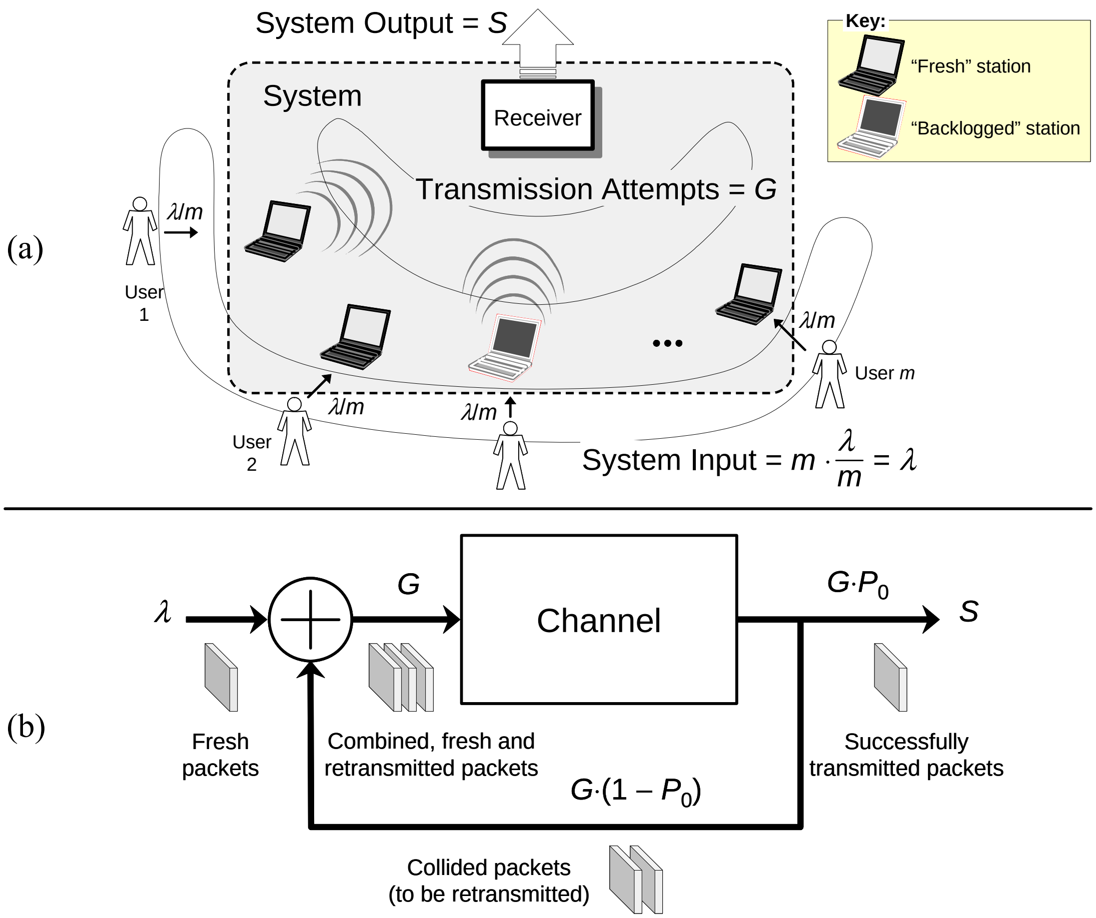

Figure 1-26: (a) ALOHA system representation. (b) Modeled as a feedback system.

packets will be lost; all nodes receive *instantaneous feedback* about the success or failure of the transmission. In other words, the acknowledgement is received *immediately* upon packet transmission, without any propagation delays.

This system can be modeled as in Figure 1-26. For a reasonable throughput, we would expect $0 < \lambda < 1$ because the system can successfully carry at most one packet per slot, i.e., only one node can "talk" (or, transmit) at a time. Also, for the system to function, the departure rate of packets out from the system should equal the arrival rate in equilibrium. In equilibrium, on one hand, the departure rate cannot physically be greater than the arrival rate; on the other hand, if it is smaller than the arrival rate, all the nodes will eventually become backlogged.

The following simplified derivation yields a reasonable approximation. In addition to the new packets, the backlogged nodes generate retransmissions of the packets that previously suffered collisions. If the retransmissions are sufficiently randomized, it is plausible to approximate the total number of transmission attempts per slot, retransmissions and new transmissions combined, as a Poisson random variable with some parameter $G > \lambda$.

The probability of successful transmission (i.e., throughput S) is the probability of an arrival times the probability that the packet does not suffer collision; because these are independent events, the joint probability is the product of their probabilities. The probability of an arrival is

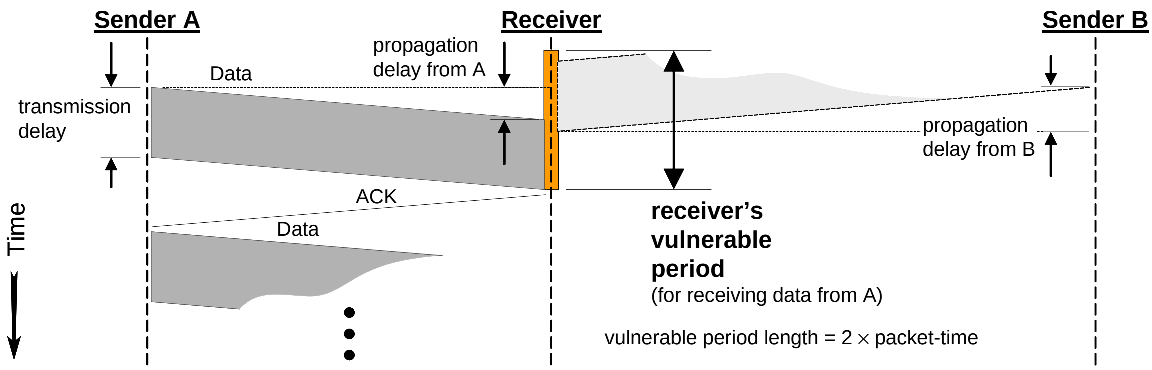

Figure 1-27: The receiver's vulnerable period during which collisions are possible.

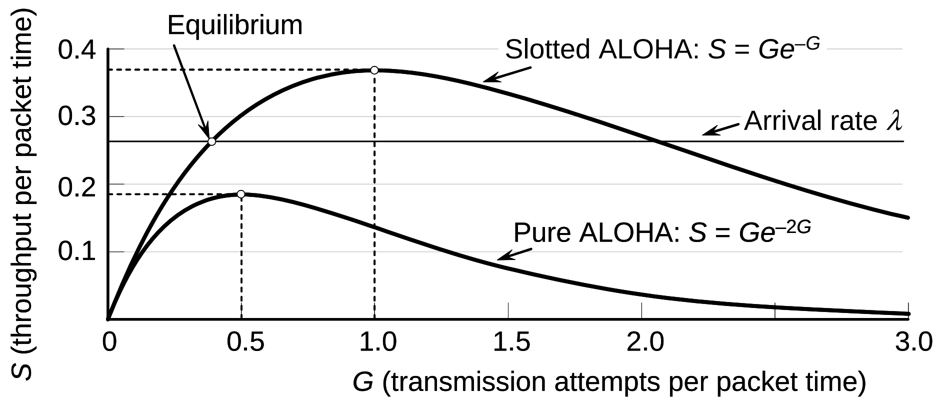

Figure 1-28: Efficiency of the ALOHA MAC protocol. (In the case of Slotted ALOHA, the packet time is equal to the *slot time*.)

$P_a = \tau \cdot G$, where $\tau = 1$ is the slot duration and G is the total arrival rate on the channel (new and backlogged packets, combined).

The packet will not suffer collision if no other senders have transmitted their packets during the so-called "vulnerable period" or "window of vulnerability." We define receiver's **vulnerable period** as the time during which other transmissions may cause collision with sender's transmission. For pure ALOHA, the vulnerable period is two slots long $[t - 1, t + 1)$, as illustrated in Figure 1-27. Any transmission that started within one packet-time before this transmission or during this transmission will overlap with this transmission and result in a collision. For slotted ALOHA, the vulnerable period lasts one slot $[t, t + 1)$, assuming that all stations are synchronized and can start transmission only at slot intervals. From the Poisson distribution formula (see Appendix A), $P_0 = P\{A(t + \tau) - A(t) = 0\}$. With $\tau = 1$ for slotted ALOHA, we have

$$S = P_a \cdot P_0 = (1 \cdot G) \cdot P\{A(t+1) - A(t) = 0\} = G \cdot e^{-G} \tag{1.12}$$

For pure ALOHA, $\tau = 2$, so $S = G \cdot e^{-2G}$. In equilibrium, the arrival rate (system input), λ, to the system should be the same as the departure rate (system output), $S = G \cdot e^{-G}$. The reader should recall Figure 1-26(a), and this relationship is illustrated in Figure 1-28.

We see that for slotted ALOHA, the maximum possible throughput of $1/e \approx 0.368$ occurs at $G = 1$. This is reasonable, because if $G < 1$, too many idle slots are generated, and if $G > 1$, too many collisions are generated. At $G = 1$, the packet departure rate is one packet per packet time (or, per

Table 1-1: Characteristics of three basic CSMA protocols when the channel is sensed idle or busy. If a transmission was unsuccessful, all three protocols perform backoff and repeat.

CSMA Protocol	Sender's listening-and-transmission rules
Nonpersistent	If medium is idle, transmit. If medium is busy, wait random amount of time and sense channel again.
1-persistent	If medium is idle, transmit (i.e., transmit with probability 1). If medium is busy, *continue sensing* until channel is idle; then transmit immediately (i.e., transmit with probability 1).
p-persistent	If medium is idle, transmit with probability p. If medium is busy, *continue sensing* until channel is idle; then transmit with probability p.

slot), the fraction $1/e$ of which are newly arrived packets and $1-\frac{1}{e}$ are the successfully retransmitted backlogged packets.

Carrier Sense Multiple Access Protocols (CSMA)

Problems related to this section: Problem 1.17 → ?

The key problem with the ALOHA protocol is that it employs a very simple strategy for coordinating the transmissions: a node transmits a new packet as soon as it is created, and in case of collision, it retransmits with a retransmission probability.

An improved coordination strategy is to have the nodes "listen before they talk." That is, the sender listens to the channel before transmitting and transmits only if the channel is detected as idle. Listening to the channel is known as **carrier sense**, which is why this strategy has the name *carrier sense multiple access* (CSMA).

The medium is decided *idle* if there are no transmissions for time duration the parameter β time units, because this is the propagation delay between the most distant stations in the network. The time taken by the electronics for detection should also be added to the propagation time when computing channel-sensing time, but it is usually ignored as negligible.

The key issues with a listen-before-talk approach are:

(1) When to listen and, in case the channel is found busy, whether to keep listening until it becomes idle or stop listening and try later

(2) Whether to transmit immediately upon finding the channel idle or slightly delay the transmission

Upon finding the channel busy, the node might listen persistently until the end of the ongoing transmission. Another option is to listen periodically. Once the channel becomes idle, the node might transmit immediately, but there is a danger that some other nodes also waited ready for transmission, which would lead to a collision. Another option is, once the channel becomes idle, to hold the transmission briefly for a random amount of time, and only if the channel remains idle, start transmitting the packet. This reduces the chance of a collision significantly, although it does not remove it, because both nodes might hold their transmissions for the same amount of

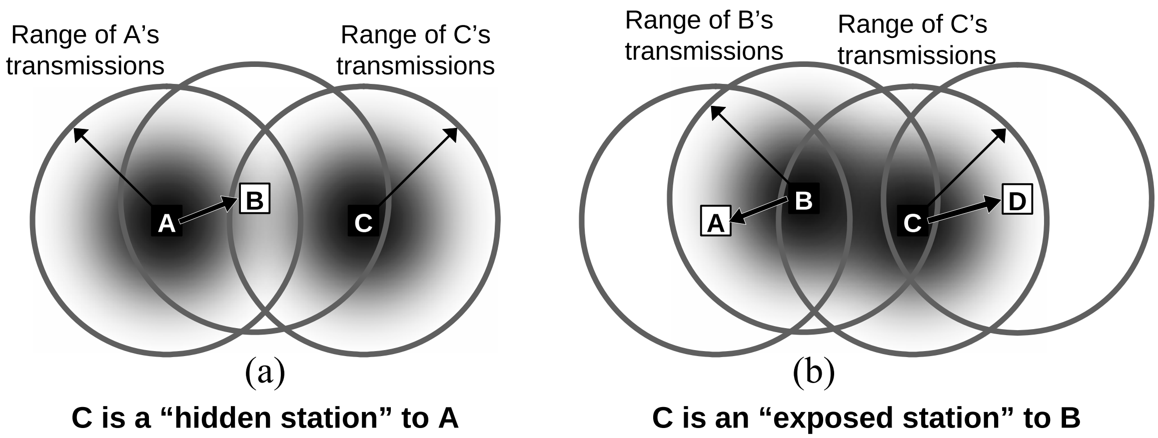

Figure 1-29: (a) Hidden station problem: *C* cannot hear *A*'s transmissions. (b) Exposed station problem: *C* defers transmission to *D* because it hears *B*'s transmission.

time. Several CSMA protocols that make different choices regarding the listening and transmission start are shown in Table 1-1. For each of the protocols in Table 1-1, when the sender discovers that a transmission was unsuccessful (by a retransmission timer timeout), the sender behaves the same way: it inserts a randomly distributed retransmission delay (backoff) and repeats the listening-and-transmission procedure.

The efficiency of CSMA is better than that of ALOHA because of CSMA's shorter vulnerable period: The stations will not initiate transmission if they sense a transmission already in progress. Notice that nonpersistent CSMA is less greedy than 1-persistent CSMA in the sense that, upon observing a busy channel, it does not continually sense it with intention of seizing it immediately upon detecting the end of the previous transmission (Table 1-1). Instead, nonpersistent CSMA waits for a random period and then repeats the procedure. Consequently, this protocol leads to better channel utilization but longer delays than 1-persistent CSMA.

Wireless broadcast networks show some phenomena not present in wireline broadcast networks. The air medium is partitioned into broadcast regions, rather than being a single broadcast medium. This is simply due to the exponential propagation loss of the radio signal, as discussed earlier in Section 1.1.2. As a result, two interesting phenomena arise: (*i*) not all stations within a partition can necessarily hear each other; and, (*ii*) the broadcast regions can overlap. The former causes the hidden station problem and the latter causes the exposed station problem.

Unlike the wireline broadcast medium, the transitivity of connectivity does not apply. In wireline broadcast networks, such as Ethernet, if station *A* can hear station *B* and station *B* can hear station *C*, then station *A* can hear station *C*. This is not always the case in wireless broadcast networks, as seen in Figure 1-29(a). In the **hidden station problem**, station *C* cannot hear station *A*'s transmissions and may mistakenly conclude that the medium is available. If *C* does start transmitting, it will interfere at *B*, wiping out the frame from *A*. Generally, a station *X* is considered to be hidden from another station *Y* in the same receiver's area of coverage if the transmission coverages of the transceivers at *X* and *Y* do not overlap. A station that can sense the transmission from both the source and receiver nodes is called **covered station**.

Different air partitions can support multiple simultaneous transmissions, which are successful as long as each receiver can hear at most one transmitter at a time. In the **exposed station problem**, station *C* defers transmission to *D* because it hears *B*'s transmission, as illustrated in Figure 1-29(b). If *C* senses the medium, it will hear an ongoing transmission and falsely conclude that it may not send to *D*, when in fact such a transmission would cause bad reception only in the zone between *B* and *C*, where neither of the intended receivers is located. Thus, the carrier sense mechanism is insufficient to detect all transmissions on the wireless medium.

Hidden and exposed station problems arise only for CSMA-type protocols. ALOHA, for instance, does not suffer from such problems because it does not perform channel sensing before transmission (i.e., it does not listen before talking). Under the hidden stations scenario, the performance of CSMA degenerates to that of ALOHA, because carrier-sensing mechanism essentially becomes useless. With exposed stations it becomes worse because carrier sensing prevents the exposed stations from transmission, where ALOHA would not mind the busy channel.

CSMA/CD

Problems related to this section: ? → ?

Persistent and nonpersistent CSMA protocols are clearly an improvement over ALOHA because they ensure that no station begins to transmit when it senses the channel busy. Another improvement is for stations to abort their transmissions as soon as they detect a collision[4]. Quickly terminating damaged packets saves time and bandwidth. This protocol is known as *CSMA with Collision Detection*, or CSMA/CD, which is a variant of 1-persistent CSMA. It works as follows (Figure 1-30):

[4] In networks with wired media, the station compares the signal that it places on the wire with the one observed on the wire to detect collision. If these signals are not the same, a collision has occurred.

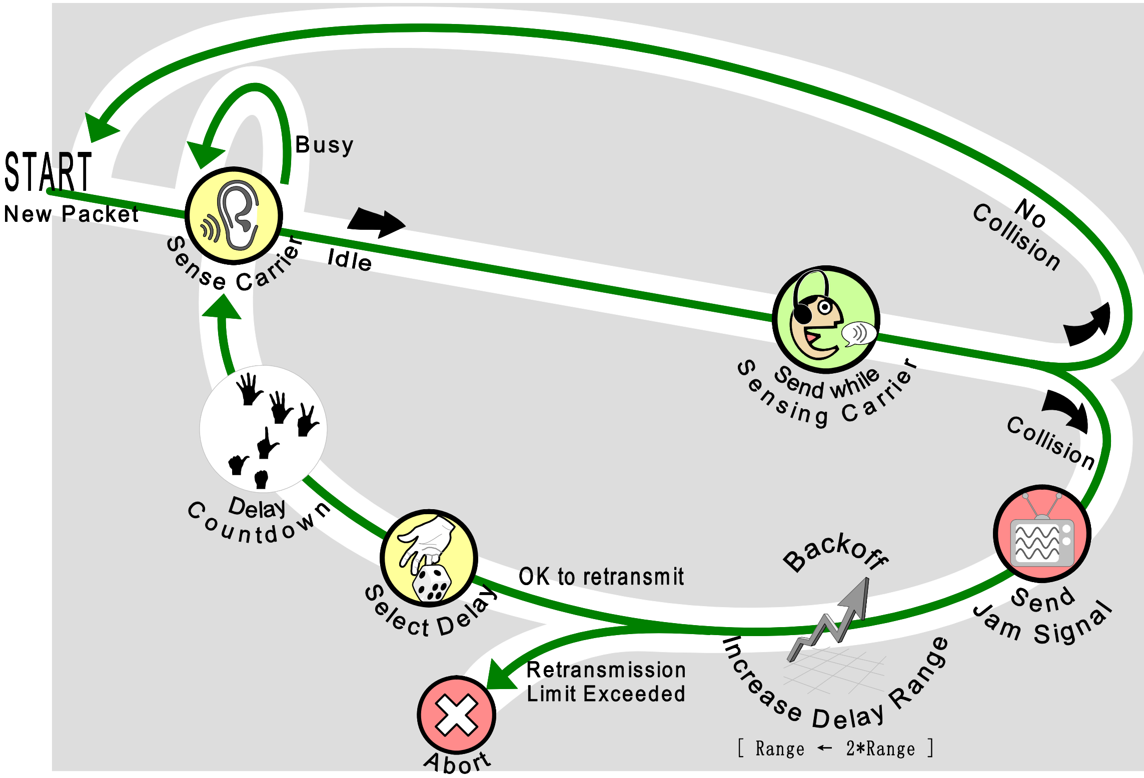

Figure 1-30: The sender's state diagram for CSMA/CD protocol.

1. Wait until the channel is idle.
2. When the channel is idle, transmit immediately *and* sense the carrier during the transmission (or, "listen *while* talking").
3. If you detect collision, abort the ongoing packet transmission, double the backoff range, choose a random amount of backoff delay, wait for this amount of delay, and go to step 1.

A given station can experience a collision during the initial part of its transmission (the **collision window**) before its transmitted signal has had time to propagate to all stations on the CSMA/CD medium. Once the collision window has passed, a transmitting station is said to have *acquired the medium*; subsequent collisions are avoided because all other stations can be assumed to have noticed the signal and to be deferring to it. The time to acquire the medium is thus based on the round-trip propagation time. If the station transmits the complete frame successfully and has additional data to transmit, it will again listen to the channel before attempting a transmission (Figure 1-30).

The collision detection process is illustrated in Figure 1-31. At time t_0 both stations are listening ready to transmit. The CSMA/CD protocol requires that the transmitter detect the collision before it has stopped transmitting its frame. Therefore, the transmission time of the smallest frame must be larger than one round-trip propagation time, i.e., 2β, where β is the propagation constant described in Figure 1-24. The station that detects collision must transmit a **jam signal**, which carries a special binary pattern to inform the other stations that a collision occurred. The jam pattern consists of 32 to 48 bits. The transmission of the jam pattern ensures that the collision lasts long enough to be detected by all stations on the network.

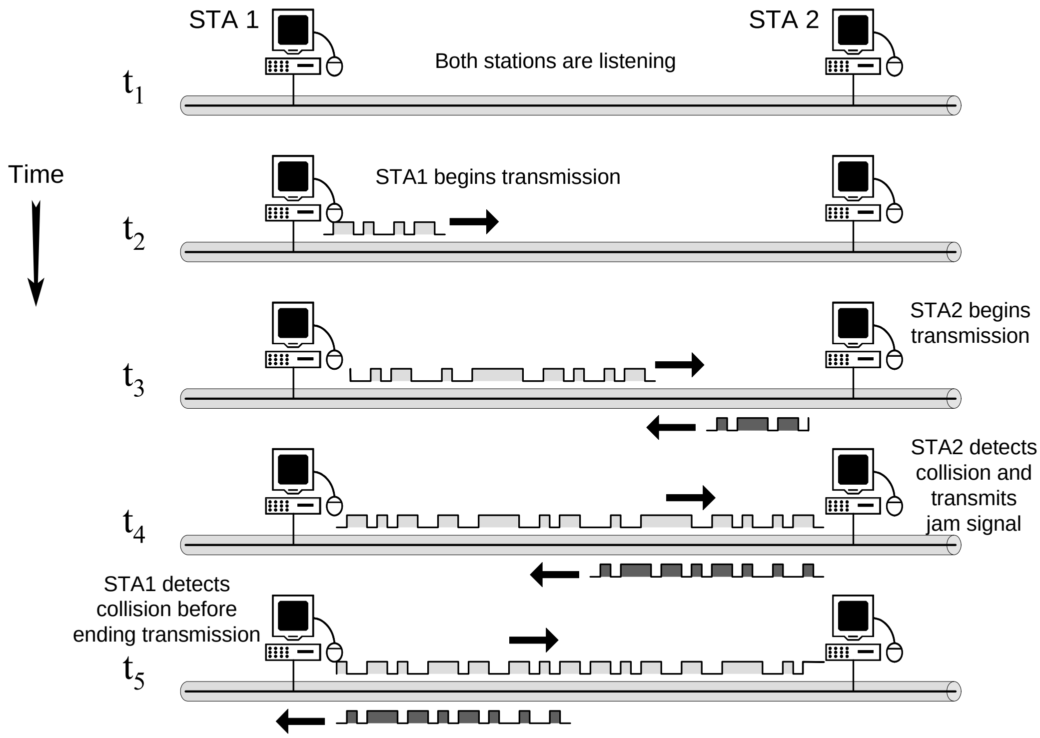

Figure 1-31: Collision detection by CSMA/CD stations.

It is important to realize that collision detection is an analog process. The station's hardware must listen to the cable while it is transmitting. If what it reads back is different from what it is putting out, it knows that a collision is occurring.

After k collisions, a random number of slot times is chosen from the backoff range $[0, 2^k - 1]$. After the first collision, each sender might wait 0 or 1 slot times. After the second collision, the senders might wait 0, 1, 2, or 3 slot times, and so forth. As the number of retransmission attempts increases, the number of possibilities for the choice of delay increases. The backoff range is usually *truncated*, which means that after a certain number of increases, the retransmission timeout reaches a ceiling and the exponential growth stops. For example, if the ceiling is set at k=10, then the maximum delay is 1023 slot times. In addition, as shown in Figure 1-30, the number of attempted retransmissions is limited, so that after the maximum allowed number of retransmissions the sender gives up and aborts the retransmission of this frame. The sender resets its backoff parameters and retransmission counters at the end of a successful transmission or if the transmission is aborted.

Notice that CSMA/CD achieves reliable transmission without acknowledgements. If the sender does not detect collision, this means that the sender has not detected any errors during the transmission. Therefore, it simply assumes that the receiver received the same signal (i.e., the frame was received error free), and there is no need for an acknowledgement.

Here is an example:

Example 1.1 Illustration of a Timing Diagram for CSMA/CD

Consider a local area network of three stations using the CSMA/CD protocol shown in Figure 1-30. At the end of a previous transmission, station-1 and station-3 each have one frame to transmit, while

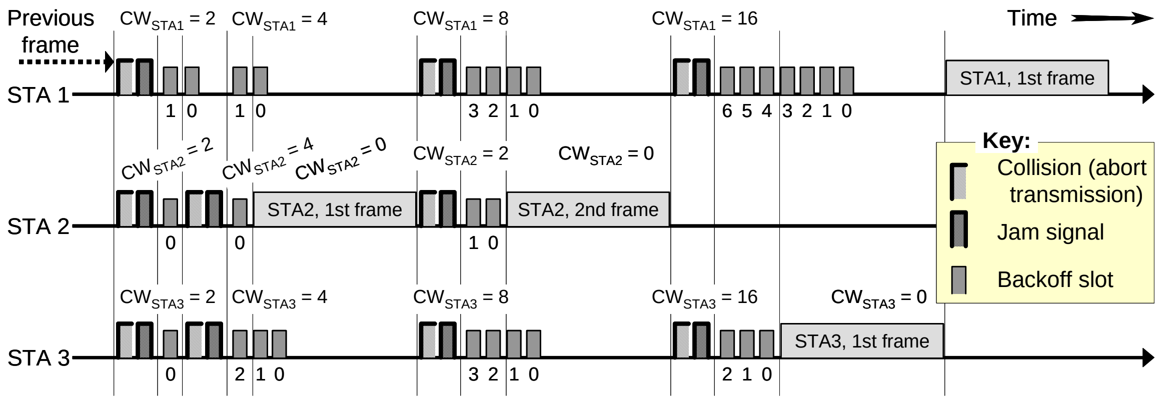

Figure 1-32: Example of three CSMA/CD stations transmission with collision and backoff.

station-2 has two frames. Assume that all frames are of the same length. After the first collision assume that the randomly selected backoff values are: STA1 = 1; STA2 = 0; STA3=0. Next, after the second collision, the backoff values are: STA1 = 1; STA2 = 0; STA3=2. Then, after the third collision, the backoff values are: STA1 = 3; STA2 = 1; STA3=3. Finally, after the fourth collision the backoff values are: STA1 = 6; (STA2 is done by now); STA3=2. Show the timing diagram and indicate the contention window (CW) sizes.

The solution is shown in Figure 1-32. Initially, all three stations attempt transmission and there is a collision; they all detect the collision, abort their transmissions in progress, and send the jam signal. After this, all three stations set their contention window (*CW*) size to 2 and randomly choose their delay periods from the set {0, …, *CW*} = {0, 1}. As given in the problem statement, station-1 chooses its backoff delay as 1, while stations 2 and 3 booth choose their backoff delay as 0. This leads to the second collision. After the second backoff delay, station-2 succeeds in transmitting its first frame and resets its backoff parameters (including the contention window *CW*) to their default values. The other two stations keep the larger ranges of the contention window because they have not successfully transmitted their frames yet. This gives station-2 an advantage after the third collision. Because it chooses the backoff delay from a shorter range of values (*CW*=2), it is more likely to select a small value and, therefore, again succeed in transmitting another frame.

To derive the performance of the CSMA/CD protocol, we will assume a network of m stations with heavy and constant load, where all stations are always ready to transmit. We make a simplifying assumption that there is a constant retransmission probability in each slot. If each station transmits during a contention slot with probability p, the probability A that some station will acquire the channel in that slot is

$$A = m \cdot p \cdot (1 - p)^{m-1}$$

A is maximized when $p = 1/m$, with $A \to 1/e$ as $m \to \infty$. Next, we calculate the average number of contention slots that a station wastes before it succeeds in transmitting a packet. The probability that the station will suffer collision $(j - 1)$ times as succeed on the j^{th} attempt (i.e., that the contention interval has exactly j slots in it) is $A \cdot (1 - A)^{j-1}$. Therefore, the average number of slots per contention is given as the expected value

$$\sum_{j=0}^{\infty} j \cdot A \cdot (1 - A)^{j-1} = \frac{1}{A}$$

Because each slot is $2\cdot\beta$ long, the mean contention interval, w, is $2\cdot\beta/A$. Assuming optimal p, the average number of contention slots is never more than e, so w is at most $2\cdot\beta\cdot e \approx 5.4\times\beta$. If an average frame takes $t_x = L/R$ seconds to transmit, then the channel efficiency is

$$\eta_{\text{CSMA/CD}} = \frac{L/R}{L/R + 2\cdot\beta/A} = \frac{1}{1 + 2\cdot\beta\cdot e\cdot R/L} \tag{1.13}$$

where A is substituted with the optimal value $1/e$.

We will see in Section 1.5.2 how IEEE 802.3 LAN, known as Ethernet, uses CSMA/CD.

CSMA/CA

Problems related to this section: Problem 1.20 → Problem 1.23

In wireless LANs, it is not practical to do collision detection because of two main reasons:

1. Implementing a collision detection mechanism would require the implementation of a full duplex radio, capable of transmitting and receiving at once. Unlike wired LANs, where a transmitter can simultaneously monitor the medium for a collision, in wireless LANs the transmitter's power overwhelms a collocated receiver. The dynamic range of the signals on the medium is very large. This is mainly result of the propagation loss, where the signal drops exponentially from its source (recall Figure 1-7!). Thus, a transmitting station cannot effectively distinguish incoming weak signals from noise and the effects of its own transmission.
2. In a wireless environment, we cannot assume that all stations hear each other, which is the basic assumption of the collision detection scheme. Again, due to the propagation loss we have the following problem. The fact that the transmitting station senses the medium free does not necessarily mean that the medium is free around the receiver area. (This is the known as the *hidden station problem*, as described in Figure 1-29.)

As a result, when a station transmits a frame, it has no idea whether the frame collided with another frame until it receives an acknowledgement from the receiver (or times out due to the lack of an acknowledgement). In this situation, collisions have a greater effect on performance than with CSMA/CD, where colliding frames can be quickly detected and aborted while the transmission is in progress. Thus, it makes sense to try to avoid collisions, if possible, and a popular scheme for this is *CSMA/Collision Avoidance*, or CSMA/CA. CSMA/CA is essentially p-persistence, with the twist that when the medium becomes idle, a station must wait for a time period to learn about the fate of the previous transmission before contending for the medium. Figure 1-33 shows sender's state diagram. After a frame was transmitted, the maximum time until a station detects a collision is twice the propagation time of a signal between the stations that are farthest apart plus the detection time. Thus, the station needs at least $2\times\beta$ to ensure that the station is always capable of determining if another station has accessed the medium at the start of the previous slot. The interval between frames (or, packets) needed for the carrier-sense mechanism to determine that the medium is idle and available for transmission is called a **backoff slot**.

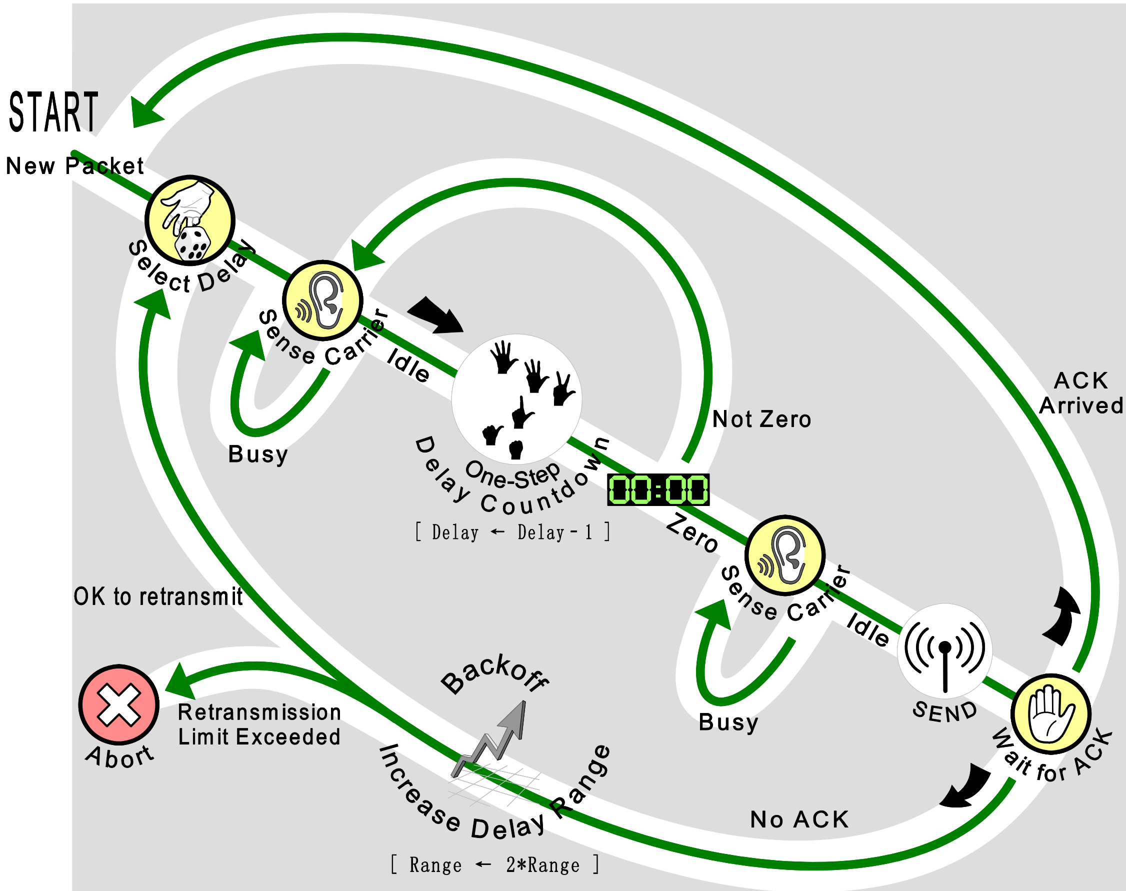

Figure 1-33: The sender's state diagram for CSMA/CA protocol.

When a station wants to transmit data, it first senses the medium whether it is busy. If the medium is busy, the station enters the **access deferral state**. The station continuously senses the medium, waiting for it to become idle. When the medium becomes idle, the station first sets a **contention timer** to a time interval randomly selected in the range [0, *CW*−1], where *CW* is a predefined **contention window** length. Notice that unlike CSMA/CD (Figure 1-30), CSMA/CA station performs carrier sensing after every slot counted down, i.e., it is listening during the contention window (Figure 1-33). In other words, during the backoff procedure, if the station senses the channel as idle for the duration of a backoff slot, the station decrements the counter by one. If the channel is sensed as busy, the station freezes the countdown and waits for the channel to become idle. The station can transmit the frame after it counts down to zero.

After transmitting a frame, the station waits for the receiver to send an ACK. If no ACK is received, the frame is assumed lost to collision, and the source tries again, choosing a contention timer at random from an interval twice as long as the one before (*binary exponential backoff*). The decrementing counter of the timer guarantees that the station will transmit, unlike a p-persistent approach where for every slot the decision of whether or not to transmit is based on a fixed probability p or q_r. Thus regardless of the timer value a station starts at, it always counts down to zero. If the station senses that another station has begun transmission while it was waiting for the expiration of the contention timer, it does not reset its timer, but merely freezes it, and restarts the countdown when the frame completes transmission. In this way, stations that happen to choose a longer timer value get higher priority in the next round of contention.

As it can be seen, CSMA/CA deliberately introduces delay in transmission in order to avoid collision. Avoiding collisions increases the protocol efficiency in terms of the percentage of frames that get successfully transmitted (useful throughput). Notice that efficiency measures only the ratio of the successful transmission to the total number of transmissions. However, it does not specify the delays that result from the deferrals introduced to avoid the collisions. **Error! Reference source not found.** shows the qualitative relationship for the average packet delays, depending on the packet arrival rate.

We will see in Section 1.5.3 how IEEE 802.11 wireless LAN, known as Wi-Fi, uses CSMA/CA.

1.4 Routing and Addressing

In general networks, arbitrary source-destination node pairs communicate via intermediary network nodes. These intermediary nodes are called switches or routers and their main purpose is to bring packets to their destinations. A good routing protocol will also do it in an efficient way, meaning via the shortest path or the path that is in some sense optimal. The data-carrying capacity of the resulting source-to-destination path directly depends on the efficiency of the routing protocol employed.

Bridges, Switches, and Routers

A **packet switch** is a network device with several incoming and outgoing links that forwards packets from incoming to outgoing links. Each attachment to a network is known as a **network interface** or **network port**. When a packet is received by a switch, the appropriate outgoing port is decided based on the packet's guidance information (contained in the packet header).

Two general approaches are used to interconnect multiple networks: *bridges* or *routers*. **Bridges** are simple networking devices that are used for interconnecting local area networks (LANs) that use identical protocols for the physical and link layers of their protocol stack. The terms "bridge" and "switch" are often used synonymously. Because bridged networks use the same protocols, the amount of processing required at the bridge is minimal. There are also more sophisticated bridges, which are capable of mapping from one link-layer format to another. More information on bridges and switches is available in Section 1.5.2.

Routers are general-purpose packet switches that can interconnect arbitrary networks. A **router** has two important functions: (1) *routing*, which is the process of finding and maintaining optimal paths between source and destination nodes; and, (2) *forwarding* (or *switching*), which is the process of relaying incoming data packets along the routing path. A router is a switch that builds its forwarding table by routing algorithms. Routing often searches for the shortest path, which in abstract graphs is a graph distance between the nodes. Shortest path can be determined in different ways, such as:

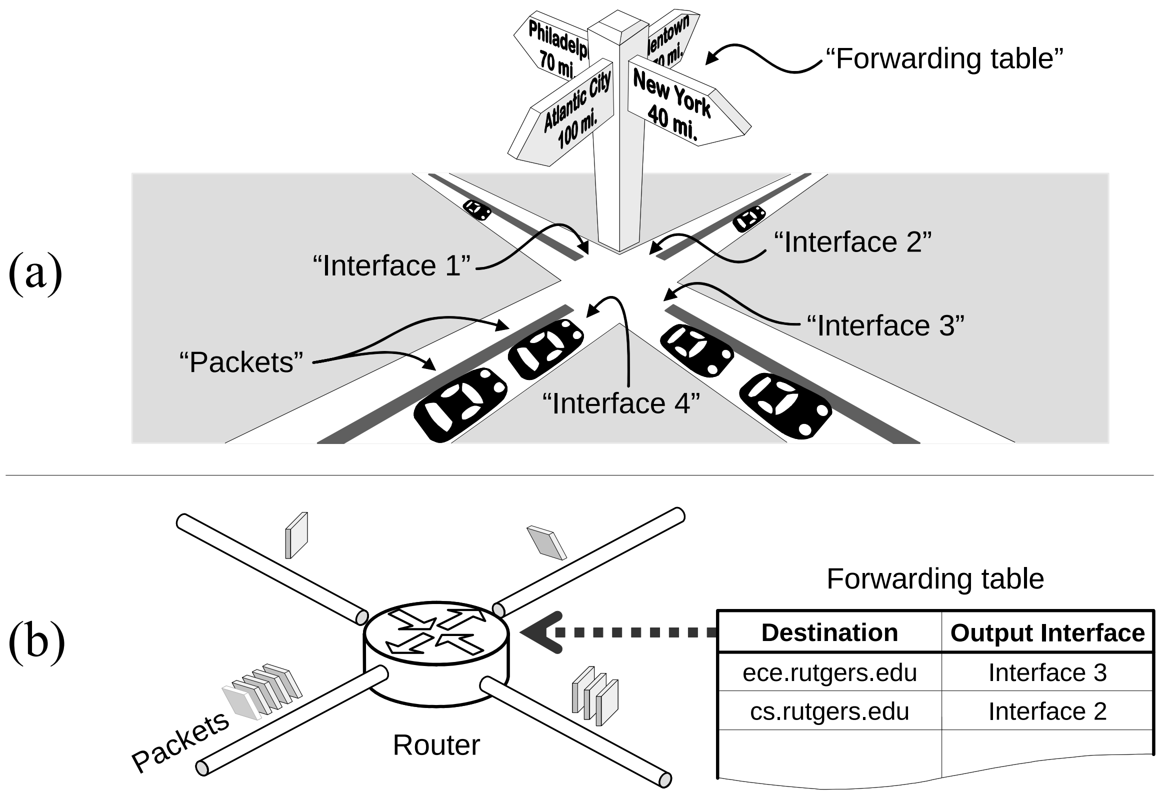

Figure 1-34: A router can be thought of as a crossroads, with connecting points corresponding to the router interfaces. When a car (packet) enters the intersection, it is directed out by looking up the forwarding table.

- Knowing the graph topology, calculate the shortest path
- Send “boomerang” probes on round trips to the destination along the different outgoing paths. Whichever returns back the first is the one that carries the information about the shortest path

Figure 1-34 illustrates an analogy between a crossroads and a router. Similar to a road sign, the router maintains a forwarding table that directs the incoming packets to the appropriate exit interfaces, depending on their final destination. Of course, as the road signs on different road intersections list different information depending on intersection’s location relative to the roadmap, so the routing tables in different routers list different information depending on router’s location relative to the rest of the network.

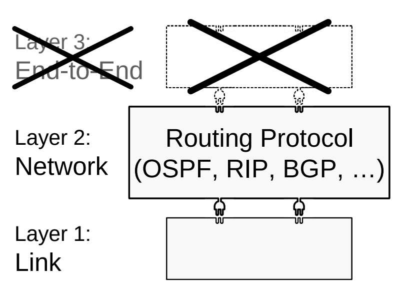

A **router** is a network device that interconnects two or more computer networks, where each network may be using a different link-layer protocol. The two major problems of delivering packets in networks from an arbitrary source to an arbitrary location are:

- How to build the forwarding tables in all network nodes
- How to do forwarding (efficiently)

Usually, a requirement is that the path that a packet takes from a source to a destination should be in some sense *optimal*. There are different optimality metrics, such as quickest, cheapest, or most secure delivery. Later, in Sections 1.4.2 and 1.4.3, we will learn about some algorithms for finding optimal paths, known as *routing algorithms*.

Pseudo code of a routing protocol module is given in Listing 1-2.

Listing 1-2: Pseudo code of a routing protocol module.

```
public class RoutingProtocol extends Thread {
    // specifies how frequently this node advertises its routing info
    public static final int ADVERTISING_PERIOD = 100;

    // link layer protocol that provides services to this protocol
    private ProtocolLinkLayer linkLayerProtocol;

    // associative table of neighboring nodes
    //     (associates their addresses with this node's interface cards)
    private HashMap neighbors;

    // information received from other nodes
    private HashMap othersRoutingInfo;

    // this node's routing table
    private HashMap myRoutingTable;

    // constructor
    public RoutingProtocol(
        ProtocolLinkLayer linkLayerProtocol
    ) {
        this.linkLayerProtocol = linkLayerProtocol;

        populate myRoutingTable with costs to my neighbors;
    }

    // thread method; runs in a continuous loop and sends routing-info advertisements
    //     to all the neighbors of this node
    public void run() {
        while (true) {
            try { Thread.sleep(ADVERTISING_PERIOD); }
            catch (InterruptedException e) {
                for (all neighbors) {
                    Boolean status = linkLayerProtocol.send();
                    // If the link was down, update own routing & forwarding tables
                    //     and send report to the neighbors
                    if (!status) {
                    }
                }
            }
        }
    }

    // upcall method (called from the layer below this one, in a bottom-layer thread!)
    //     the received packet contains an advertisement/report from a neighboring node
    public void handle(byte[] data) throws Exception {
```

```
            // reconstruct the packet as in Listing 1-1 (Section 1.1.4) for a generic handle()
            //          but there is no handover to an upper-layer protocol;
            // update my routing table based on the received report
            synchronized (routingTable) { // critical region
            } // end of the critical region

            // update the forwarding table of the peer forwarding protocol
            //          (note that this protocol is running in a different thread!)
            call the method setReceiver() in Listing 1-1
        }
}
```

The code description is as follows: … to be described …

As will be seen later, routing is not an easy task. Optimal routing requires a detailed and timely view of the network topology and link statuses. However, obtaining such information requires a great deal of periodic messaging between all nodes to notify each other about the network state in their local neighborhoods. This is viewed as overhead because it carries control information and reduces the resources available for carrying user information. The network engineer strives to reduce overhead. In addition, the finite speed of propagating the messages and processing delays in the nodes imply that the nodes always deal with an outdated view of the network state. Therefore, in real-world networks routing protocols always deal with a partial and outdated view of the network state. The lack of perfect knowledge of the network state can lead to a poor behavior of the protocol and/or degraded performance of the applications.

Path MTU is the smallest maximum transmission unit of any link on the current path (also known as route) between two hosts. The concept of MTU is defined in Section 1.1.3.

This section deals mainly with the control functions of routers that include building the routing tables. Later, in Section 4.1 we will consider how routers forward packets from incoming to outgoing links. This process consists of several steps and each step takes time, which introduces delays in packet delivery. Also, due to the limited size of the router memory, some incoming packets may need to be discarded for the lack of memory space. Section 4.1 describes methods to reduce forwarding delays and packet loss due to memory shortage.

1.4.1 Networks, Internets, and the IP Protocol

A *network* is a set of computers directly connected to each other, i.e., with no intermediaries. A network of networks is called *internetwork*.

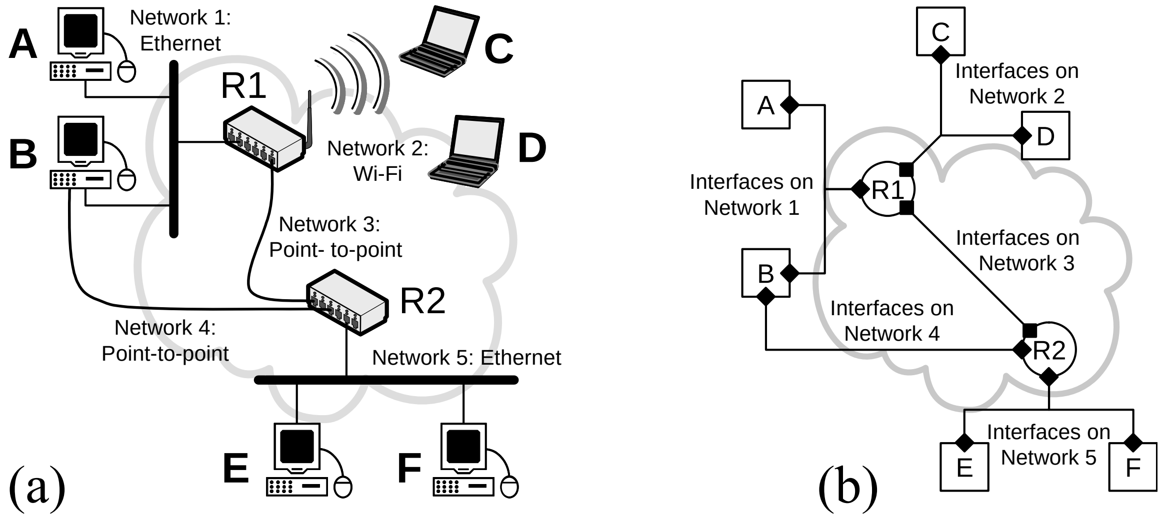

Figure 1-35: Example internetwork: (a) The physical networks include 2 Ethernets, 2 point-to-point links, and 1 Wi-Fi network. (b) Topology of the internetwork and interfaces.

Consider an example internetwork in Figure 1-35(a), which consists of five physical networks interconnected by two routers. The underlying network that a device uses to connect to other devices could be a LAN connection like Ethernet or Token Ring, a wireless LAN link such as 802.11 (known as Wi-Fi) or Bluetooth, or a dialup, DSL, or a T-1 connection. Each physical network will generally use its own frame format, and each format has a limit on how much data can be sent in a single frame (link MTU, Section 1.1.3).

Two types of network nodes are distinguished: hosts vs. routers. Each *host* usually has a single network attachment point, known as **network interface**, and therefore it cannot relay packets for other nodes. Even if a host has two or more network interfaces, such as node B in Figure 1-35(a), it is not intended to be used for transit traffic. Hosts usually do not participate in the routing algorithm. Unlike hosts, *routers* have the primary function of relaying transit traffic from other nodes. Each router has a minimum of two, but usually many more, network interfaces. In Figure 1-35(a), both routers R1 and R2 have three network attachment points (interfaces) each. Each interface on every host and router must have a network address that is globally unique.[5] A node with two or more network interfaces is said to be **multihomed**[6] (or, multiconnected). Notice that multihomed hosts do *not* participate in routing or forwarding of transit traffic. Multihomed hosts act as any other end host, except they may use different interfaces for different destinations, depending on the destination distance.

The whole idea behind a network layer protocol is to implement the concept of a "virtual network" where devices talk even though they are far away, connected using different physical network technologies. This means that the layers above the network layer do not need to worry about details, such as differences in packet formats or size limits of underlying link-layer

---

[5] This is not necessarily true for interfaces that are behind NATs, as discussed later.

[6] Most notebook computers nowadays come with two or more network interfaces, such as Ethernet, Wi-Fi, Bluetooth, etc. However, the host becomes "multihomed" only if two or more interfaces are assigned unique network addresses and they are simultaneously active on their respective physical networks.

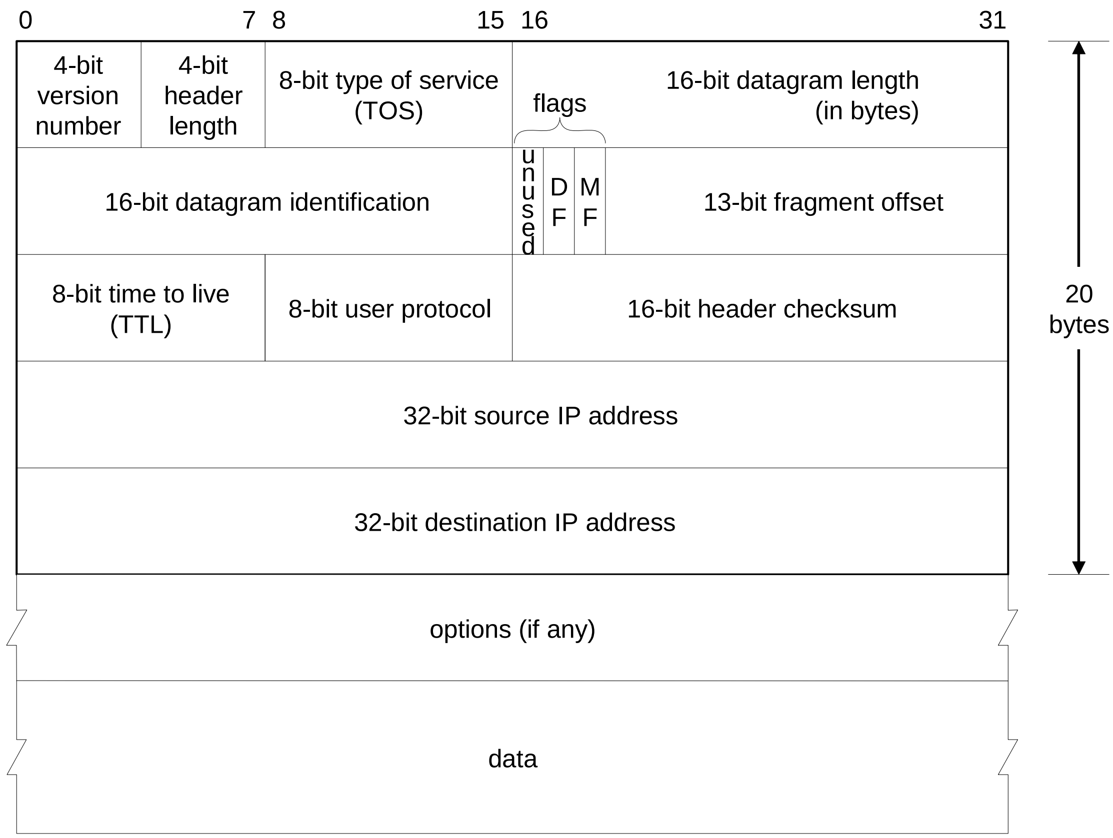

Figure 1-36: The format of IPv4 datagrams.

technologies. The network layer manages these issues seamlessly and presents a uniform interface to the higher layers. The most commonly used network layer protocol is the *Internet Protocol* (IP). The most commonly deployed version of IP is version 4 (IPv4). The next generation, IP version 6 (IPv6),[7] is designed to address the shortcomings of IPv4 and currently there is a great effort in transitioning the Internet to IPv6. IPv6 is reviewed in Section 8.1.

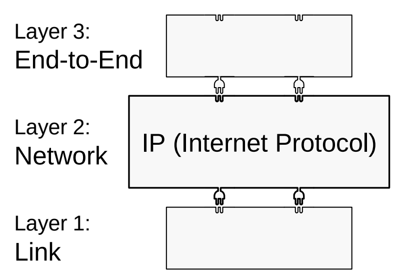

IP Header Format

Data transmitted over an internet using IP is carried in packets called IP *datagrams*. Figure 1-36 shows the format of IP version 4 datagrams. Its fields are as follows:

Version number: This field indicates version number, to allow evolution of the protocol. The value of this field for IPv4 datagrams is 4.

Header length: This field specifies the length of the IP header in 32-bit words. Regular header length is 20 bytes, so the default value of this field equals 5, which is also the minimum allowed

[7] IP version 5 designates the Stream Protocol (SP), a connection-oriented network-layer protocol. IPv5 was an experimental real-time stream protocol that was never widely used.

value. In case the options field is used, the value can be up to $4^2 - 1 = 15$, which means that the options field may contain up to $(15 - 5) \times 4 = 40$ bytes.

Type of service: This field is used to specify the treatment of the datagram in its transmission through component networks. It was designed to carry information about the desired quality of service features, such as prioritized delivery. It was never widely used as originally defined, and its meaning has been subsequently redefined for use by a technique called Differentiated Services (DS), which will be described later in Section 3.3.5.

Datagram length: Total datagram length, including both the header and data, in bytes.

Identification: This is a *sequence number* that, together with the source address, destination address, and user protocol, is intended to identify a datagram uniquely.

Flags: There are three flag bits, of which only two are currently defined. The first bit is reserved and currently unused. The **DF** (Don't Fragment) bit prohibits fragmentation when set. This bit may be useful if it is known that the destination does not have the capability to reassemble fragments. However, if this bit is set and the datagram exceeds the MTU size of the next link, the datagram will be discarded. The **MF** (More Fragments) bit is used to indicate the fragmentation parameters. When this bit is set, it indicates that this datagram is a fragment of an original datagram and this is not its last fragment.

Fragment offset: This field indicates the starting location of this fragment within the original datagram, measured in 8-byte (64-bit) units. This implies that the length of data carried by all fragments before the last one must be a multiple of 8 bytes. The reason for specifying the offset value in units of 8-byte chunks is that only 13 bits are allocated for the offset field, which makes possible to refer to 8,192 locations. On the other hand, the datagram length field of 16 bits allows for datagrams up to 65,536 bytes long. Therefore, to be able to specify any offset value within an arbitrary-size datagram, the offset units are in $65{,}536 \div 8{,}192 = 8$-byte units.

Time to live: The TTL field specifies how long a datagram is allowed to remain in the Internet, to catch packets that are stuck in routing loops. This field was originally set in seconds, and every router that relayed the datagram decreased the TTL value by one. In current practice, a more appropriate name for this field is *hop limit counter* and its default value is usually set to 64.

User protocol: This field identifies the higher-level protocol to which the IP protocol at the destination will deliver the payload. In other words, this field identifies the type of the next header contained in the payload of this datagram (i.e., after the IP header). Example values are 6 for TCP, 17 for UDP, and 1 for ICMP. A complete list is maintained at http://www.iana.org/assignments/protocol-numbers.

Header checksum: This is an error-detecting code applied to the header only. Because some header fields may change during transit (e.g., TTL, fragmentation fields), this field is reverified and recomputed at each router. The checksum is formed by taking the ones complement of the 16-bit ones-complement addition of all 16-bit words in the header. Before the computation, the checksum field is itself initialized to a value of zero.

Source IP address: This address identifies the end host that originated the datagram. Described later in Section 1.4.4.

Destination IP address: This address identifies the end host that is to receive the datagram.

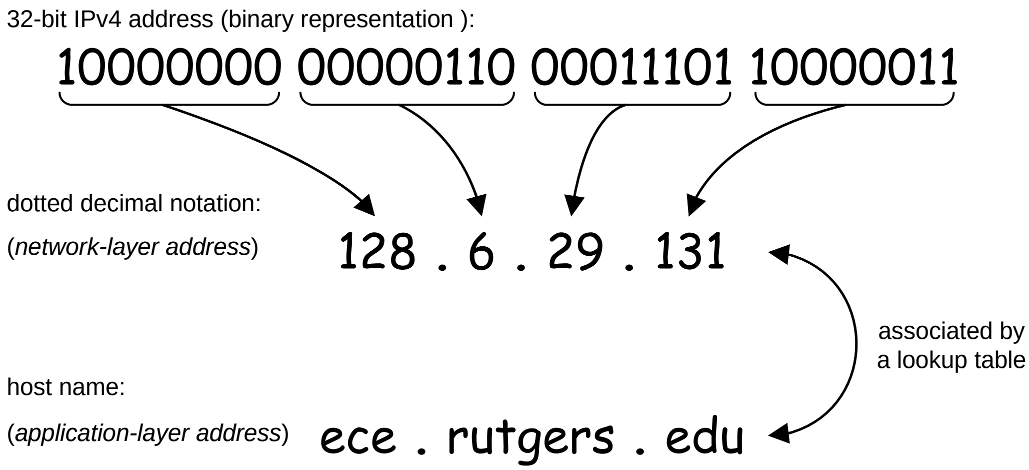

Figure 1-37: Dotted decimal notation for IP version 4 addresses.

Options: This field encodes the options requested by the sending user.

To send messages using the IP protocol, we encapsulate the data from a higher-layer ("user") protocol into IP datagrams. These datagrams must then be sent down to the link-layer protocol, where they are further encapsulated into the frames of whatever technology is going to be used to physically convey them, either directly to their destination, or indirectly to the next intermediate step in their journey to their intended recipient. (The encapsulation process is illustrated in Figure 1-13.) The link-layer protocol puts the entire IP datagram into the data portion (the payload) of its frame format, just as IP puts end-to-end layer messages, end-to-end headers and all, into its IP Data field.

Naming and Addressing

Names and addresses play an important role in all computer systems as well as any other symbolic systems. They are labels assigned to entities such as physical objects or abstract concepts, so those entities can be referred to in a symbolic language. Because computation is specified in and communication uses symbolic language, the importance of names should be clear. It is important to emphasize the importance of naming the network nodes, because if a node is not named, it does not exist! We simply cannot target a message to an unknown entity[8]. The main issues about naming include:

- Names must be *unique* so that different entities are not confused with each other
- Names must be *bound* to and *resolved* with the entities they refer to, to determine the object of computation or communication

It is common in computing and communications to differentiate between names and addresses of objects. Technically, both are addresses (of different kind), but we distinguish them for easier usage. *Names* are usually human-understandable, therefore variable length (potentially rather

[8] Many communication networks allow broadcasting messages to all or many nodes in the network. Hence, in principle the sender could send messages to nodes that it does not know of. However, this is not an efficient way to communicate and it is generally reserved for special purposes.

long) and may not follow a strict format. *Addresses* are intended for machine use, and for efficiency reasons have fixed lengths and follow strict formatting rules. For example, you could name your computers: "My office computer for development-related work" and "My office computer for business correspondence." The addresses of those computers could be: 128.6.236.10 and 128.6.237.188, respectively. Figure 1-37 illustrates the relationship between the binary representation of an IP address, its dotted-decimal notation, and the associated name. One could say that "names" are *application-layer addresses* and "addresses" are *network-layer addresses*. Notice that in dotted-decimal notation the maximum decimal number is 255, which is the maximum number that can be represented with an 8-bit field. The mapping between the names and addresses is performed by the Domain Name System (DNS), described in Section 8.4.

Distinguishing names and addresses is useful for another reason: this separation allows keeping the same name for a computer that needs to be labeled differently when it moves to a different physical place (see Mobile IP in Section 8.3.4). For example, the name of your friend may remain the same in your email address book when he or she moves to a different company and changes their email address. Of course, the name/address separation implies that there should be a mechanism for name-to-address binding and address-to-name resolution.

Two most important address types in contemporary networking are:

- Link-layer address of a device, also known as *medium access control (MAC) address*, which is a physical address for a given *network interface card* (NIC), also known as *network adaptor* or *line card*. These addresses are standardized by the IEEE group in charge of a particular physical-layer communication standard, assigned to different vendors, and hardwired into the physical devices.

- Network-layer address of a device, which is a logical address and can be changed by the end user. This address is commonly referred to as *IP address*, because IP is by far the most common network-layer protocol. Network-layer addresses are standardized by the Internet Engineering Task Force (http://www.ietf.org).

Notice that a quite independent addressing scheme is used for telephone networks and it is governed by the International Telecommunications Union (http://www.itu.int).

People designed postal addresses with a structure that facilitates human memorization and post-service delivery of mail. So, a person's address is structured hierarchically, with country name on top of the hierarchy, followed by the city name, postal code, and the street address. One may wonder whether there is anything to be gained from adopting a similar approach for network computer naming. After all, computers deal equally well with numbers and do not need mnemonic techniques to help with memorization and recall. It turns out that in very large networks, the address structure can assist with more efficient message routing to the destination. Section 1.4.4 describes how IPv4 addresses are structured to assist routing.

Datagram Fragmentation and Reassembly

Problems related to this section: Problem 1.24

The Internet Protocol's main responsibility is to deliver data between devices on different networks, i.e., across an internetwork. For this purpose, the IP layer encapsulates data received

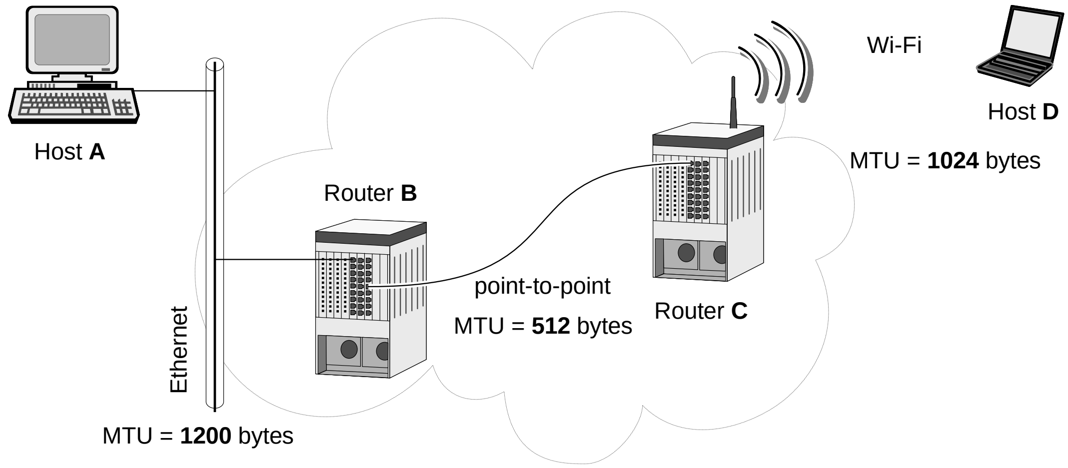

Figure 1-38: Example scenario for IP datagram fragmentation.

from higher layers into IP datagrams for transmission. These datagrams are then passed down to the link layer where they are sent over physical network links.

In Section 1.1.3, we saw that underlying network technology imposes the upper limit on the frame (packet) size, known as maximum transmission unit (MTU). As the datagram is forwarded along the source-destination path, each hop may use a different physical network, with a different maximum underlying frame size. If an IP datagram is larger than the MTU of the underlying network, it may be necessary to break up the datagram into several smaller datagrams. This process is called **fragmentation**. The fragment datagrams are then sent individually and reassembled at the destination into the original datagram.

IP is designed to manage datagram size in a seamless manner. It matches the size of the IP datagram to the size of the underlying link-layer frame size, and performs fragmentation and reassembly so that the upper-layer protocols are not aware of this process. Here is an example:

Example 1.2 Illustration of IP Datagram Fragmentation

In the example scenario shown in Figure 1-38, an application on host *A*, say email client, needs to send a JPEG image to the receiver at host *D*. Assume that the sender uses the TCP protocol (described in Chapter 2), which in turn uses IP as its network-layer protocol. The first physical network is Ethernet (Section 1.5.2), which for illustration is configured to limit the size of the payload it sends to 1,200 bytes. The second network uses a Point-to-Point protocol that limits the payload size 512 bytes and the third network is Wi-Fi (Section 1.5.3) with the payload limit equal to 1024 bytes.

Figure 1-39 illustrates the process by which IP datagrams are fragmented by the source device and possibly routers along the path to the destination. As we will learn in Chapter 2, TCP learns from IP about the MTU of the first link and prepares the TCP packets to fit this limit, so the host's IP layer does not need to perform any fragmentation. However, router *B* needs to break up the datagram into several smaller datagrams to fit the MTU of the point-to-point link. As shown in Figure 1-39, the IP layer at router *B* creates three smaller datagrams from the first datagram it receives from host *A*.

The bottom row in Figure 1-39(b) shows the contents of the fragmentation-related fields of the datagram headers (the second row of the IP header shown in Figure 1-36). Recall that the length of data carried by all fragments before the last one must be a multiple of 8 bytes and the offset values are in units of 8-byte chunks. Because of this constraint, the size of the first two datagrams created by fragmentation on router *B* is 508 bytes (20 bytes for IP header + 488 bytes of IP payload). Although the MTU allows IP datagrams of 512 bytes, this would result in a payload size of 492, which is not a

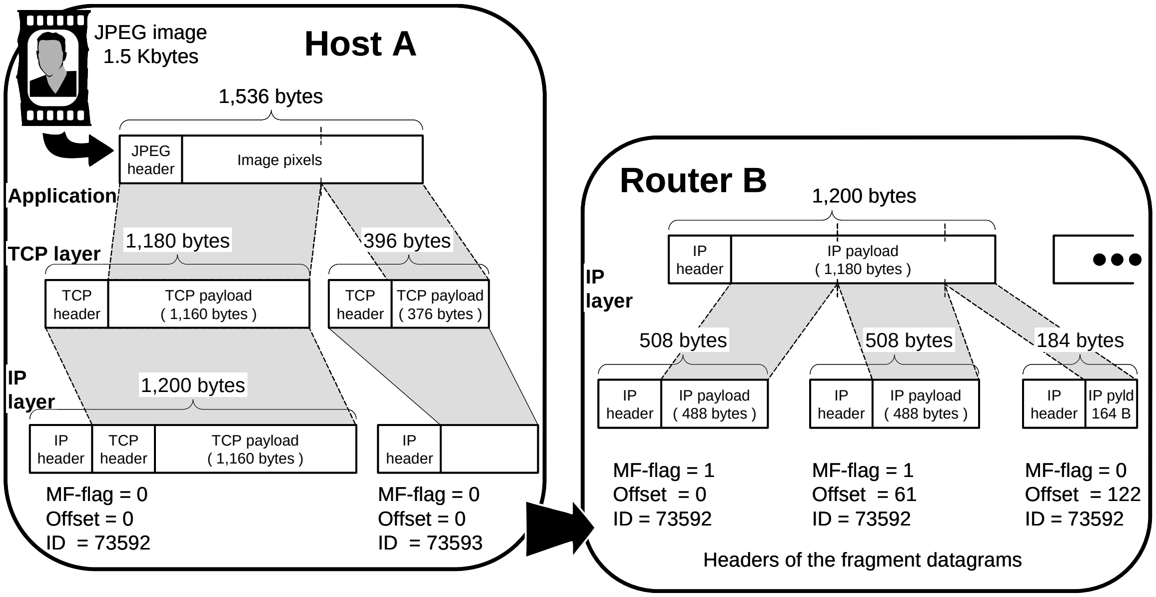

Figure 1-39: IP datagram fragmentation at Router B of the network shown in Figure 1-38. The fragments will be reassembled at the destination (Host D) in an exactly reverse process.

multiple of 8 bytes. Notice also that the offset value of the second fragment is 61, which means that this fragment starts at $8 \times 61 = 488$ bytes in the original IP datagram from which this fragment is created.

It is important to reemphasize that lower layer protocols do not distinguish any structure in the payload passed to them by an upper-layer protocol (Figure 1-13). Therefore, although in the example of Figure 1-39 the payload of IP datagrams contains both TCP header and user data, the IP does *not* distinguish any structure within the datagram payload. When the IP layer on Router *B* receives an IP datagram with IP header (20 bytes) + 1,180 bytes of payload, it removes the IP header and does not care what is in the payload. Router *B*'s IP layer splits the 1,180 bytes into three fragments, so that when it adds its own IP header in front of each payload fragment, none of the resulting IP datagrams will exceed 512 bytes in size. Router *B* then forwards the three datagrams to the next hop.

1.4.2 Link State Routing

Problems related to this section: Problem 1.25 → ?

A key problem of routing algorithms is finding the *shortest path* between any two nodes, such that the sum of the costs of the links constituting the path is minimized. The two most popular algorithms used for this purpose are Dijkstra's algorithm, used in link state routing, and Bellman-Ford algorithm, used in distance vector routing. The link state routing is presented first, followed by the distance vector routing; Section 1.4.5 describes the path vector routing, which is similar to the distance vector routing.

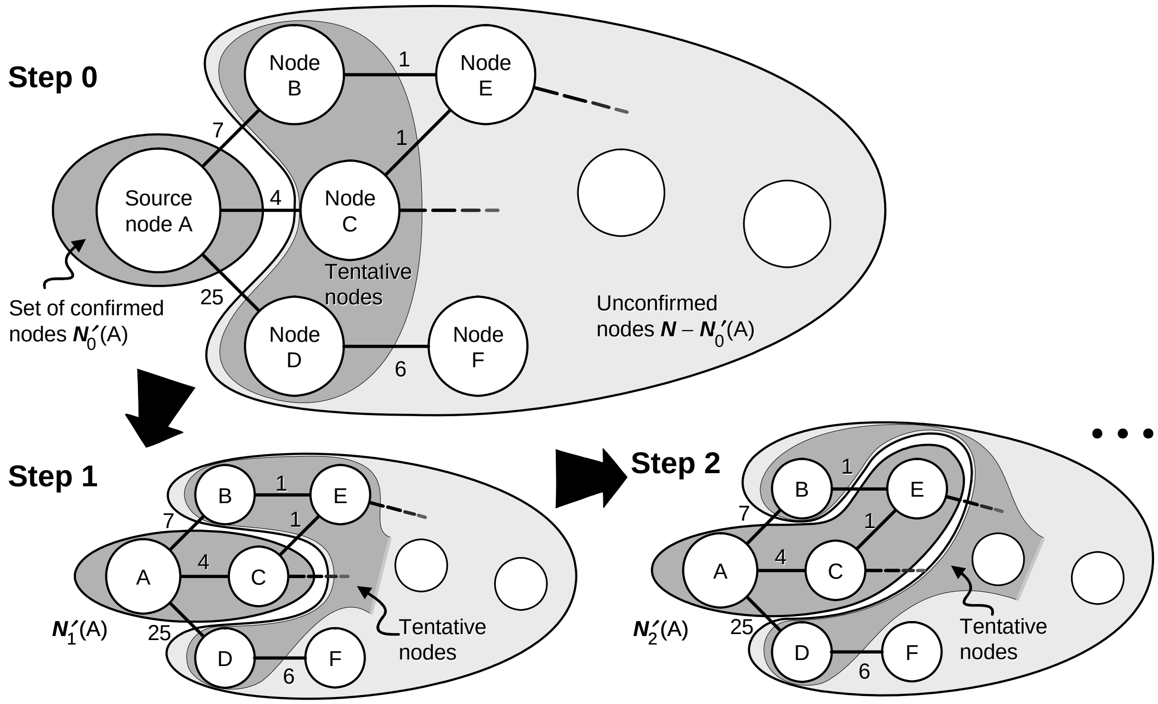

Figure 1-40: Illustration of finding the shortest path using Dijkstra's algorithm.

The key idea of the link state routing algorithm is to disseminate the information about local connectivity of each node to all other nodes in the network. Once all nodes gather the local information from all other nodes, each node knows the topology of the entire network and can independently compute the shortest path from itself to any other node in the network. This is done by iteratively identifying the closest node from the source node in the order of increasing path cost (Figure 1-40). At the k^{th} step we have the set $\boldsymbol{N'_k}(A)$ of k closest nodes to node A ("confirmed nodes") as well as the shortest distance D_X from each node X in $\boldsymbol{N'_k}(A)$ to node A. Of all paths connecting some node not in $\boldsymbol{N'_k}(A)$ ("unconfirmed nodes") with node A, there is the shortest one that passes exclusively through nodes in $\boldsymbol{N'_k}(A)$, because $c(X, Y) \geq 0$. Therefore, the $(k + 1)$st closest node should be selected among those unconfirmed nodes that are neighbors of nodes in $\boldsymbol{N'_k}(A)$. These nodes are marked as "tentative nodes" in Figure 1-40.

When a router (network node A) is initialized, it determines the link cost on each of its network interfaces. For example, in Figure 1-40 the cost of the link connecting node A to node B is labeled as "7" units, that is $c(A, B) = 7$. The node then advertises this set of link costs to *all* other nodes in the network (not just its neighboring nodes). Each node receives the link costs of all nodes in the network and, therefore, each node has a representation of the entire network. To advertise the link costs, the node creates a packet, known as *Link-State Advertisement* (LSA) or *Link-State Packet* (LSP), which contains the following information:

- The ID of the node that created the LSA
- A list of directly connected neighbors of this node, with the link cost to each one
- A sequence number for this packet
- A time-to-live for this packet

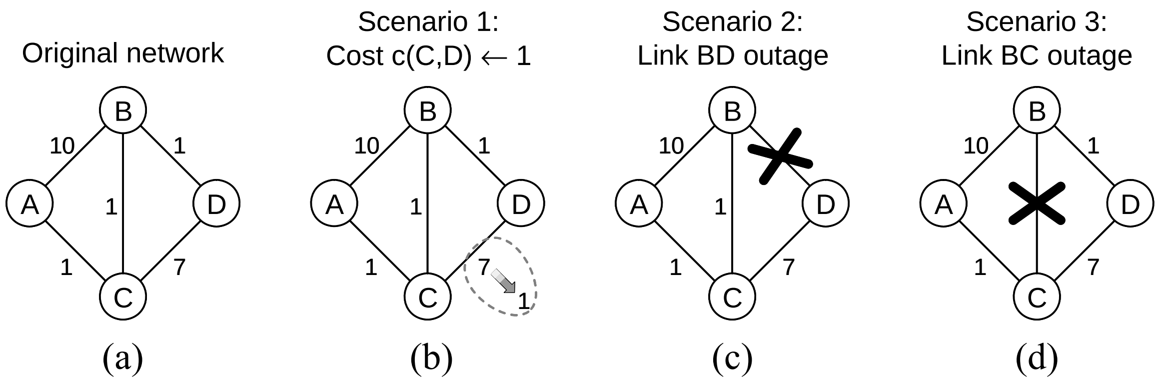

Figure 1-41: Example network used for illustrating the routing algorithms.

In the initial step, all nodes send their LSAs to all other nodes in the network using the mechanism called *broadcasting*. The shortest-path algorithm, which is described next, starts with the assumption that all nodes already exchanged their LSAs. The next step is to build a routing table, which is an intermediate step towards building a forwarding table. A **routing table** of a node (source) contains the paths and distances to all other nodes (destinations) in the network. A **forwarding table** of a node pairs different destination nodes with appropriate output interfaces of this node (recall Figure 1-34(b)).

The *link state routing algorithm* works as follows. Let $\boldsymbol{N}$ denote the set of all nodes in a network. In Figure 1-41, $\boldsymbol{N} = \{A, B, C, D\}$. The process can be summarized as an iterative execution of the following steps

1. Check the LSAs of all nodes in the confirmed set $\boldsymbol{N}'$ to update the tentative set (recall that tentative nodes are unconfirmed nodes that are neighbors of confirmed nodes)
2. Move the tentative node with the shortest path to the confirmed set $\boldsymbol{N}'$.
3. Go to Step 1.

The process stops when $\boldsymbol{N}' = \boldsymbol{N}$. Here is an example:

Example 1.3 Link State Routing Algorithm

Consider the network in Figure 1-41(a) and assume that it uses the link state routing algorithm. Starting from the initial state for all nodes, show how node *A* finds the shortest paths to all other nodes in the network. The figure below shows how node *A*'s link-state advertisement (LSA) is broadcast through the network.

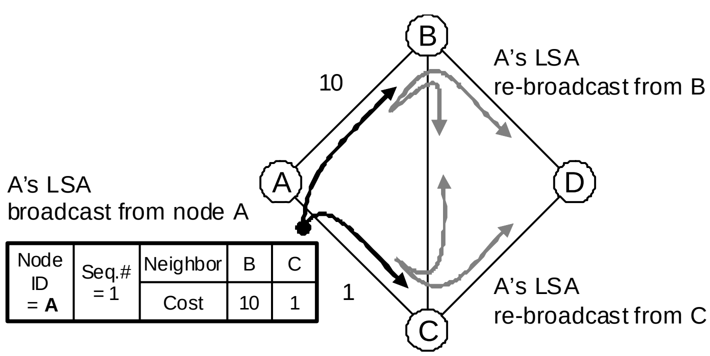

Assume that all nodes broadcast their LSAs and each node already received LSAs from all other nodes in the network before it starts the shortest path computation, as shown in this figure:

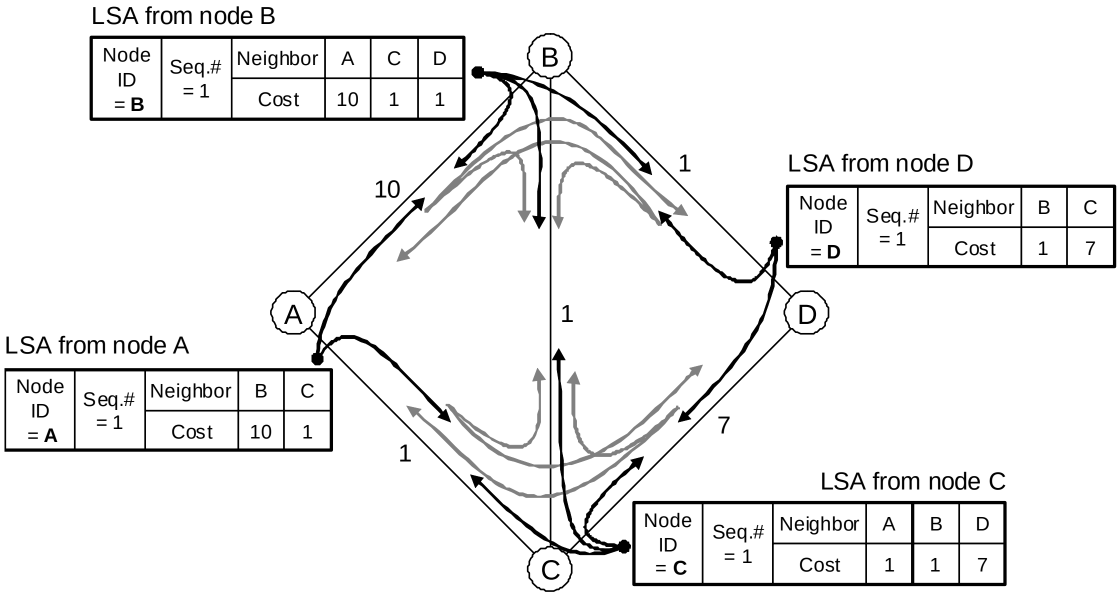

Table 1-2 shows the process of building a routing table at node *A* of the network shown in Figure 1-41(a). Each node is represented with a triplet (*Destination node ID*, *Path length*, *Next hop*). The node *x* maintains two sets (recall Figure 1-40): Confirmed(*x*) set, denoted as ***N*′**, and Tentative(*x*) set. At the end, the routing table in node *A* contains these entries: {(*A*, 0, –), (*C*, 1, *C*), (*B*, 2, *C*), (*D*, 3, *C*)}. Every other node in the network runs the same algorithm to compute its own routing table.

To account for failures of network elements, the nodes should repeat the whole procedure periodically. That is, each node periodically broadcasts its LSA to all other nodes and recomputes

Table 1-2: Steps for building a routing table at node *A* in Figure 1-41. Each node is represented with a triplet (*Destination node ID, Path length, Next hop*).

Step	Confirmed set *N*′	Tentative set	Comments
0	(*A*, 0, –)	∅	Initially, *A* is the only member of Confirmed(*A*), so examine *A*’s LSA.
1	(*A*, 0, –)	(*B*, 10, *B*), (*C*, 1, *C*)	*A*’s LSA says that *B* and *C* are reachable at costs 10 and 1, respectively. Since these are currently the lowest known costs, put on Tentative(*A*) list.
2	(*A*, 0, –), (*C*, 1, *C*)	(*B*, 10, *B*)	Move lowest-cost member (*C*) of Tentative(*A*) into Confirmed set. Next, examine LSA of newly confirmed member *C*.
3	(*A*, 0, –), (*C*, 1, *C*)	(*B*, 2, *C*), (*D*, 8, *C*)	Cost to reach *B* through *C* is 1+1=2, so replace (*B*, 10, *B*). *C*’s LSA also says that *D* is reachable at cost 7+1=8.
4	(*A*, 0, –), (*C*, 1, *C*), (*B*, 2, *C*)	(*D*, 8, *C*)	Move lowest-cost member (*B*) of Tentative(*A*) into Confirmed, then look at *B*’s LSA.
5	(*A*, 0, –), (*C*, 1, *C*), (*B*, 2, *C*)	(*D*, 3, *C*)	Because *D* is reachable via *B* at cost 1+1+1=3, replace the Tentative(*A*) entry for *D*.
6	(*A*, 0, –), (*C*, 1, *C*), (*B*, 2, *C*), (*D*, 3, *C*)	∅	Move lowest-cost member (*D*) of Tentative(*A*) into Confirmed. END.

its routing table based on the received LSAs.

Limitations: Routing Loops

Link state routing needs large amount of resources to calculate routing tables. It also creates heavy traffic because of flooding the LSA packets from each node throughout the network.

On the other hand, link state routing converges much faster to correct values after link failures than distance vector routing (described in Section 1.4.3), which suffers from the so-called counting-to-infinity problem.

Before the nodes start their routing table computation (as in Table 1-2), they all must have received the same LSAs from all other nodes in the network. If not all of the nodes are working from *exactly* the same map, routing loops can form. A **routing loop** is a subset of network nodes configured so that data packets may wander aimlessly in the network, making no progress towards their destination, and causing traffic congestion for all other packets. In the simplest form of a routing loop, two neighboring nodes each think the other is the best next hop to a given destination. Any packet headed to that destination arriving at either node will loop between these two nodes. Routing loops involving more than two nodes are also possible.

The reason for routing loops formation is simple: because each node computes its shortest-path tree and its routing table without interacting in any way with any other nodes, then if two nodes start with different maps, it is easy to have scenarios in which routing loops are created.

The most popular practical implementation of link-state routing is *Open Shortest Path First* (OSPF) protocol, reviewed in Section 8.2.2.

1.4.3 Distance Vector Routing

Problems related to this section: Problem 1.27 → Problem 1.29

The key idea of the distance vector routing algorithm is that each node assumes that its neighbors already know the shortest path to each destination node. The node then selects the neighbor for which the overall distance (from the source node to its neighbor, plus from the neighbor to the destination) is minimal. The process is repeated iteratively until all nodes settle to a stable solution. This algorithm is also known by the names of its inventors as *Bellman-Ford algorithm*. Figure 1-42 illustrates the process of finding the shortest path in a network using Bellman-Ford algorithm. The straight lines indicate single links connecting the neighboring nodes. The wiggly lines indicate the shortest paths between the two end nodes (other nodes along these paths are not shown). The bold line indicates the overall shortest path from source to destination.

Next, we describe the *distance vector routing algorithm*. Let $\mathbf{N}$ denote the set of all nodes in a network. In Figure 1-41, $\mathbf{N} = \{A, B, C, D\}$. The two types of quantities that this algorithm uses are:

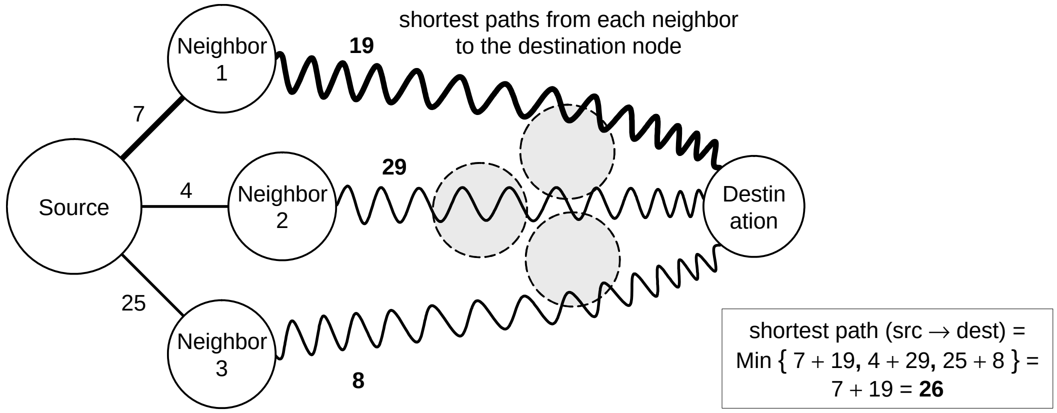

Figure 1-42: Illustration of finding the shortest path using Bellman-Ford algorithm. The thick line (crossing Neighbor 1) represents the shortest path from Source to Destination.

(i) *Link cost* assigned to an individual link directly connecting a pair of nodes (routers). Link costs are given to the algorithm either by having the network operator manually enter the cost values or by having an independent program determine these costs. For example, in Figure 1-41 the cost of the link connecting the nodes A and B is labeled as "10" units, that is $c(A, B) = 10$.

(ii) *Node distance* for an arbitrary pair of nodes, which represents the lowest sum of link costs for all links along all the possible paths between this node pair. The distance from node X to node Y is denoted as $D_X(Y)$. These will be computed by the routing algorithm.

The **distance vector** of node X is the vector of distances from node X to all other nodes in the network, denoted as $DV(X) = \{D_X(Y);\ Y \in \boldsymbol{N}\}$. When determining the minimum-cost path (i.e., distance), it is important to keep in mind that we are not interested in how people would solve this problem. Rather, we wish to know how a group of computers can solve such a problem. Computers (routers) cannot rely on what we people see by looking at the network's graphical representation; computers must work only with the information exchanged in messages.

Let $\eta(X)$ symbolize the set of neighboring nodes of node X. For example, in Figure 1-41 $\eta(A) = \{B, C\}$ because B and C are the only nodes directly linked to node A. The distance vector routing algorithm runs at every node X and calculates the distance to every other node $Y \in \boldsymbol{N}$, $Y \neq X$, using the following formula:

$$D_X(Y) = \min_{V \in \eta(X)} \{c(X,V) + D_V(Y)\} \tag{1.14}$$

To apply this formula, every node must receive the distance vector from all other nodes in the network. Every node maintains a *table of distance vectors*, which includes its own distance vector and distance vectors of its neighbors. Initially, the node assumes that the distance vectors of its neighbors are filled with infinite elements. Here is an example:

Example 1.4 Distributed Distance Vector Routing Algorithm

Consider the original network in Figure 1-41(a) and assume that it uses the distributed distance vector routing algorithm. Starting from the initial state for all nodes, show the first few steps until the routing algorithm reaches a stable state.

For node A, the routing table initially looks as follows.

Routing table at node *A*: Initial

Distance to

From	A	B	C
A	0	10	1
B	∞	∞	∞
C	∞	∞	∞

+

Received Distance Vectors

	A	B	C	D
From B	10	0	1	1

	A	B	C	D
From C	1	1	0	7

=

Routing table at node *A*: After 1st exchange

Distance to

From	A	B	C	D
A	0	2	1	8
B	10	0	1	1
C	1	1	0	7

Notice that node A only keeps the distance vectors of its immediate neighbors, B and C, and not that of any other nodes, such as D. Initially, A may not even know that D exists. Next, each node sends its distance vector to its immediate neighbors and, as a result, A receives distance vectors from B and C. For the sake of simplicity, let us assume that at every node all distance vector packets arrive simultaneously. Of course, this is not the case in reality, but asynchronous arrivals of routing packets do not affect the algorithm operation. When a node receives an updated distance vector from its neighbor, the node overwrites the neighbor's old distance vector in its routing table with the new one. As shown in the figure above, A overwrites the initial distance vectors for B and C. In addition, A re-computes its own distance vector according to Eq. (1.14), as follows:

$$D_A(B) = \min\{c(A,B) + D_B(B),\ c(A,C) + D_C(B)\} = \min\{10+0,\ 1+1\} = 2$$

$$D_A(C) = \min\{c(A,B) + D_B(C),\ c(A,C) + D_C(C)\} = \min\{10+1,\ 1+0\} = 1$$

$$D_A(D) = \min\{c(A,B) + D_B(D),\ c(A,C) + D_C(D)\} = \min\{10+1,\ 1+7\} = 8$$

The new values for A's distance vector are shown in the rightmost table in the above figure.

Similar computations will take place on all other nodes and the whole process is illustrated in Figure 1-43. The end result is as shown in Figure 1-43 as the second column entitled "After 1st exchange." Because for every node the newly computed distance vector is different from the previous one, Figure 1-43 shows that each node sends its new distance vector to its immediate neighbors. The cycle repeats for every node until there is no difference between the new and the previous distance vector. As shown in Figure 1-43, this happens after three exchanges.

A distance-vector routing protocol requires that each router informs its neighbors of topology changes *periodically* and, in some cases, when a change is detected in the topology of a network (*triggered updates*). Routers can detect link failures by periodically testing their links with "heartbeat" or HELLO packets. However, if the router crashes, then it has no way of notifying neighbors of a change. Therefore, distance vector protocols must make some provision for *timing out routes* when periodic routing updates are missing for the last few update cycles.

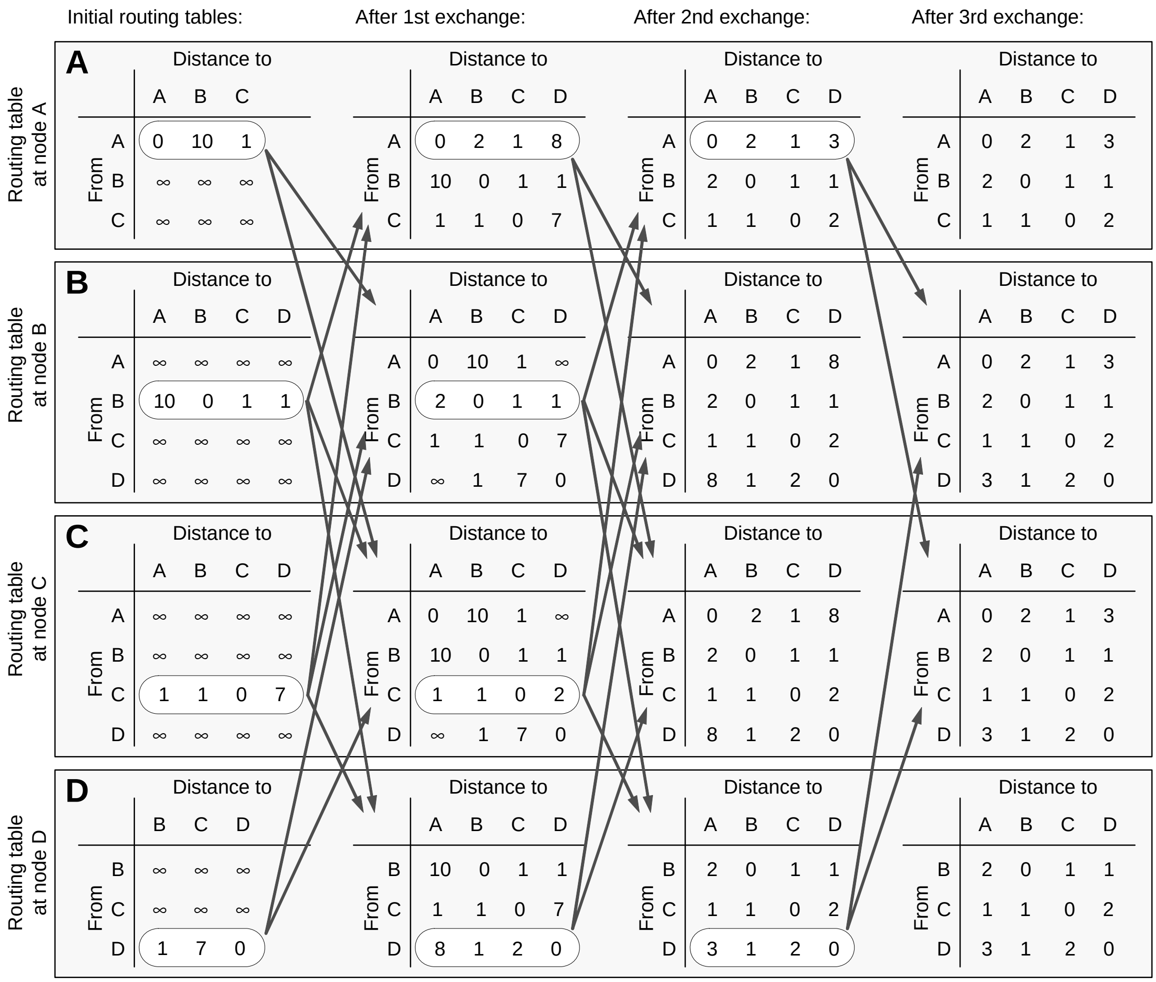

Figure 1-43: Distance vector (DV) algorithm for the original network in Figure 1-41.

Compared to link-state protocols, which require a router to inform all the other nodes in its network about topology changes, distance-vector routing protocols have less computational complexity and message overhead (because each node informs only its own neighbors).

Limitations: Routing Loops and Counting-to-Infinity

Distance vector routing works well if nodes and links are always up, but it suffers from several problems when links fail and become restored. The problems happen because the node does not reveal the information it used to compute its distance vector when it distributes the vector to the neighbors. As a result, remote routers do not have sufficient information to determine whether their choice of the next hop will cause routing loops to form. Although reports about lowering link costs (good news) are adopted quickly, reports about increased link costs (bad news) only spread in slow increments. This problem is known as the “counting-to-infinity problem.”

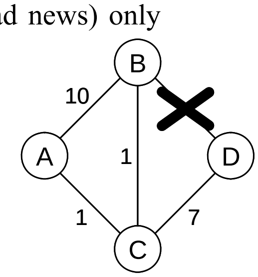

Consider Scenario 2 in Figure 1-41(c) reproduced here in the figure on the right, where after the network stabilizes, the link *BD* fails. Before the failure, the distance vector of the node *B* will be as shown in the figure below. After *B* detects the link failure, it sets its own distance to D as ∞. (Notice that B cannot use old distance vectors it obtained earlier from its neighbors to recompute its new distance vector, because it does not know if they are valid anymore.) If B sends immediately its new distance vector to C,[9] C would figure out that D is unreachable, because its previous best path led via B and now it became unavailable. However, it may happen that C just sent its periodic update (unchanged from before the link BD failure) to B and B receives it after discovering the failure of BD but before sending out its own update. Node B then recomputes its new distance to node D as

$$D_B(D) = \min\{c(B,A) + D_A(D),\ c(B,C) + D_C(D)\} = \min\{10+3,\ 1+2\} = 3$$

and B choses C as the next hop to D. Because we humans can see the entire network topology, we know that C has the distance to D equal to 2 going via the link BC followed by the link CD. However, because C received from B only the numeric values of the distances, not the paths over which these distances are computed, C does *not* know that itself lays on B's shortest path to D!

Routing table at node *B* **before** *BD* outage

Distance to

From	A	B	C	D
A	0	2	1	3
B	2	0	1	1
C	1	1	0	2
D	3	1	2	0

1. *B* detects *BD* outage
2. *B* sets *c*(*B*, *D*) = ∞
3. *B* recomputes its distance vector
4. *B* obtains 3 as the shortest distance to *D*, via *C*

Routing table at node *B* **after** *BD* outage

Distance to

From	A	B	C	D
A	0	2	1	3
B	2	0	1	**3**
C	1	1	0	2
D	3	1	2	0

Given the above routing table, when B receives a data packet destined to D, it will forward the packet to C. However, C will return the packet back to B because for C, B is the next hop on the shortest path from C to D. The packet will bounce back and forth between these two nodes forever (or until their forwarding tables are changed). This phenomenon is called a **routing loop**, because packets may wander aimlessly in the network, making no progress towards their destination.

Because B's distance vector has changed, it reports its new distance vector to its neighbors A and C (triggered update). After receiving B's new distance vector, C will determine that its new shortest path to D measures 4, via C. Now, because C's distance vector changed, it reports its new distance vector to its neighbors, including B. The node B now recomputes its new distance vector and finds that the shortest path to D measures 5, via C. B and C keep reporting the changes until they realize that the shortest path to D is via C because C still has a functioning link to D with the cost equal to 7. This process of incremental convergence towards the correct distance is very slow compared to other route updates, causing the whole network not noticing a router or link outage for a long time, and was therefore named **counting-to-infinity problem**.

A simple solution to the counting-to-infinity problem is known as **hold-down timers**. When a node detects a link failure, it reports to its neighboring nodes that an attached network has gone down. The neighbors immediately start their hold-down timers to ensure that this route will not be mistakenly reinstated by an advertisement received from another router that has not yet learned

[9] Node B will also notify its neighbor A, but for the moment we ignore A because A's path to D goes via C, and on, via B. Hence, A will not be directly affected by this situation.

about this route being unavailable. Until the timer elapses, the router ignores updates regarding this route. Router accepts and reinstates the invalid route if it receives a new update with a better metric than its own or the hold-down timer has expired. At that point, the network is marked as reachable again and the routing table is updated. Typically, the hold-down timer is greater than the total convergence time, providing time for accurate information to be learned, consolidated, and propagated through the network by all routers.

Another solution to the counting-to-infinity problem is known as **split-horizon routing**. The key idea is that it is never useful to send information about a route back in the direction from which it came. Therefore, a router never advertises the cost of a destination to its neighbor *N*, if *N* is the next hop to that destination. The split-horizon rule helps prevent two-node routing loops. In the above example, without split horizons, *C* continues to inform *B* that it can get to *D*, but it does not say that the path goes through *B* itself. Because *B* does not have sufficient intelligence, it picks up *C*'s route as an alternative to its failed direct connection, causing a routing loop. Conversely, with split horizons, *C never* advertises the cost of reaching *D* to *B*, because *B* is *C*'s next hop to *D*. Although hold-downs should prevent counting-to-infinity and routing loops, split horizon provides extra algorithm stability.

An improvement of split-horizon routing is known as **split horizon with poisoned reverse**. Here, the router advertises its full distance vector to all neighbors. However, if a neighbor is the next hop to a given destination, then the router replaces its actual distance value with an infinite cost (meaning "destination unreachable"). In a sense, a route is "poisoned" when a router marks a route as unreachable (infinite distance). Routers receiving this advertisement assume the destination network is unreachable, causing them to look for an alternative route or remove this destination from their routing tables. In the above example, *C* would *always* advertise the cost of reaching *D* to *B* as equal to ∞, because *B* is *C*'s next hop to *D*.

In a single-path internetwork (chain-of-links configuration), split horizon with poisoned reverse has no benefit beyond split horizon. However, in a multipath internetwork, split horizon with poisoned reverse greatly reduces counting-to-infinity and routing loops. The idea is that increases in routing metrics generally indicate routing loops. Poisoned reverse updates are then sent to remove the route and place it in hold-down. Counting-to-infinity can still occur in a multipath internetwork because routes to networks can be learned from multiple sources. None of the above methods works well in general cases. The core problem is that when *X* tells *Y* that it has a path to somewhere, *Y* has no way of knowing whether it itself is on the path.

The most popular practical implementation of link-state routing is *Routing Information Protocol* (RIP), reviewed in Section 8.2.1.

1.4.4 IPv4 Address Structure and CIDR

Problems related to this section: Problem 1.31 → Problem 1.33

Section 1.4.1 briefly mentions that the structure of network addresses should be designet to assist with message routing along the path to the destination. This may not be obvious at first, so let us consider again the analogy between a router and a crossroads (Figure 1-34). Suppose you are driving from Philadelphia to Bloomfield, New Jersey (Figure 1-44). If the sign on the road intersection contained all small towns in all directions, you can imagine that it would be very

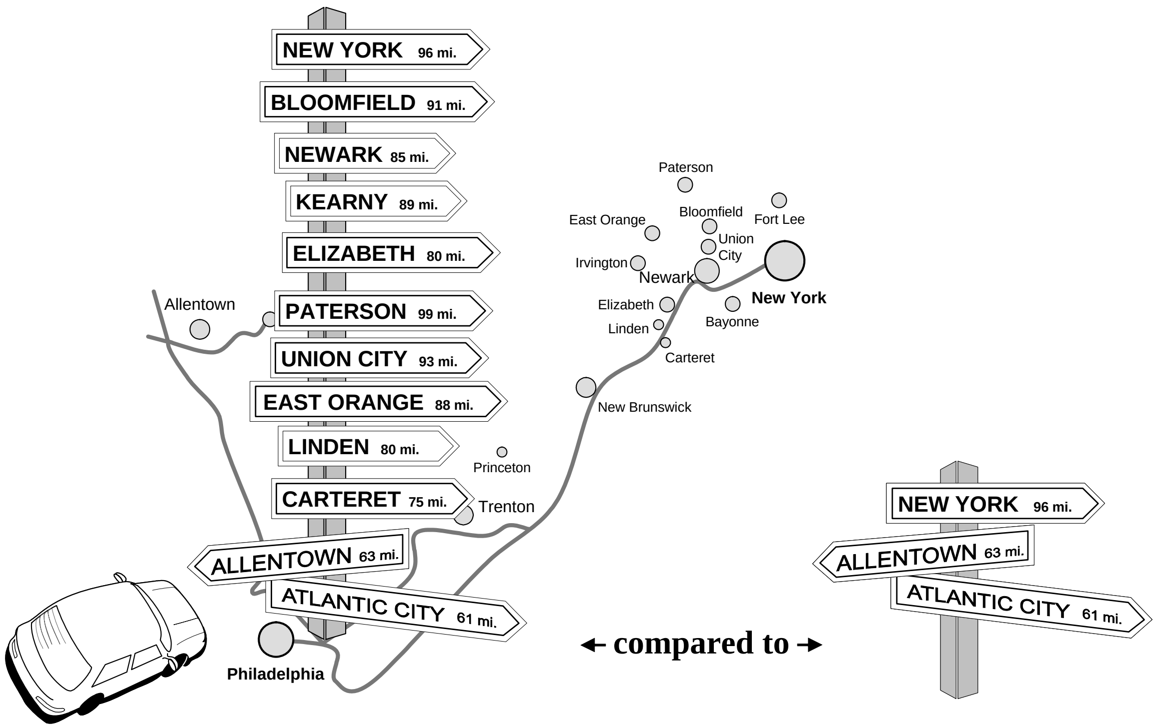

Figure 1-44: Illustration of the problem with the forwarding-table size. Real world road signs contain only a few destinations (lower right corner) to keep them manageable.

difficult to build and use such "forwarding tables." The intersections would be congested by cars looking-up the long table and trying to figure out which way to exit out of the intersection. The problem is solved by listing only the major city names on the signs. Notice that in this case "New York" represents the *entire region* around the city, including all small towns in the region. That is, you do not need to pass through New York to reach Bloomfield, New Jersey. On your way, as you are approaching New York, at some point there will be another crossroads with a sign for Bloomfield. Therefore, *hierarchical address structure* gives a hint about the location that can be used to simplify routing.

Large computer networks, such as the Internet, encounter a similar problem with building and using forwarding tables. The solution has been to divide the network address into two parts: a fixed-length "region" portion (in the most significant bits) and an "intra-region" address. These two parts combined represent the actual network address. In this model, forwarding is simple: The router first looks at the "region" part of the destination address; if it sees a packet with the destination address *not* in this router's region, it does a lookup on the "region" portion of the address and forwards the packet onwards. Conversely, if the destination address *is* in this router's region, it does a lookup on the "intra-region" portion and forwards the packet on. This structuring of network-layer addresses dramatically reduces the size of the forwarding tables. The data in the forwarding table for routes outside the router's region is at most equal to the number of regions in the entire network, typically much smaller than the total number of possible addresses.

The idea of hierarchical structuring can be extended to a multi-level hierarchy, starting with individual nodes at level 0 and covering increasingly larger regions at higher levels of the addressing hierarchy. In such a network, as a packet approaches its destination, it would be forwarded more and more precisely until it reaches the destination node. The key issues in designing such hierarchical structure for network addresses include:

- Should the hierarchy be *uniform*, for example so that a region at level i+1 contains twice as many addresses as a region at level i. In other words, what is the best granularity for quantizing the address space at different levels, and should the hierarchy follow a regular or irregular pattern?
- Should the hierarchy be *statically defined* or could it be *dynamically adaptive*? In other words, should every organization be placed at the same level regardless of how many network nodes it manages? If different-size organizations are assigned to different levels, what happens if an organization outgrows its original level or merges with another organization? Should organization's hierarchy (number of levels and nodes per level) remain forever fixed once it is designed?

The original solution for structuring IPv4 addresses (standardized with RFC-791 in 1981) decided to follow a uniform pattern for structuring the network addresses and opted for a statically defined hierarchy. IPv4 addresses were standardized to be 32-bits long, which gives a total of 2^{32} = 4,294,967,296 possible network addresses. At that time, the addresses were grouped into four classes, each class covering different number of addresses. In computer networks, "regions" correspond to sub-networks, or simply networks, within an internetwork (Section 1.4.1). Depending on the class, the first several bits correspond to the "network" identifier and the remaining bits to the "host" identifier (Figure 1-45). Class A addresses start with a binary "0" and have the next 7 bits for network number and the last 24 bits for host number. Class B addresses start with binary "10", use the next 14 bits for network number, and the last 16 bits for host number (e.g., Rutgers has a Class B network, with addresses in dotted-decimal notation of the form `128.6.*`). Class C addresses start with binary "110" and have the next 21 bits for network number, and the last 8 bits for host number. A special class of addresses is Class D, which are used for IP multicast (described in Section 3.3.2). They start with binary "1110" and use the next 28 bits for the group address. Multicast routing is described later in Section 3.3.2. Addresses that start with binary "1111" are reserved for experiments.

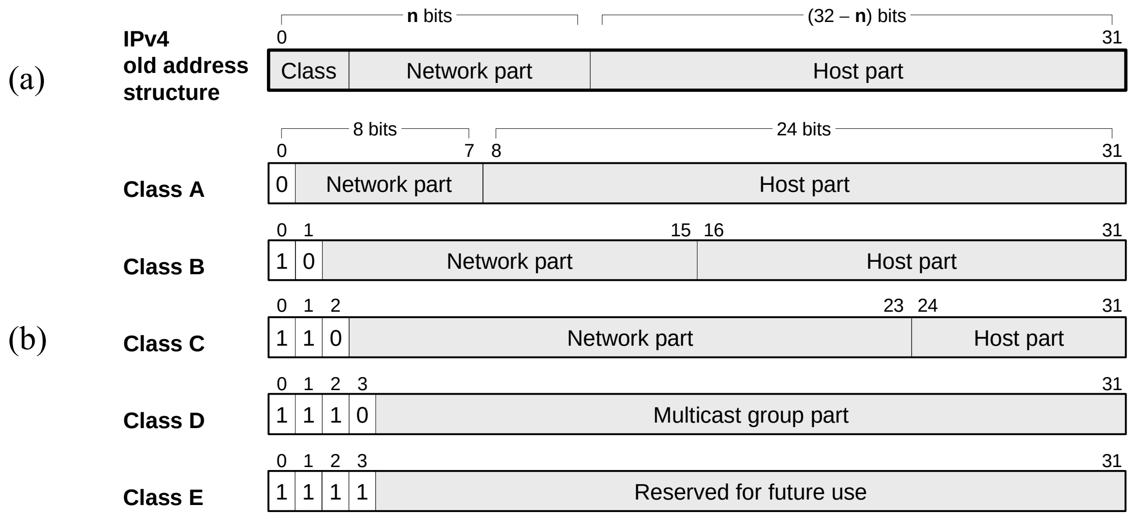

Figure 1-45: (a) Class-based structure of IPv4 addresses (deprecated). (b) The structure of the individual address classes.

The router-forwarding task in IPv4 is a bit more complicated than for an unstructured addressing scheme that requires an exact-match. For every received packet, the router examines its destination address and determines whether it belongs to the same region as this router's addresses. If so, it looks for an exact match; otherwise, it performs a fixed-length lookup depending on the class.

In the original design of IPv4, address space was partitioned in regions of three sizes: Class A networks had a large number of addresses, 2^{24} = 16,777,216, Class B networks had 2^{16} = 65,536 addresses each, and Class C networks had only 2^{8} = 128 addresses each. For example, the Rutgers University IP addresses belong to Class B because the network part starts with bits `10` (Figure 1-37), so the network part of the address is: `10000000 00000110` or `128.6.*` in dotted-decimal notation. The address space has been managed by IETF and organizations requested and obtained a set of addresses belonging to a class. As the Internet grew, most organizations were assigned Class B addresses, because their networks were too large for a Class C address, but not large enough for a Class A address. Unfortunately, large part of the address space went unused. For example, if an organization had slightly more than 128 hosts and acquired a Class B address, almost 65,400 addresses went unused and could not be assigned to another organization.

Figure 1-46 lists special IPv4 addresses.

CIDR Scheme for Internet Protocol (IPv4) Addresses

By 1991, it became clear that the 2^{14} = 16,384 Class B addresses would soon run out and a different approach was needed. It was observed that addresses from the enormous Class C space were rarely allocated and the solution was proposed to assign new organizations contiguous subsets of Class C addresses instead of a single Class B address. This allowed for a refined granularity of address space assignment. In this way, the allocated set of Class C addresses could be much better matched to the organization needs than with whole Class B sets. This solution *optimizes the common case*. The common case is that most organizations require at most a few

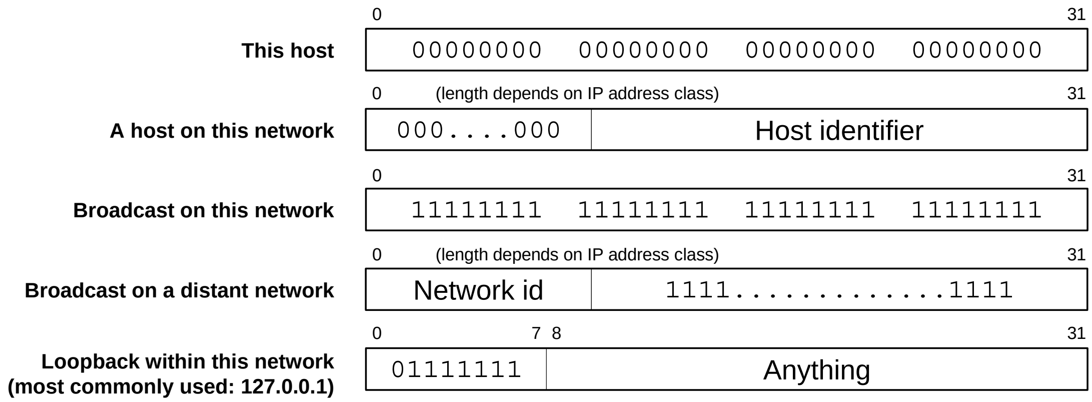

Figure 1-46: Special IP version 4 addresses.

thousand addresses, and this need could not be met with individual Class C sets, while an entire Class B represented a too coarse match to the need. A middle-road solution was needed.

Routing protocols that work with aggregated Class C address sets are said to follow **Classless Interdomain Routing** or **CIDR** (pronounced "cider"). CIDR not only solved the problem of address shortages, but also by aggregating Class C sets into contiguous regions, it reduced the forwarding table sizes because routers aggregate routes based on IP prefixes in a classless manner. Instead of having a forwarding-table entry for every individual address, the router now keeps a single entry for a subset of addresses (see analogy in Figure 1-44).

The CIDR-based addressing works as follows. An organization is assigned a region of the address space defined by two numbers, A and m. The assigned address region is denoted A/m. A is called the **prefix** and it is a 32-bit number (often written in dotted decimal notation) denoting the address space, while m is called the **mask** and it is a decimal number between 1 and 32. Therefore, when a network is assigned A/m, it means that it gets the $2^{(32-m)}$ addresses, all sharing the first m bits of A. For example, the network "192.206.0.0/21" corresponds to the $2^{(32-21)}$ = 2048 addresses in the range from 192.206.0.0 to 192.206.7.255.

SIDEBAR 1.2: Hierarchy without Topological Aggregation

♦ There are different ways to organize addresses hierarchically. Internet addresses are aggregated *topologically*, so that addresses in the same physical subnetwork share the same address prefix (or suffix). Another option is to partition the address space by manufacturers of networking equipment. The addresses are still globally unique, but not aggregated by proximity (i.e., network topology). An example is the Ethernet link-layer address, described in Section 1.5.2. Each Ethernet attachment adaptor has assigned a globally unique address, which has two parts: a part representing the manufacturer's code, and a part for the adaptor number. The manufacturer code is assigned by a global authority, and the adaptor number is assigned by the manufacturer. Obviously, each Ethernet adaptor on a given subnetwork may be from a different manufacturer, and noncontiguous subnetworks may have adaptors from the same manufacturer. However, this type of hierarchy is not suitable for routing purposes because it does not scale to networks with tens of millions of hosts, such as the Internet. Ethernet addresses cannot be aggregated in routing tables, and large-scale networks cannot use Ethernet addresses to identify destinations. Equally important, Ethernet addresses cannot be *summarized* and exchanged by the routers participating in the routing protocols. Therefore, topological aggregation of network addresses is the fundamental reason for the scalability of the Internet's network layer. (See more discussion in Section 8.3.1.)

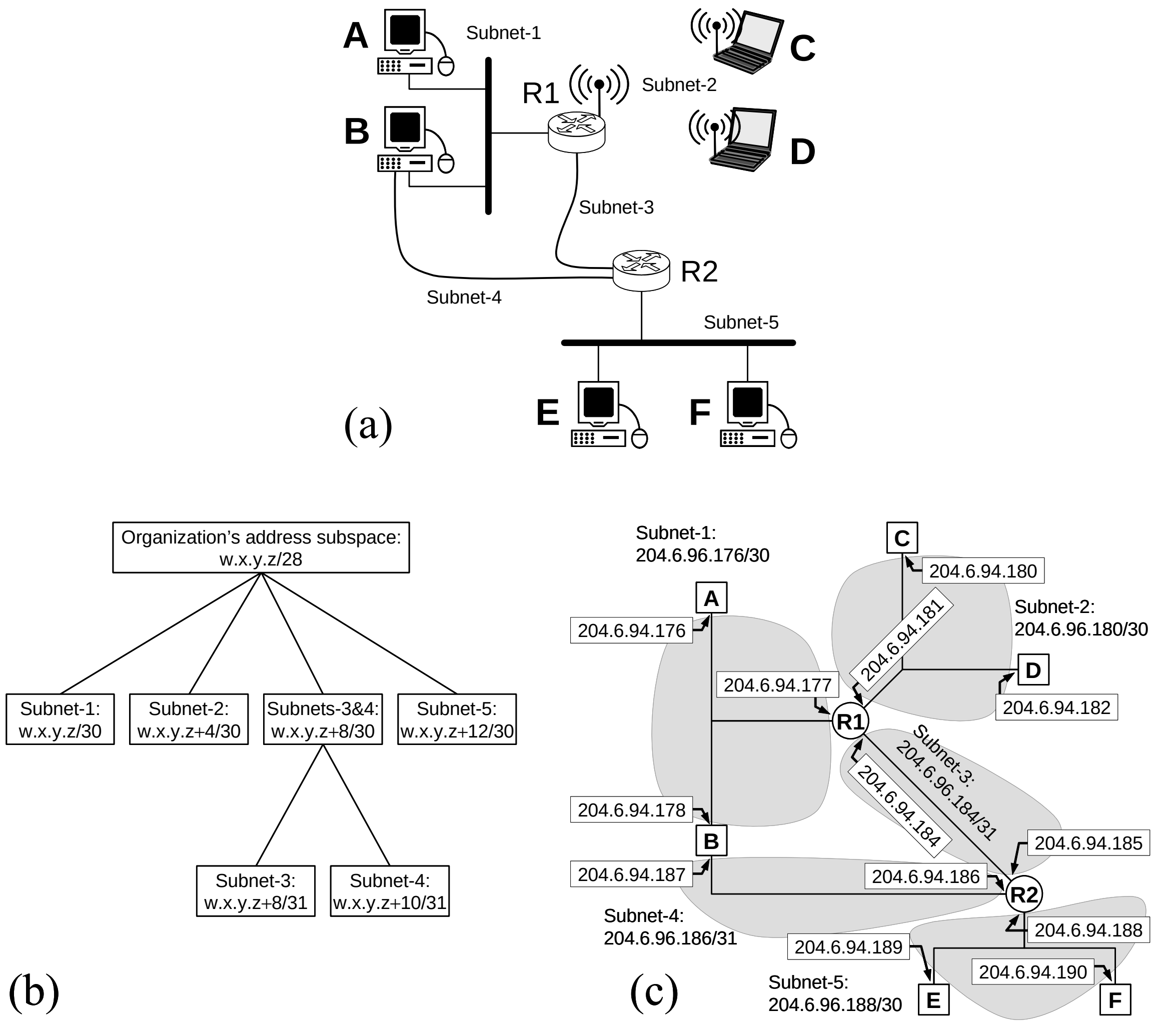

Figure 1-47: (a) Example internetwork with five physical networks reproduced from Figure 1-35 above. (b) Desired hierarchical address assignment under the CIDR scheme. (c) Example of an actual address assignment.

Suppose for the sake of illustration that you are administering your organization's network as shown in Figure 1-35, reproduced here in Figure 1-47(a). Assume that you know that this network will remain fixed in size, and your task is to acquire a set of network addresses and assign them optimally to the hosts. Your first task is to determine how many addresses to request. As seen in Section 1.4.1, both routers R1 and R2 have 3 network interfaces each. Because your internetwork has a total of 13 interfaces (3 + 3 for routers, 2 for host B and 5 × 1 for other hosts), you need 13 unique IP addresses. However, you would like to structure your organization's network hierarchically, so that each subnet is in its own address space, as shown in Figure 1-47(b). Subnets 3 and 4 have only two interfaces each, so they need 2 addresses each. Their assignments will have the mask $m = 31$. You can group these two in a single set with $m = 30$. Subnets 1, 2, and 5 have three interfaces each, so you need at least 2 bits (4 addresses) for each and their masks will equal $m = 30$. Therefore, you need 4 × 4 addresses (of which three will be unused) and your address region will be of the form *w.x.y.z*/28, which gives you $2^{(32-28)} = 2^4 = 16$ addresses. Let us assume that the actual address subspace assignment that you acquired is

Table 1-3: CIDR hierarchical address assignment for the internetwork in Figure 1-47.

Subnet	Subnet mask	Network prefix	Interface addresses
1	204.6.94.176/30	11001100 00000110 01011110 101100--	A: 204.6.94.176
			R1-1: 204.6.94.177
			B-1: 204.6.94.178
2	204.6.94.180/30	11001100 00000110 01011110 101101--	C: 204.6.94.180
			R1-2: 204.6.94.181
			D: 204.6.94.182
3	204.6.94.184/31	11001100 00000110 01011110 1011100-	R1-3: 204.6.94.184
			R2-1: 204.6.94.185
4	204.6.94.186/31	11001100 00000110 01011110 1011101-	R2-2: 204.6.94.186
			B-2: 204.6.94.187
5	204.6.94.188/30	11001100 00000110 01011110 101111--	R2-3: 204.6.94.188
			E: 204.6.94.189
			F: 204.6.94.190

204.6.94.176/28. Then you could assign the individual addresses to the network interfaces as shown in Table 1-3 as well as in Figure 1-47(c).

1.4.5 Autonomous Systems and Path Vector Routing

Problems related to this section: Problem 1.35 → ?

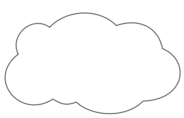

Figure 1-35 presents a naïve view of the Internet, where many hosts are mutually connected via intermediary nodes (routers or switches) that live inside the "network cloud." This would imply that the cloud is managed by a single administrative organization and all nodes cooperate to provide the best service to the consumers' hosts. In reality the Internet is composed of many independent networks (or, "clouds"), each managed by a different organization driven by its own commercial or political interests. (The reader may also wish to refer to Figure 1-3 to get a sense of complexity of the Internet.) Each individual administrative domain is known as an **autonomous system** (AS). Given their divergent commercial interests, these administrative domains are more likely to compete (for profits) than to collaborate in harmony with each other.

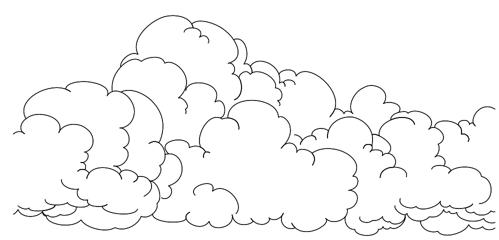

Both distance vector and link state routing protocols have been used for *interior routing* (or, *internal routing*). That is, they have been used inside individual administrative domains or autonomous systems. However, both protocols become ineffective in large networks composed of many domains (autonomous systems). The scalability issues of both protocols were discussed earlier. In addition, they do not provide mechanisms for an administrative entity to represent its economic interests as part of the routing protocol. Economic interests can be described using *logical rules* that express the routing *policies* to reflect the economic interests. For this purpose, we need *exterior routing* (or, *external routing*) protocols for routing between different autonomous systems.

We first review the challenges posed by interacting autonomous domains and then present the path vector routing algorithm that can be used to address some of those issues.

Autonomous Systems: Peering Versus Transit

An Autonomous System (AS) can independently decide whom to exchange traffic with on the Internet, and it is not dependent upon a third party for access. Networks of Internet Service Providers (ISPs), hosting providers, telecommunications companies, multinational corporations, schools, hospitals, and even individuals can be Autonomous Systems; all one needs is a unique **Autonomous System Number (ASN)** and a block of IP addresses. A central authority (http://iana.org/) assigns ASNs and assures their uniqueness. At the time of this writing (2010), the Internet consists of over 25,000 Autonomous Systems. Most organizations and individuals do not interconnect autonomously to other networks, but connect via an ISP. One could say that an end-user is "buying transit" from their ISP.

Figure 1-48 illustrates an example of several Autonomous Systems. In order to get traffic from one end-user to another end-user, ASs need to have an interconnection mechanism. These interconnections can be either *direct* between two networks or *indirect* via one or more intermediary networks that agree to transport the traffic. Most AS connections are indirect, since it is nearly impossible to interconnect directly with all networks on the globe. In order to make it from one end of the world to another, the traffic will often be transferred through several indirect interconnections to reach the end-user. The economic agreements that allow ASs to interconnect directly and indirectly are known as "peering" or "transit," and they are the two mechanisms that underlie the interconnection of networks that form the Internet.

A **peering** agreement (or, *swap* contract) is a voluntary interconnection of two or more autonomous systems for exchanging traffic between the customers of each AS. This is often done so that neither party pays the other for the exchanged traffic; rather, each derives revenue from its own customers. Therefore, it is also referred to as "settlement-free peering."

In a **transit** agreement (or, *pay* contract), one autonomous system agrees to carry the traffic that flows between another autonomous system and all other ASs. Since no network connects directly to all other networks, a network that provides transit will deliver some of the traffic indirectly via one or more other transit networks. A transit provider's routers will announce to other networks that they can carry traffic to the network that has bought transit. The transit provider receives a "transit fee" for the service.

The transit fee is based on a reservation made up-front for a certain speed of access (in Mbps) or the amount of bandwidth used. Traffic from (upstream) and to (downstream) the network is included in the *transit fee*; when one buys 10Mbps/month from a transit provider, this includes 10 up and 10 down. The traffic can either be limited to the amount reserved, or the price can be calculated afterward (often leaving the top five percent out of the calculation to correct for aberrations). Going over a reservation may lead to a penalty.

An economic agreement between ASs is implemented through (*i*) a physical interconnection of their networks, and (*ii*) an exchange of routing information through a common routing protocol. This section reviews the problems posed by autonomous administrative entities and requirements for a routing protocol between Autonomous Systems. Section 8.2.3 describes the protocol used in the current Internet, called Border Gateway Protocol (BGP), which meets these requirements.

The Internet is intended to provide *global reachability* (or, end-to-end reachability), meaning that any Internet user can reach any other Internet user as if they were on the same network. To be

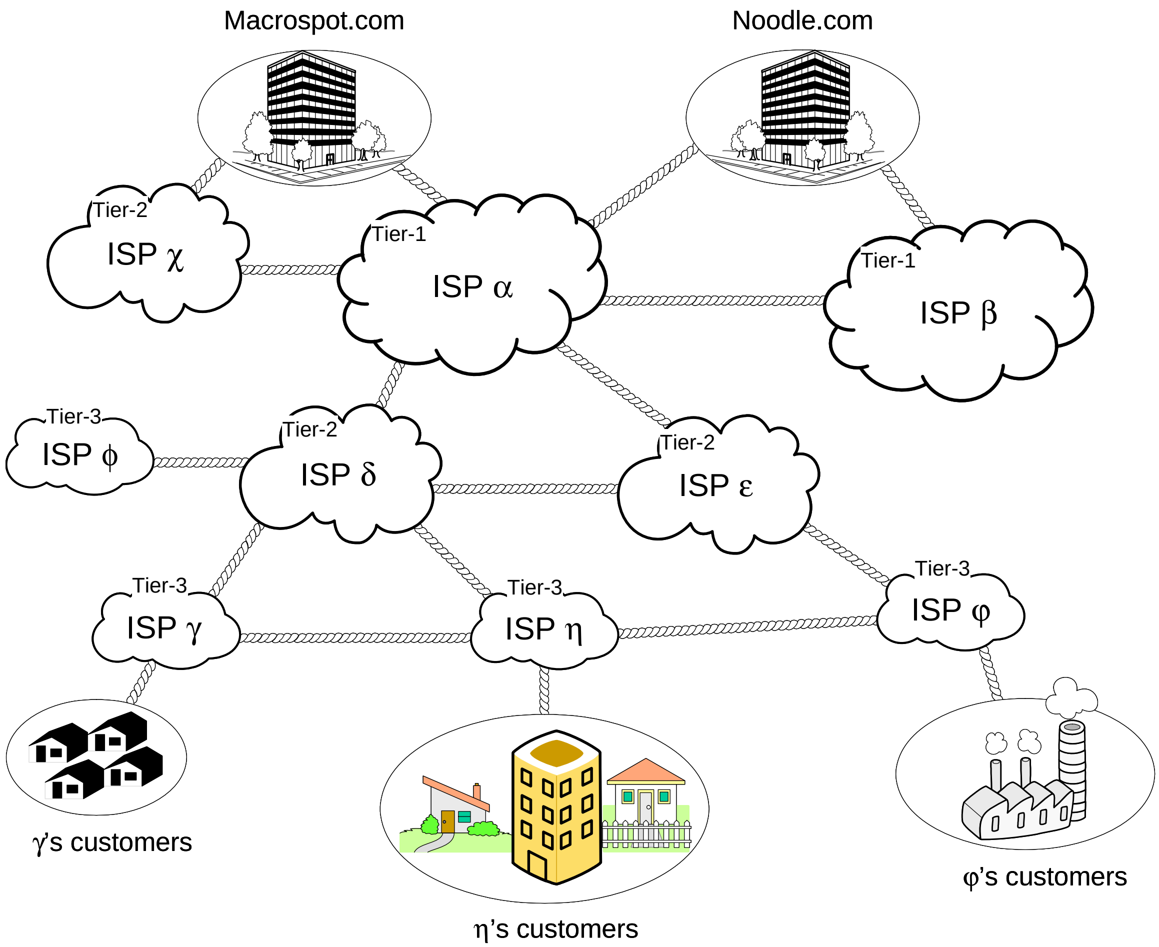

Figure 1-48: An example collection of Autonomous Systems with physical interconnections.

able to reach any other network on the Internet, Autonomous System operators work with each other in following ways:

- *Sell transit* (or Internet access) service to that AS ("transit provider" sells transit service to a "transit customer"),
- *Peer* directly with that AS, or with an AS who sells transit service to that AS, or
- *Pay* another AS for *transit* service, where that "transit provider" must in turn also sell, peer, or pay for access.

Therefore, any AS connected to the Internet must either pay another AS for transit, or peer with every other AS that also does not purchase transit.

Consider the example in Figure 1-48. Tier-1 Internet Service Providers (ISPα and ISPβ) have global reachability information and can see all other networks and, because of this, their forwarding tables do not have default entries. They are said to be *default-free*. At present (2010) there are about 10 Tier-1 ISPs in the world. The different types of ASs (mainly by their size) lead to different business relationships between them. ISPs enter peering agreements mostly with other ISPs of the similar size (reciprocal agreements). Therefore, a Tier-1 ISP would form a peering agreement with other Tier-1 ISPs, and sell transit to lower tiers ISPs. Similarly, a Tier-2 (regional

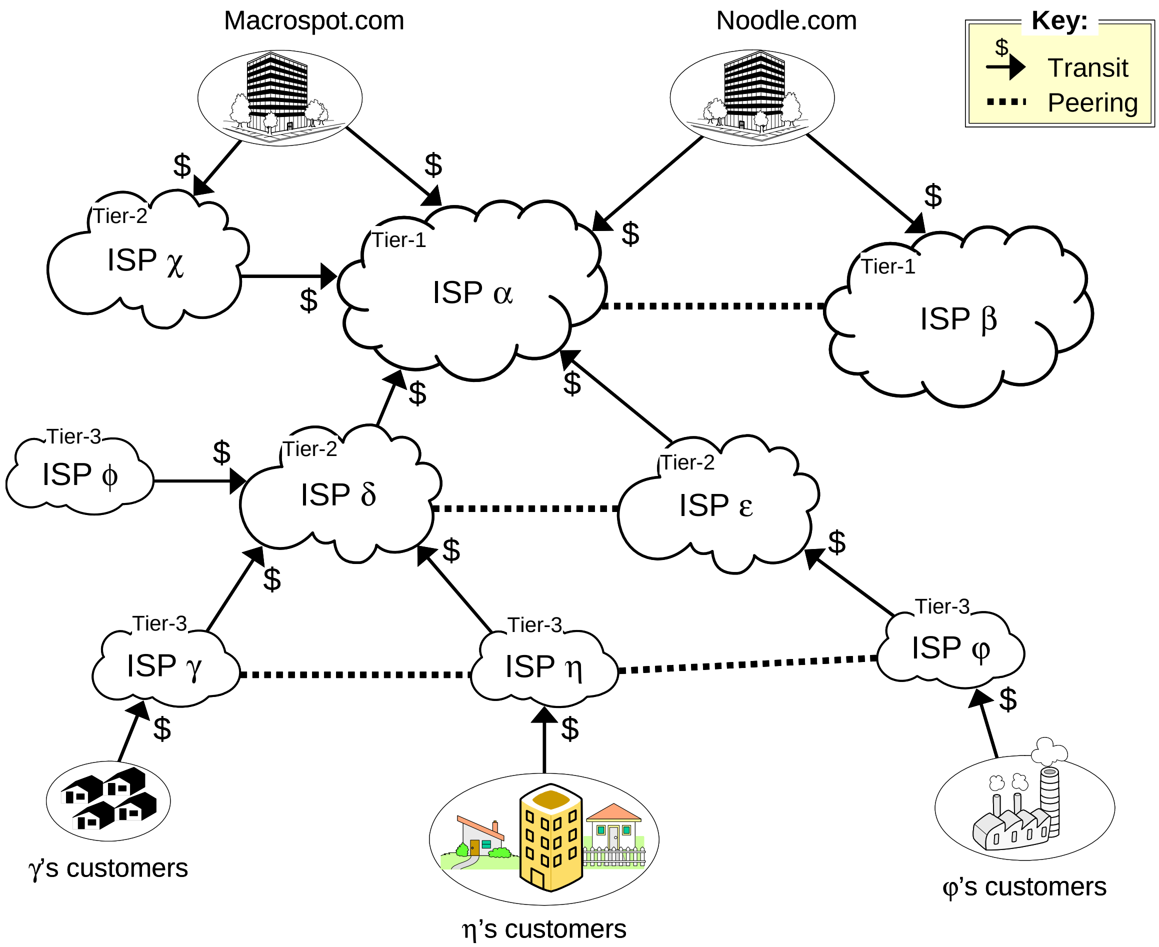

Figure 1-49: Feasible business relationships for the example ASs in Figure 1-48.

or countrywide) ISP would form a peering agreement with other Tier-2 ISPs, pay for transit service to a Tier-1 ISP, and sell transit to lower Tier-3 ISPs (local). As long as the traffic ratio of the concerned ASs is not highly asymmetrical (e.g., up to 4-to-1 is a commonly accepted ratio), there is usually no financial settlement for peering.

Transit relationships are preferable because they generate revenue, whereas peering relationships usually do not. However, peering can offer reduced costs for transit services and save money for the peering parties. Other less tangible incentives ("mutual benefit") include:

- Increased redundancy (by reducing dependence on one or more transit providers) and improved performance (attempting to bypass potential bottlenecks with a "direct" path),
- Increased capacity for extremely large amounts of traffic (distributing traffic across many networks) and ease of requesting for emergency aid (from friendly peers).

Figure 1-49 shows reasonable business relationships between the ISPs in Figure 1-48. ISPϕ cannot peer with another Tier-3 ISP because it has a single physical interconnection to a Tier-2 ISPδ. An Autonomous System that has only a single connection to one other AS is called **stub AS**. The two large corporations at the top of Figure 1-49 each have connections to more than one other AS but they refuse to carry transit traffic; such an AS is called **multihomed AS**. ISPs

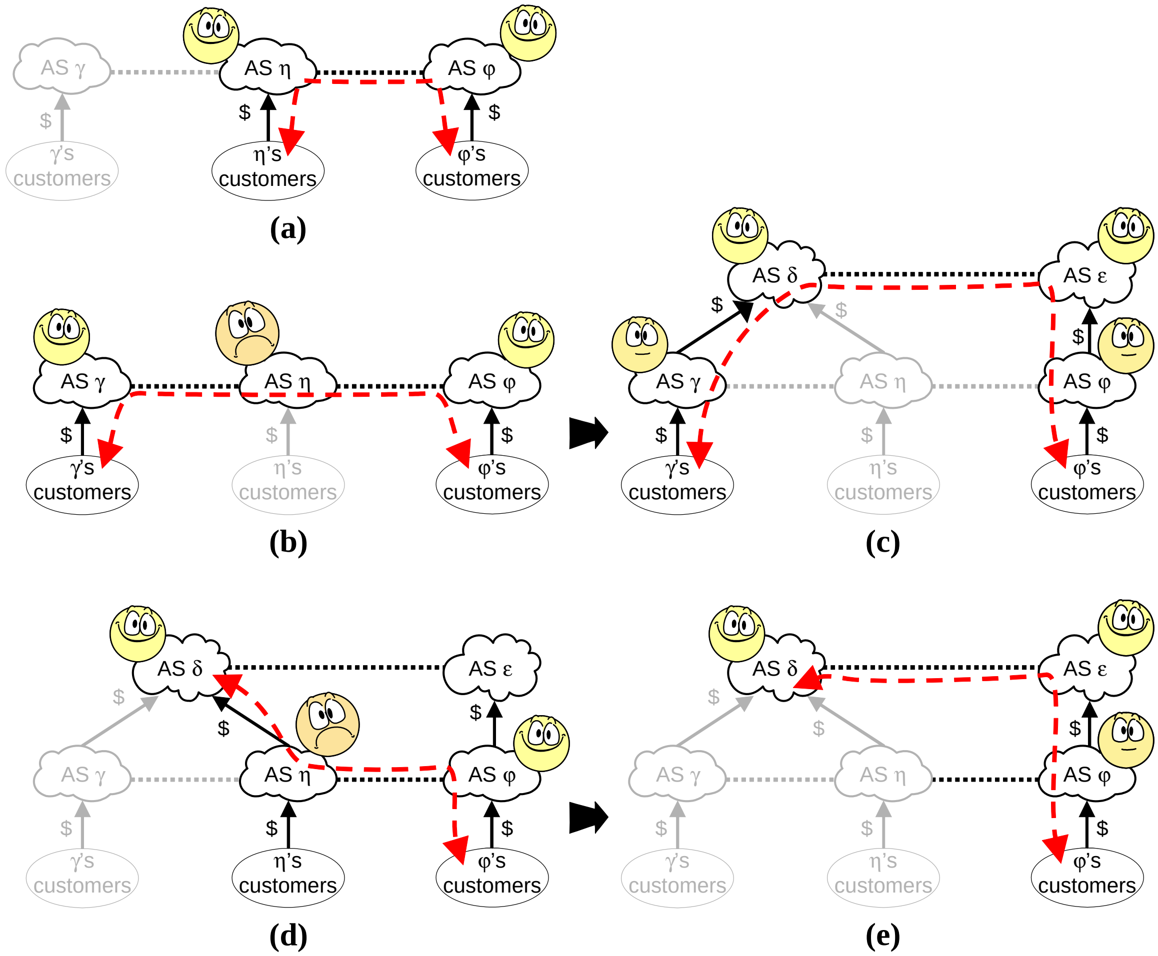

Figure 1-50: Providing selective transit service to make or save money.

usually have connections to more than one other AS and they are designed to carry both transit and local traffic; such an AS is called **transit AS**.

When two providers form a peering link, the traffic flowing across that link incurs a cost on the network it enters. Such a cost may be felt at the time of network provisioning: in order to meet the negotiated quantity of traffic entering through a peering link, a provider may need to increase its network capacity. A network provider may also see a cost for entering traffic on a faster timescale; when the amount of incoming traffic increases, congestion on the network increases, and this leads to increased operating and network management costs. For this reason, each AS needs to decide carefully what kind of transit traffic it will support.

Each AS is in one of three types of business relationships with the ASs to which is has a direct physical interconnection: transit provider, transit customer, or peer. To its paying customers, the AS wants to provide unlimited transit service. However, to its provider(s) and peers it probably wishes to provide a *selective transit service*. Figure 1-50 gives examples of how conflicting interests of different parties can be resolved. The guiding principle is that ASs will want to avoid highly asymmetrical relationships without reciprocity. In Figure 1-50(a), both ASη and ASφ benefit from peering because it helps them to provide global reachability to their own customers. In Figure 1-50(b), ASγ and ASφ benefit from using transit service of ASη (with whom both of them are peers), but ASη may lose money in this arrangement (because of degraded service to its own customers) without gaining any benefit. Therefore, ASη will not carry transit traffic between its peers. An appropriate solution is presented in Figure 1-50(c), where ASγ and ASφ use their transit providers (ASδ and ASε, respectively), to carry their mutual transit traffic. ASδ and ASε are peers and are happy to provide transit service to their transit customers (ASγ and ASφ). Figure 1-50(d) shows a scenario where higher-tier ASδ uses its transit customer ASη to gain reachability of ASφ. Again, ASη does not benefit from this arrangement, because it pays ASδ for transit and does not expect ASδ in return to use its transit service for free. The appropriate solution is shown in Figure 1-50(e) (which is essentially the same as Figure 1-50(c)).

To implement these economic decisions and prevent unfavorable arrangements, ASs design and enforce *routing policies*. An AS that wants avoid providing transit between two neighboring ASs, simply does not advertise to either neighbor that the other can be reached via this AS. The neighbors will not be able to "see" each other via this AS, but via some other ASs. Routing policies for selective transit can be summarized as:

- To its transit customers, the AS should make visible (or, reachable) all destinations that it knows of. That is, all routing advertisements received by this AS should be passed on to own transit customers;
- To its peers, the AS should make visible only its own transit customers, but not its other peers or its transit provider(s), to avoid providing unrecompensed transit;
- To its transit providers, the AS should make visible only its own transit customers, but not its peers or its other transit providers, to avoid providing unrecompensed transit.

In the example in Figure 1-49, Tier-1 ISPs (ASα and ASβ) can see all the networks because they peer with one another and all other ASs buy transit from them. ASγ can see ASη and its customers directly, but not ASφ through ASη. ASδ can see ASφ through its peer ASε, but not via its transit customer ASη. Traffic from ASφ to ASϕ will go trough ASε (and its peer ASδ), but not through ASη.

To illustrate how routers in these ASs implement the above economic policies, let us imagine example routers as in Figure 1-51. Suppose that a router in ASϕ sends an update message advertising the destination prefix `128.34.10.0/24`. The message includes the **routing path vector** describing how to reach the given destination. The path vector starts with a single AS number {ASϕ}. A border router (router *K*) in ASδ receives this message and disseminates it to other routers in ASδ. Routers in ASδ prepend their own AS number to the message path vector to obtain {ASδ, ASϕ} and redistribute the message to the adjacent ASs. Because ASη does not have economic incentive to advertise a path to ASϕ to its peer ASφ, it sends an update message with path vector containing only the information about ASη's customers. On the other hand, ASε has economic incentive to advertise global reachability to its own transit customers. Therefore, routers in ASε prepend their own AS number to the routing path vector from ASδ to obtain {ASε, ASδ, ASϕ} and redistribute the update message to ASφ. Routers in ASφ update their routing and forwarding tables based on the received path vector. Finally, when a router in ASφ needs to send a data packet to a destination in the subnet `128.34.10.0/24` (in ASϕ) it sends the packet first to the next hop on the path to ASϕ, which is ASε.

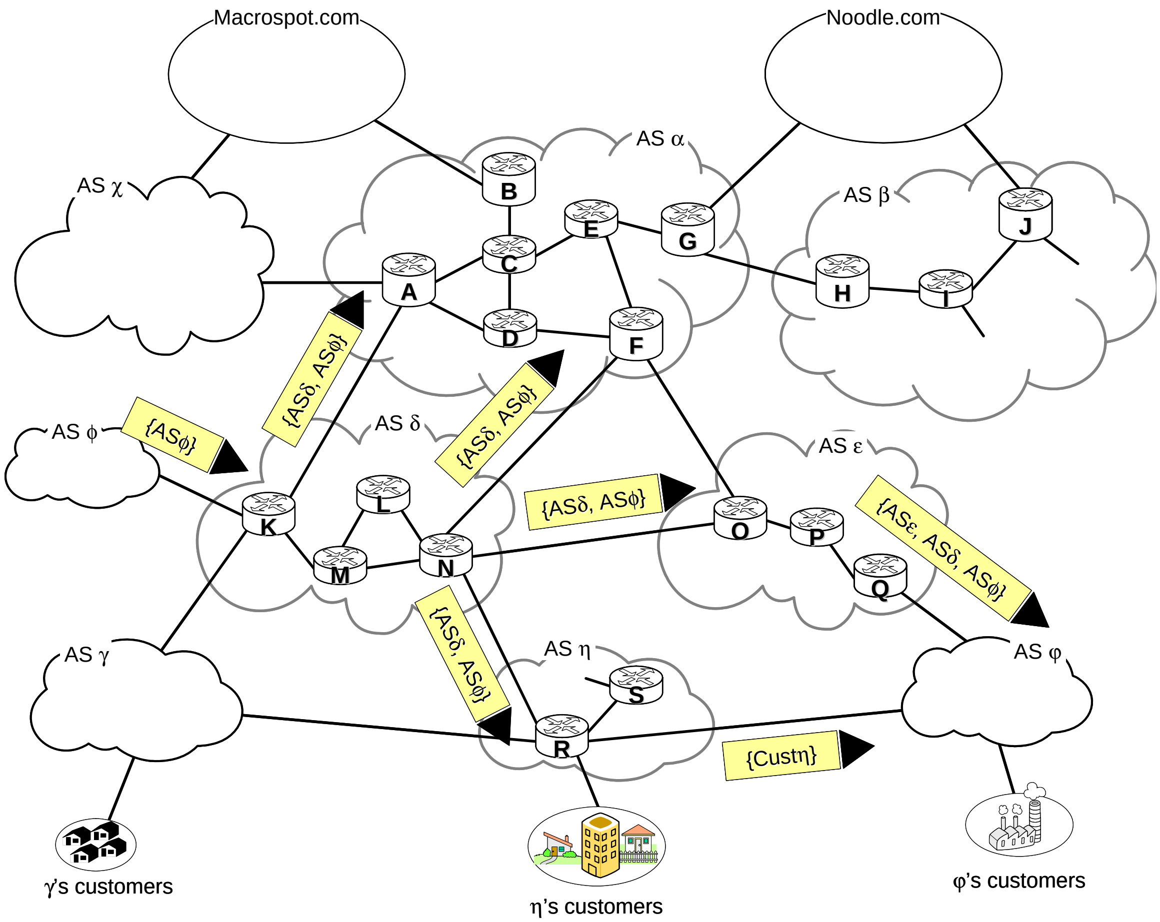

Figure 1-51: Example of routers within the ASs in Figure 1-48. Also shown is how a routing update message from ASϕ propagates to ASφ.

Path Vector Routing

Path vector routing is used for *inter-domain* or *exterior routing* (routing between different Autonomous Systems). The path vector algorithm is somewhat similar to the distance vector algorithm (Section 1.4.3). Each border (or edge) router in a given AS advertises the destinations it can reach to its neighboring routers (in different ASs). However, instead of advertising the networks in terms of a destination address and the distance to that destination, the networks are advertised as destination addresses with path descriptions to reach those destinations. A route is defined as a pairing between a destination and the attributes of the path to that destination, thus the name, **path vector routing**. The *path vector* contains a complete path as a sequence of ASs to reach the given destination. The path vector is carried in a special path attribute that records the sequence of ASs through which the reachability message has passed. The path that contains the smallest number of ASs becomes the preferred path to reach the destination.

At predetermined times, each node advertises its own network address and a copy of its path vector down every attached link to its immediate neighbors. An example is shown in Figure 1-51, where a router in ASϕ sends a scheduled update message. After a router receives path vectors from its neighbors, it performs path selection by merging the information received from its neighbors with that already in its existing path vector. The path selection is based on some kind of path metric, similar to distance vector routing algorithm (Section 1.4.3). Again, Eq. (1.14) is applied to compute the "shortest" path. Here is an example:

Example 1.5 Path Vector Routing Algorithm

Consider the network topology in Figure 1-41(a) (reproduced below) and assume that it uses the path vector routing algorithm. Instead of router addresses, the path vector works with Autonomous System Numbers (ASNs). Starting from the initial state for all nodes, show the first few steps until the routing algorithm reaches a stable state.

The solution is similar to that for the distributed distance vector routing algorithm (Example 1.4). For AS α, the initial routing table is as in the leftmost table below. The notation ⟨*d* | χ, ξ, ζ⟩ symbolizes that the path from the AS under consideration to AS ζ is *d* units long, and χ and ξ are the ASs along the path to ζ. If the path metric simply counts the number of hops, then the path-vector packets do not need to carry the distance *d*, because it can be determined simply by counting the ASs along the path.

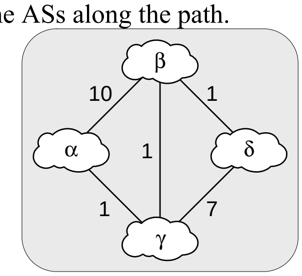

Routing table at AS α: Initial

From \ Path to	α	β	γ
α	⟨0 | α⟩	⟨10 | β⟩	⟨1 | γ⟩
β	⟨∞ | ⟩	⟨∞ | ⟩	⟨∞ | ⟩
γ	⟨∞ | ⟩	⟨∞ | ⟩	⟨∞ | ⟩

+

Received Path Vectors (1st exchange)

	α	β	γ	δ
From β	⟨10 | α⟩	⟨0 | β⟩	⟨1 | γ⟩	⟨1 | δ⟩

	α	β	γ	δ
From γ	⟨1 | γ⟩	⟨1 | β⟩	⟨0 | γ⟩	⟨7 | δ⟩

Routing table at node α: After 1st exchange:

=

From \ Path to	α	β	γ	δ
α	⟨0 | α⟩	⟨2 | γ, β⟩	⟨1 | γ⟩	⟨8 | γ, δ⟩
β	⟨10 | α⟩	⟨0 | β⟩	⟨1 | γ⟩	⟨1 | δ⟩
γ	⟨1 | α⟩	⟨1 | β⟩	⟨0 | γ⟩	⟨7 | δ⟩

Again, AS α only keeps the path vectors of its immediate neighbors, β and γ, and not that of any other ASs, such as δ. Initially, α may not even know that δ exists. Next, each AS advertises its path vector to its immediate neighbors, and α receives their path vectors. When an AS receives an updated path vector from its neighbor, the AS overwrites the neighbor's old path vector with the new one. In addition, *A* re-computes its own path vector according to Eq. (1.14), as follows:

$$D_{\alpha}(\beta) = \min\{c(\alpha,\beta) + D_{\beta}(\beta),\ c(\alpha,\gamma) + D\gamma(\beta)\} = \min\{10+0,\ 1+1\} = 2 \qquad \Rightarrow \text{path: } \langle 2 \mid \gamma, \beta\rangle$$

$$D_{\alpha}(\gamma) = \min\{c(\alpha,\beta) + D_{\beta}(\gamma),\ c(\alpha,\gamma) + D_{\gamma}(\gamma)\} = \min\{10+1,\ 1+0\} = 1 \qquad \Rightarrow \text{path: } \langle 1 \mid \gamma\rangle$$

$$D_{\alpha}(\delta) = \min\{c(\alpha,\beta) + D_{\beta}(\delta),\ c(\alpha,\gamma) + D_{\gamma}(\delta)\} = \min\{10+1,\ 1+7\} = 8 \qquad \Rightarrow \text{path: } \langle 8 \mid \gamma, \delta\rangle$$

The new values for α's path vector are shown in the above table at the right. Notice that α's new path to β is via γ and the corresponding table entry is ⟨2 | γ, β⟩.

Similar computations will take place on all other nodes and the whole process is illustrated in [Figure XYZ]. The end result is as shown in [Figure XYZ] as column entitled "After 1st exchange." Because for every node the path vector computed after the first exchange is different from the previous one, each node advertises its path new vector to its immediate neighbors. The cycle repeats for every node until there is no difference between the new and the previous path vector. As shown in [Figure XYZ], this happens after three exchanges.

To implement routing between Autonomous Systems, each Autonomous System must have one or more border routers that are connected to networks in two or more ASs (its own network and a neighboring AS network). Such a node is called a **speaker node** or **gateway router**. For example, in Figure 1-51 the speaker nodes in ASα are routers *A*, *B*, *F*, and *G*; in ASβ the speaker nodes are routers *H* and *J*; and in ASδ the speakers are routers *K* and *N*. A speaker node creates a routing table and advertises it to adjoining speaker nodes in the neighboring Autonomous Systems. The idea is the same as with distance vector routing, except that only speaker nodes in each Autonomous System can communicate with routers in other Autonomous Systems (i.e., speaker nodes in those ASs). The speaker node advertises the path, not the metric of the links, in its AS or other ASs. In other words, there are no weights attached to the links in a path vector, but there is an overall cost associated with each path.

Integrating Inter-Domain and Intra-Domain Routing

Administrative entities that manage different Autonomous Systems have different concerns for routing messages within their own Autonomous System as opposed to routing messages to other Autonomous Systems or providing transit service for them. Within an Autonomous System, the key concern is how to route data packets from the origin to the destination in the *most efficient manner*. For this purpose, *intra-domain* or *interior routing protocols*, such as those based on distance-vector routing (Section 1.4.3) or link-state routing (Section 1.4.2). These protocols are known as **Interior Gateway Protocols (IGPs)**. Unlike this, the key concern of any given Autonomous System is how to route data packets from the origin to the destination in the manner that is *most profitable for this AS*. These protocols are known as **Exterior Gateway Protocols**[10] and are based on path-vector routing, described above. This duality in routing goals and solutions means that each border router (or, speaker node) will maintain two different routing tables: one obtained by the interior routing protocol and the other by the exterior routing protocol.

A key problem is how the speaker node should integrate its dual routing tables into a meaningful forwarding table. The speaker that first receives path information about a destination in another AS simply adds a new entry into its forwarding table. However, the problem is how to exchange the routing information with all other routers within this AS and achieve a consistent picture of the Internet viewed by all of the routers within this ASs. The goal is that, for a given data packet, each router in this AS should make the same forwarding decision (as if each had access to the routing tables of all the speaker routers within this AS). Each speaker node must exchange its routing information with all other routers within its own AS (known as *internal peering*). This includes both other speakers in the same AS (if there are any), as well as the remaining non-speaker routers. For example, in Figure 1-51 speaker *K* in ASδ needs to exchange routing information with speaker *N* (and vice versa), as well as with non-speaker routers *L* and *M*. Notice again that only speaker routers run both IGP and Exterior Gateway Protocol (and each maintains two routing tables); non-speaker routers run only IGP and maintain a single routing table.

The forwarding table contains pairs $\langle$*destination*, *output-port*$\rangle$ for all possible destinations. The output port corresponds to the IP address of the next hop router to which the packet will be forwarded. Recall that each router performs the longest CIDR prefix match on each packet's

---

[10] In the literature, the acronym EGP is *not* used for a generic Exterior Gateway Protocol, because EGP refers to an actual protocol, described in RFC-904, now obsolete, that preceded BGP (Section 8.2.3).

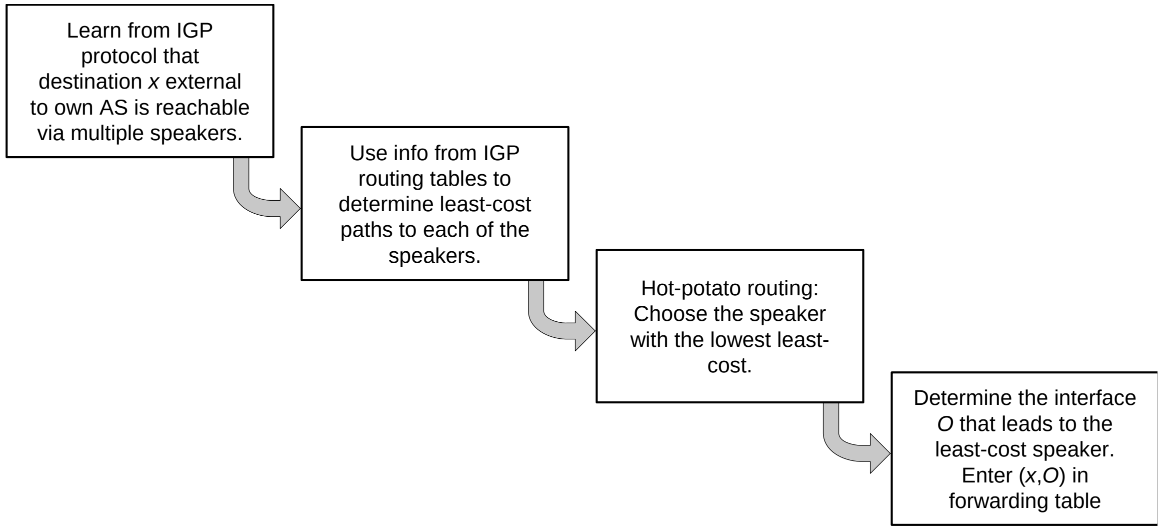

Figure 1-52: Integrating an external destination to a router's forwarding table.

destination IP address (Section 1.4.4). All forwarding tables must have a default entry for addresses that cannot be matched, and only routers in Tier-1 ISPs are default-free because they know prefixes to all networks in the global Internet.

Consider again the scenario shown in Figure 1-51 where ASφ advertises the destination prefix `128.34.10.0/24`. If the AS has a single speaker node leading outside the AS, then it is easy to form the forwarding table. For example, in Figure 1-51 ASη has a single speaker router *R* that connects it to other ASs. If non-speaker router *S* in ASη receives a packet destined to ASφ, the packet will be forwarded along the shortest path (determined by the IGP protocol running in ASη) to the speaker *S*, which will then forward it to *N* in ASδ. Consider now a different situation where router *B* in ASα receives a packet destined to ASφ. *B* should clearly forward the packet to another speaker node, but which one? As seen in Figure 1-51, both *A* or *F* will learn about ASφ but via different routes. To solve the problem, when a speaker node learns about a destination outside its own AS, it must disseminate this information to all routers within its own AS. This dissemination is handled by the AS's interior gateway protocol (IGP).

When a router learns from an IGP advertisement about a destination outside its own AS, it needs to add the new destination into its forwarding table. This applies to both non-speaker routers and speaker routers that received this IGP advertisement from the fellow speaker router (within the same AS), which first received the advertisement via exterior gateway protocol from a different AS. One approach that is often employed in practice is known as **hot-potato routing**. In hot-potato routing, the Autonomous System gets rid of the packet (the "hot potato") as quickly as possible (more precisely, as inexpensively as possible). This is achieved by having the router send the packet to the speaker node that has the lowest router-to-speaker cost among all speakers with a path to the destination. Figure 1-52 summarizes the steps taken at a router for adding the new entry to its forwarding table. In Figure 1-51, when *B* receives a packet for ASφ it will send it to *A* or *F* based on the lowest cost *within* ASα only, rather than overall lowest cost to the destination.

The most popular practical implementation of path vector routing is *Border Gateway Protocol* (BGP), currently in version 4 (BGP4). Section 8.2.3 describes how BGP4 meets the above requirements.

1.5 Link-Layer Protocols and Technologies

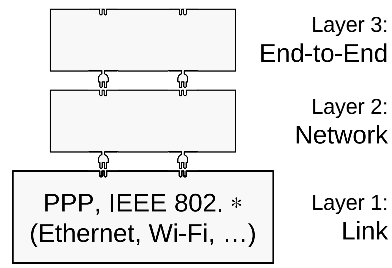

In packet-switched networks, blocks of data bits (generally called packets) are exchanged between the communicating nodes. That is, the nodes send packets rather than continuous bit-streams. At the link layer, packets are called *frames*. The key function of the link layer is transferring frames from one node to an adjacent node over a communication link. This task is complex because there are a great variety of communication link types. The key characteristics of a link include *data rate*, *duplexity* (half or full duplex), and *multiplicity* of the medium (i.e., point-to-point or shared broadcast). The link-layer services include:

• *Framing* is encapsulating a network-layer datagram into a link-layer frame by adding the header and the trailer. It is particularly challenging for a receiving node to recognize where an arriving frame begins and ends. For this purpose, special control bit-patterns are used to identify the start and end of a frame. On both endpoints of the link, receivers are continuously hunting for the start-of-frame bit-pattern to synchronize on the start of the frame. Having special control codes, in turn, creates the problem of *data transparency* (the need to avoid confusion between control codes and data) and requires *data stuffing* (described earlier in Figure 1-14).

• *Medium access control* (MAC) allows sharing a broadcast medium (Section 1.3.3) MAC addresses are used in frame headers to identify the sender and the receiver of the frame. MAC addresses are different from IP addresses and require a special mechanism for translation between different address types (Section 8.3.1). Point-to-point protocols do not need MAC.

• *Reliable delivery* between adjacent nodes includes error detection and error recovery. The techniques for error recovery include forward error correction code (Section 1.2) and retransmission by ARQ protocols (Section 1.3).

• *Connection liveness* is the ability to detect a link outage that makes impossible to transfer data over the link. For example, a wire could be cut, or a metal barrier could disrupt the wireless link. The link-layer protocol should signal this error condition to the network layer.

• *Flow control* is pacing between adjacent sending and receiving nodes to avoid overflowing the receiving node with messages at a rate it cannot process. A link-layer receiver is expected to be able to receive frames at the full datarate of the underlying physical layer. However, a higher-layer receiver may not be able receive packets at this full datarate. It is usually left up to the higher-layer receiver to throttle the higher-layer sender. (An example for the TCP protocol will be seen in Section 2.1.3.) Sometimes the link layer may also participate in flow control. A simple way of exerting backpressure on the upper-layer protocol is shown in Listing 1-1 (Section 1.1.4) at the start of the method `send()`, where an exception is thrown if the buffer for storing the unacknowledged packets is full.

There are two types of communication links: (1) *point-to-point link* with one sender and one receiver on the link, and no medium access control (MAC) or explicit MAC addressing; and, (2)

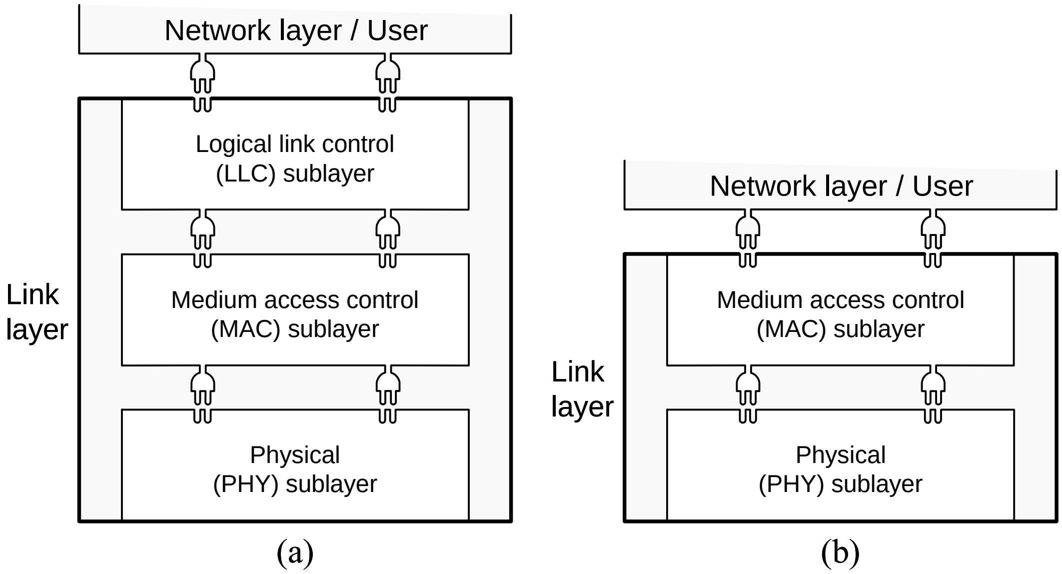

Figure 1-53: Sublayers of the link layer for broadcast communication links.

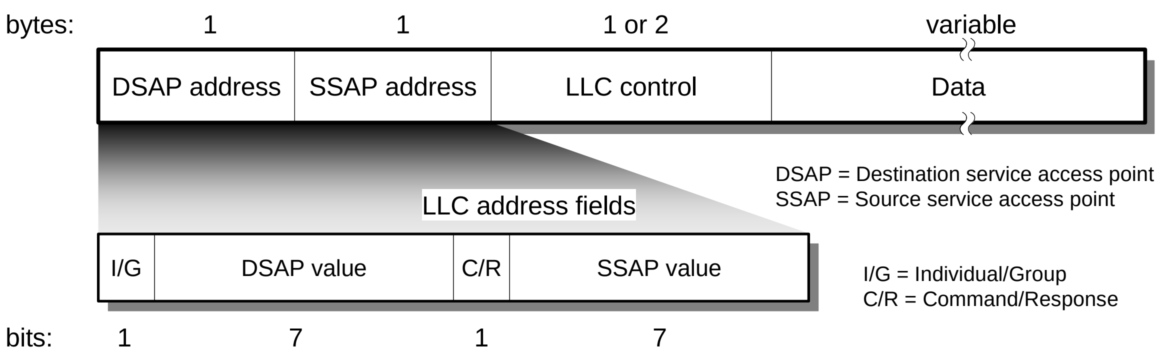

Figure 1-54: Packet format for Logical Link Control (LLC) protocol.

broadcast link over a shared wire or air medium. Point-to-point link is easier to work with than a broadcast link because broadcast requires coordination of many stations for accessing the medium. The basics of medium access control are already described in Section Section 1.3.3 and more will be covered later in this section.

Because broadcast links are so complex, it is common to subdivide the link layer of the protocol stack into three sublayers (Figure 1-53): logical-link-control (LLC) sublayer, medium access control (MAC) sublayer, and physical (PHY) sublayer. In the OSI reference model (Section 1.1.4), Layer 2 is subdivided into two sublayers: LLC and MAC sublayer. The network layer may directly use the services of a MAC sublayer (Figure 1-53(b)), or it may interact with a logical-link-control (LLC) sublayer (Figure 1-53(a)). We will see examples of both approaches later in this section. The IP protocol (Section 1.4.1) usually directly interacts with a MAC sublayer.

IEEE specified the 802.2 standard for LLC, which is the common standard for all broadcast links specified by the IEEE Working Group 802, such as Ethernet (Section 1.5.2) and Wi-Fi (Section 1.5.3) broadcast links. 802.2 LLC hides the differences between various kinds of IEEE 802 links by providing a single frame format and service interface to the network layer. 802.2 LLC also provides options for reliable delivery and flow control. Figure 1-54 shows the LLC packet format.

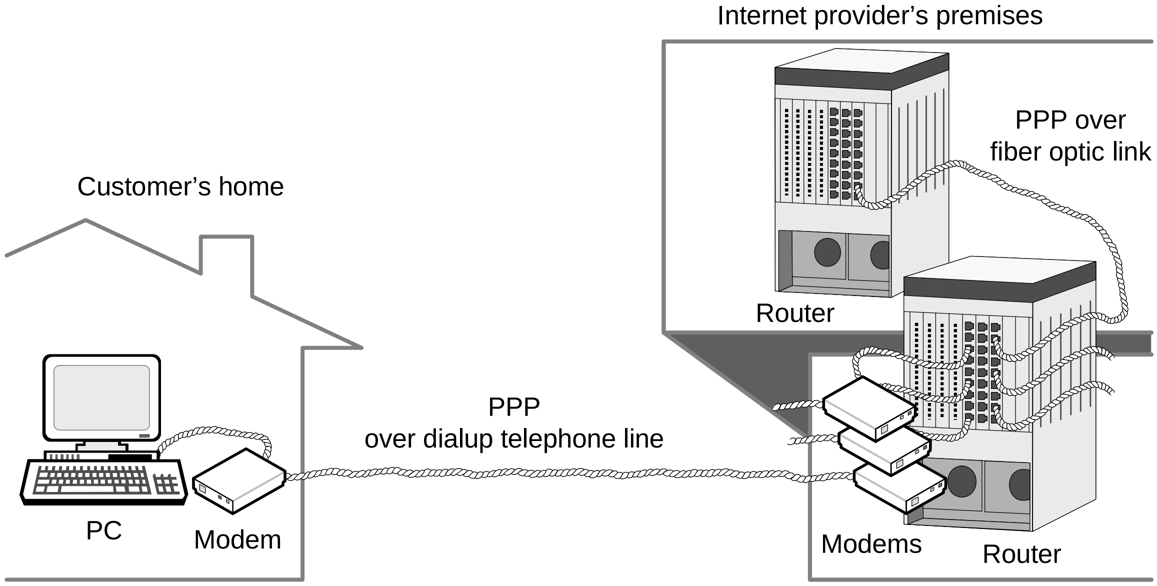

Figure 1-55: Point-to-point protocol (PPP) provides link-layer connectivity between a pair of network nodes over many types of physical networks.

The two address fields specify the destination and source users of LLC, where the "user" is usually an upper-layer protocol, such as IP (Figure 1-53). The LLC user addresses are referred to as "service access points" (SAPs), which is the OSI terminology for the user of a protocol layer. The DSAP address field identifies one or more destination users for which the LLC packet data is intended. This field corresponds to the `receivingProtocol` field in Listing 1-1. The SSAP address field identifies the upper-layer protocol that sent the data.

Section 1.5.1 reviews a link-layer protocol for point-to-point links. Sections 1.5.2 and 1.5.3 review link-layer protocol for broadcast links: Ethernet for wire broadcast links and Wi-Fi for wireless broadcast links. Within a single building, broadcast local-area networks such as Ethernet or Wi-Fi are commonly used for interconnection. However, most of the wide-area (long distance) network infrastructure is built up from point-to-point leased lines.

1.5.1 Point-to-Point Protocol (PPP)

Problems related to this section: Problem 1.40

Figure 1-55 illustrates two typical scenarios where point-to-point links are used. The first is for telephone dialup access, where a customer's PC calls up an Internet service provider's (ISP) router and then acts as an Internet host. When connected at a distance, each endpoint needs to be fitted with a *modem* to convert analog communications signals into a digital data stream. Figure 1-55 shows modems as external to emphasize their role, but nowadays computers have built-in modems. Another frequent scenario for point-to-point links is connecting two distant routers that belong to the same or different ISPs (right-hand side of Figure 1-55). Two most popular point-to-point link-layer protocols are **PPP** (point-to-point protocol), which is *byte-oriented*, viewing each frame as a collection of bytes; and **HDLC** (high-level data link control), which is *bit-oriented*. PPP, although derived from HDLC, is simpler and includes only a subset of HDLC functionality.

Visit http://en.wikipedia.org/wiki/HDLC for more details on High-Level Data Link Control (HDLC)

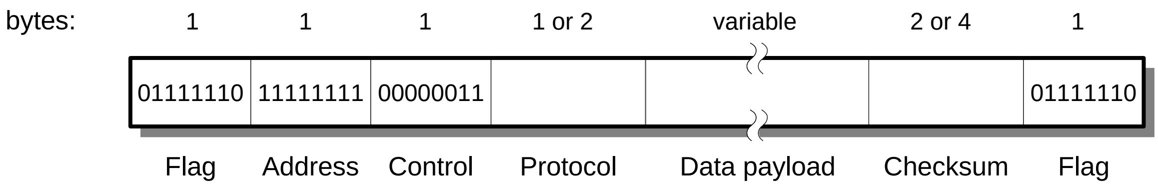

Figure 1-56: Point-to-point protocol (PPP) frame format.

(This book does not cover HDLC and the reader should check the bibliography in Section 1.7 for relevant references.)

The format of a PPP frame is shown in Figure 1-56. The PPP frame always begins and ends with a special character (called "flag"). The Flag makes it possible for the receiver to recognize the boundaries of an arriving frame. Notice that the PPP frame header does not include any information about the frame length, so the receiver recognizes the end of the frame when it encounters the trailing Flag field. The second field (Address) normally contains all ones (the broadcast address of HDLC), which indicates that all stations should accept this frame. Because there are only two hosts attached to a PPP link, PPP uses the broadcast address to avoid having to assign link-layer addresses. The third field (Control) is set to a default value 00000011. This value indicates that PPP is run in connectionless mode, meaning that frame sequence numbers are *not* used and out-of-order delivery is acceptable.

Because the Address and Control fields are always constant in the default configuration, the nodes can negotiate an option to omit these fields and reduce the overhead by 2 bytes per frame.

The Protocol field is used for demultiplexing at the receiver: it identifies the upper-layer protocol (e.g., IP) that should receive the payload of this frame. The code for the IP protocol is hexadecimal 21_{16}. The reader may wish to check Listing 1-1 (Section 1.1.4) and see how the method `handle()` calls `upperProtocol.handle()` to handle the received payload.

The Payload field is variable length, up to some negotiated maximum; if not negotiated, the default length of 1500 bytes is used. After Payload comes the Checksum field, which is by default 2 bytes, but can be negotiated to a 4-byte checksum. PPP checksum only detects errors, but has no error correction/recovery.

Figure 1-57 summarizes the state diagram for PPP; the actual finite state machine of the PPP protocol is more complex and the interested reader should consult RFC-1661 [Simpson, 1994]. There are two key steps before the endpoints can start exchanging network-layer data packets:

1. **Establishing link connection:** during this phase, the link-layer connection is set up. The link-layer peers must configure the PPP link (e.g., maximum frame length, authentication, whether to omit the Address and Control fields). PPP's **Link Control Protocol (LCP)** is used for this purpose.

2. **Connecting to network-layer protocol:** after the link has been established and options negotiated by the LCP, PPP must choose and configure one or more network-layer protocols that will operate over the link. PPP's **Network Control Protocol (NCP)** is used for this purpose. Once the chosen network-layer protocol has been configured, datagrams can be sent over the link.

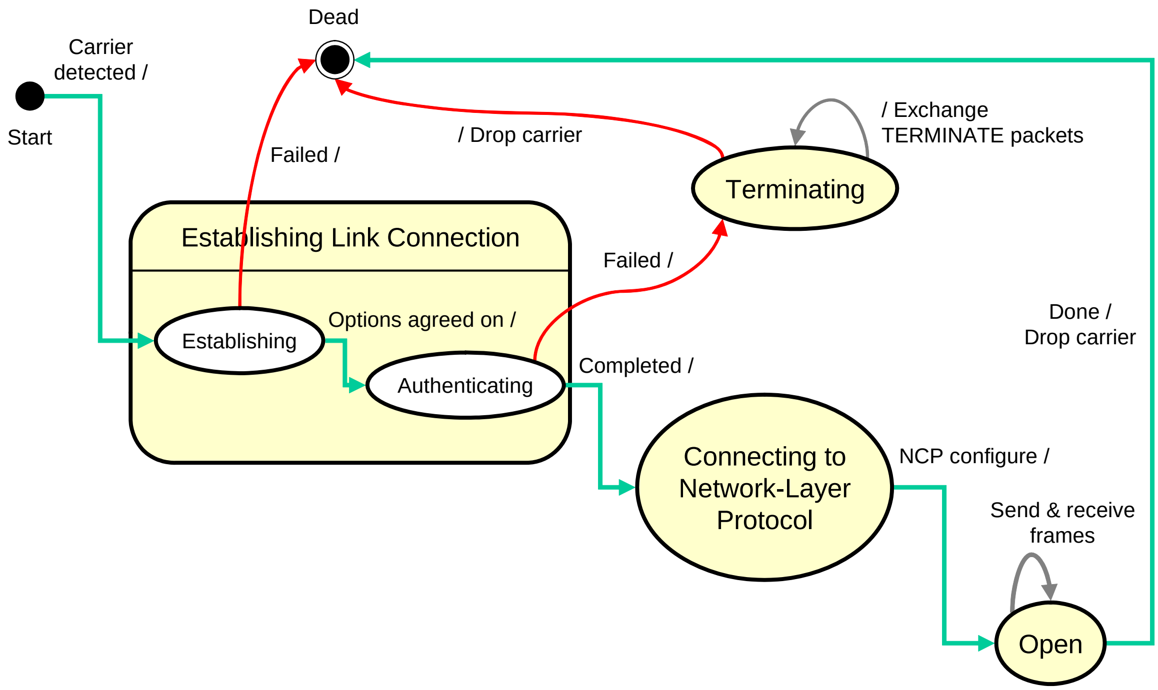

Figure 1-57: State diagram for the point-to-point protocol (PPP).

If transition through these two states is successful, the connection goes to the Open state, where data transfer between the endpoints takes place.

The Authenticating state (sub-state of Establishing Link Connection) is optional. The two endpoints may decide, during the Establishing sub-state, not to go through authentication. If they decide to proceed with authentication, they will exchange several PPP control frames.

Listing 1-1 in Section 1.1.4 shows how application calls the protocols down the protocol stack when sending a packet. However, before `send()` can be called, `lowerLayerProtocol` must be initialized. Link-layer protocol is usually built-in in the firmware of the network interface card, and the initialization happens when the hardware is powered up or user runs a special application. Therefore, NCP in step 2 above establishes the connection between the link-layer PPP protocol and the higher-layer (e.g., IP) protocol that will use its services to transmit packets.

LCP and NCP protocols send control messages encapsulated as the payload field in PPP frames (Figure 1-58). The receiving PPP endpoint delivers the messages to the receiving LCP or NCP module, which in turn configures the parameters of the PPP connection.

Although PPP frames do not use link-layer addresses, PPP provides the capability for network-layer address negotiation: endpoint can learn and/or configure each other's network address.

In summary, PPP has no error correction/recovery (only error detection), no flow control, and out-of-order delivery is acceptable. No specific protocol is defined for the physical layer in PPP.

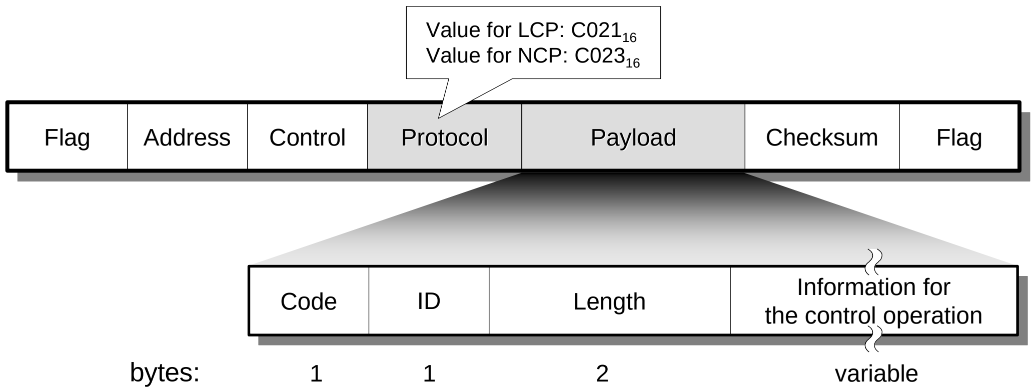

Figure 1-58: LCP or NCP packet encapsulated in a PPP frame.

1.5.2 Ethernet (IEEE 802.3)

Problems related to this section: Problem 1.41 → Problem 1.43

Ethernet is a network protocol for local area networks (LANs). The MAC protocol for Ethernet is based on the CSMA/CD protocol shown in Figure 1-30. The frame format for Ethernet is shown in Figure 1-59. The Ethernet was first standardized by DEC, Intel and Xerox, known as the DIX standard. (See Section 1.7 for an overview of Ethernet history.) When IEEE released the 802.3 standard, it adopted a slightly different frame format, as shown in Figure 1-59(b). The Type field in a DIX frame represents the upper-layer protocol that is using Ethernet as its link layer. On the other hand, an 802.3 frame carries instead the frame Length. In 802.3 frame, the upper-layer protocol is specified in the LLC frame as DSAP address (also see Figure 1-54). Because the DIX standard was widely used by the time IEEE 802.3 was released, a compromise is reached as follows. If the Type/Length field contains a number ≤1500 than it represents the frame Length and the receiver should look for the upper-layer protocol in the contained LLC packet. If the Type/Length field contains a number >1500 than it identifies the upper-layer protocol, and the data field does not contain an LLC-formatted packet, but rather a network-layer packet (e.g., an IP datagram). All versions of Ethernet up to date use this frame format.

The **Ethernet link-layer address** or **MAC-48 address** is a globally unique 6-byte (48-bit) string that comes wired into the electronics of the Ethernet attachment. An Ethernet address has two parts: a 3-byte manufacturer code, and a 3-byte adaptor number. IEEE acts as a global authority and assigns a unique manufacturer's registered identification number, while each manufacturer gives an adaptor a unique number. Although intended to be a permanent and globally unique identification, it is possible to change the MAC address on most of today's hardware, an action often referred to as *MAC spoofing*.

When a frame arrives at an Ethernet attachment, the electronics compares the destination address with its own and discards the frame if the addresses differ, unless the address is a special "broadcast address" which signals that the frame is meant for all the nodes on this network.

We know from Section 1.3.3 that for the CSMA/CD protocol, the transmission time of the smallest frame must be larger than one round-trip propagation time, i.e., 2β. This requirement limits the distance between two computers on an Ethernet LAN. The smallest frame is 64 bytes.

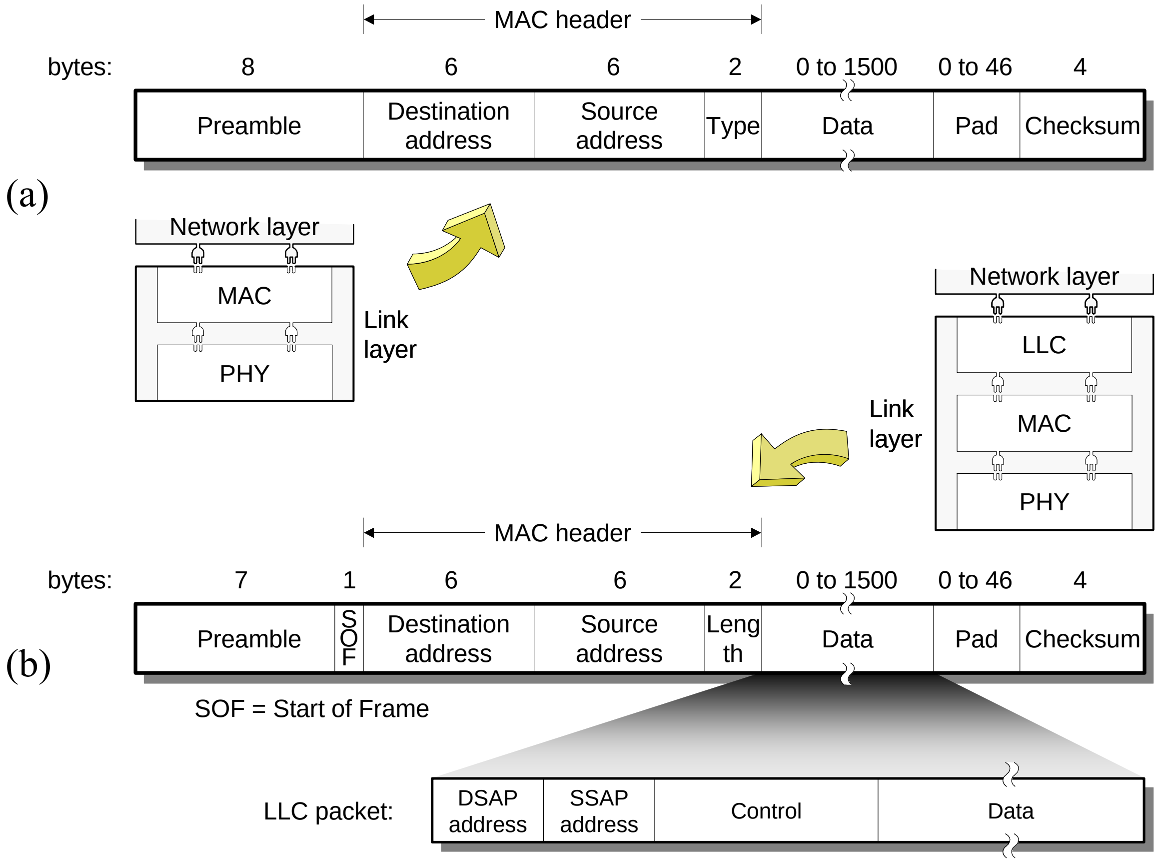

Figure 1-59: Link-layer frame format for DIX standard Ethernet Version 2.0 (a) and for IEEE standard 802.3 (b). LLC packet format is shown in Figure 1-54.

This 64-byte value is derived from the original 2500-m maximum distance between Ethernet interfaces plus the transit time across up to four repeaters plus the time the electronics takes to detect the collision. The 64 bytes correspond to 51.2 μs over a 10 Mbps link, which is larger than the round-trip time across 2500 m (about 18 μs) plus the delays across repeaters and the electronics to detect the collision.

Sensing the medium idle takes time, so there will necessarily be an idle period between transmissions of Ethernet frames. This period is known as the **interframe space** (IFS), interframe

Table 1-4: Parameter values for the Ethernet MAC protocol (CSMA/CD).

Parameter	**Data rate**		
	Up to and including 100 Mbps	1 Gbps	10 Gbps
Backoff slot time	512 bit times	4096 bit times	not applicable
Interpacket gap / IFS	96 bits	96 bits	96 bits
Attempts limit	16	16	not applicable
Backoff limit	10	10	not applicable
Jam size	32 bits	32 bits	not applicable
Maximum frame size	1518 bytes	1518 bytes	1518 bytes
Minimum frame size	512 bits (64 bytes)	512 bits (64 bytes)	512 bits (64 bytes)

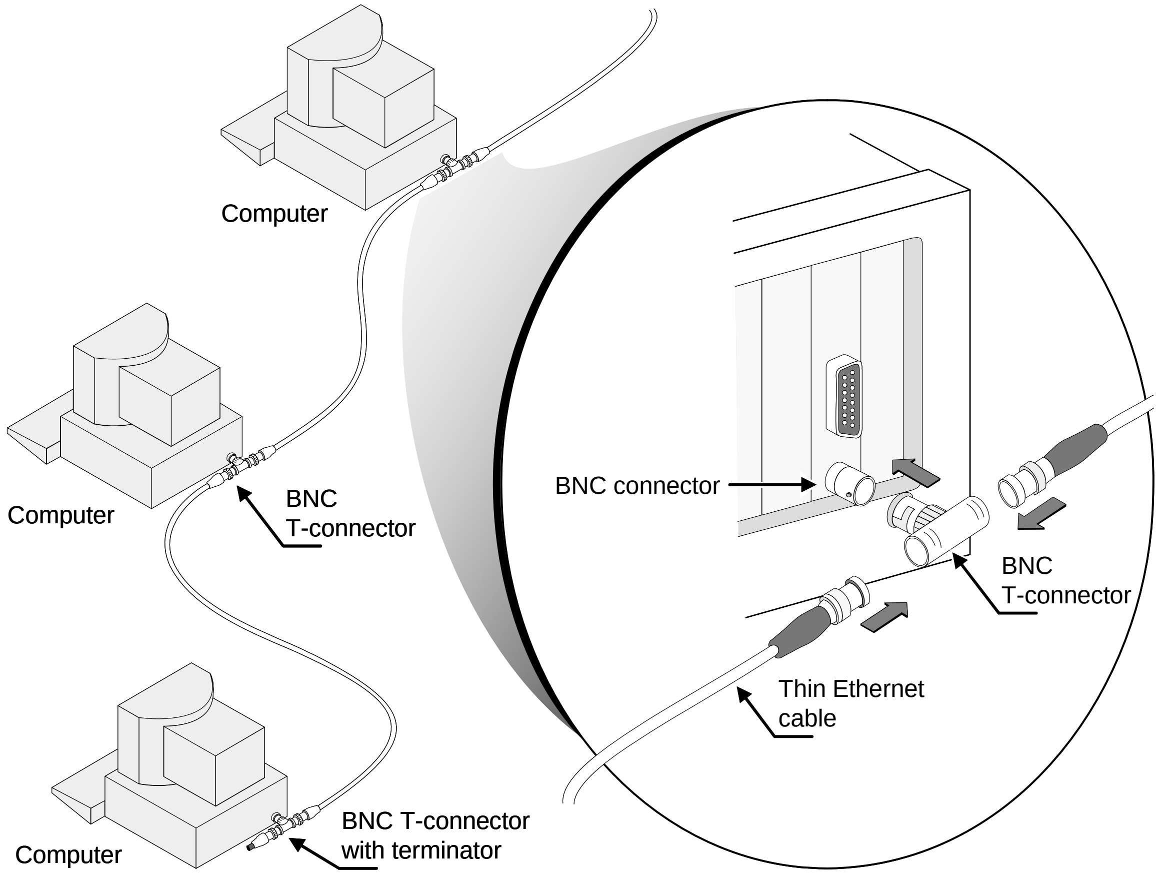

Figure 1-60: Thin coaxial cable Ethernet represents a bus-based design.

gap, or interpacket gap. It is the spacing between two non-colliding frames, from start of idle after the last bit of the FCS field of the first frame to the first bit of the Preamble of the subsequent frame. In other words, if an Ethernet network adapter senses that there is no signal energy entering the adapter from the channel for IFS-bit times, it declares the channel idle and starts to transmit the frame. The minimum interframe space is 96-bit times (the time it takes to transmit 96 bits of raw data on the medium), which is 9.6 μs for 10 Mbps Ethernet, 960 ns for 100 Mbps (fast) Ethernet, 96 ns for 1 Gbps (gigabit) Ethernet, and 9.6 ns for 10 Gbps (10 gigabit) Ethernet.

The Ethernet specification for a bus-based design allows no more than 1,024 hosts and it can span only a geographic area of 2,500 m. Table 1-4 lists some important parameters of the Ethernet MAC protocol for different data rates of the physical sublayer.

Evolution of Ethernet

Ethernet has evolved over the past 35 years since it was invented. This evolution was shaped by physical characteristics of communication links, such as *data rate*, *duplexity* (half or full duplex), and *multiplicity* of the medium (i.e., point-to-point or shared broadcast). Ethernet operation is specified for data rates from 1 Mbps to 10 Gbps using a common MAC protocol (CSMA/CD). In 1997, IEEE Std 802.3x specified *full duplex* operation. The CSMA/CD MAC protocol specifies shared medium (half duplex) operation where frame collisions can occur, as well as full duplex

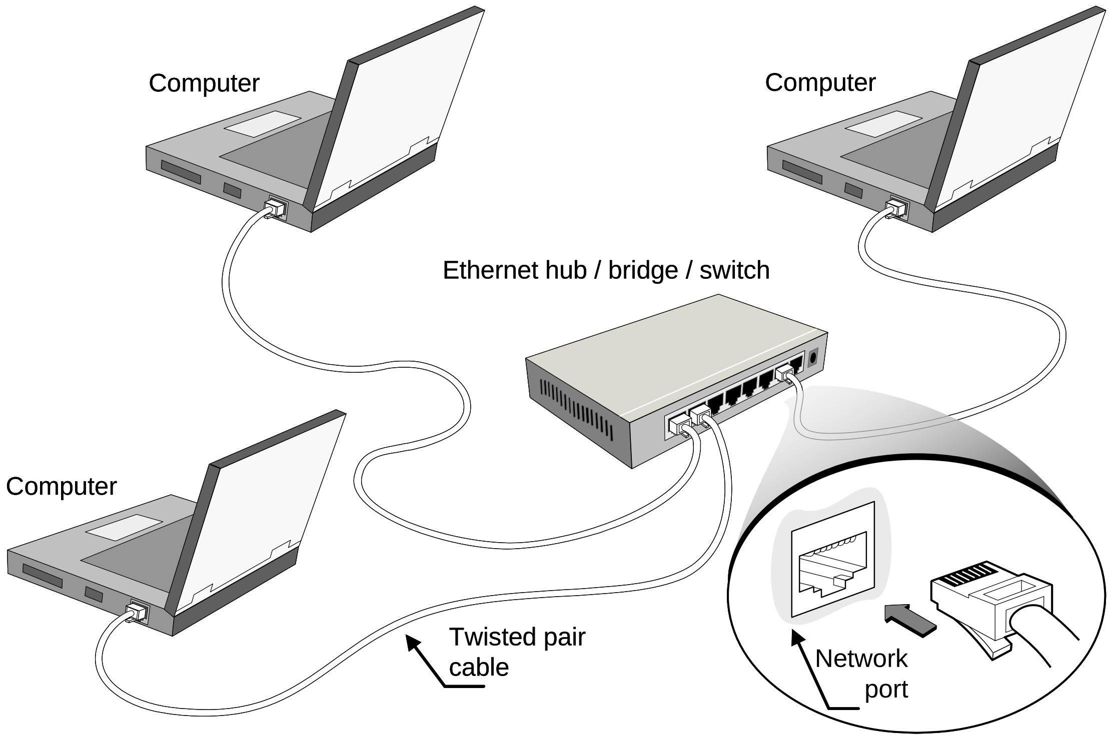

Figure 1-61: Bridged or switched Ethernet represents a star-based (hub-and-spokes) design.

operation that operates without collisions. Ethernet Physical Sublayer (PHY) is standardized for operation over coaxial, twisted-pair or fiber-optic cables.

Ethernet was first standardized for operation over coaxial cables (Figure 1-60). A cable (ether) with multiple devices attached to it in parallel is called a *multidrop cable*. This is also known as **bus**-based design for Ethernet. The multidrop cable with all stations attached to it are called a **collision domain**. If two or more stations in a collision domain transmit simultaneouly, their frames will collide and will not be successfully received. First appeared the so-called *Thick Ethernet* (or, 10Base5) which used a thick coaxial cable with markings to show where transcievers can be screwed onto the cable (2.5 meters apart). The second cable type was *Thin Ethernet* (or, 10Base2), which used standard BNC connectors to form T-junctions on the carrier cable (Figure 1-60). Multidrop-cable Ethernets were followed by a star-patterned wiring, where all computers in the LAN have a cable running to a central **hub** and incident spokes (Figure 1-61). Historically, the first instance of this design is 10Base-T. The Ethernet version notation consists of three parts, as follows:

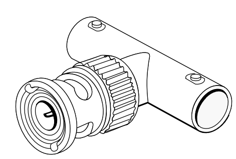

Data rate (e.g., 10 Mbps, 10 Gbps)	Baseband/Broadband transmission	Wiring type (e.g., coaxial, twisted pair or fiber optic)

For example, 10Base-T means 10 Mbps baseband transmission over unshielded twisted-pair cable (Category 5 UTP); 10GBase-X means 10 Gbps baseband over two pairs of twisted-pair cable.

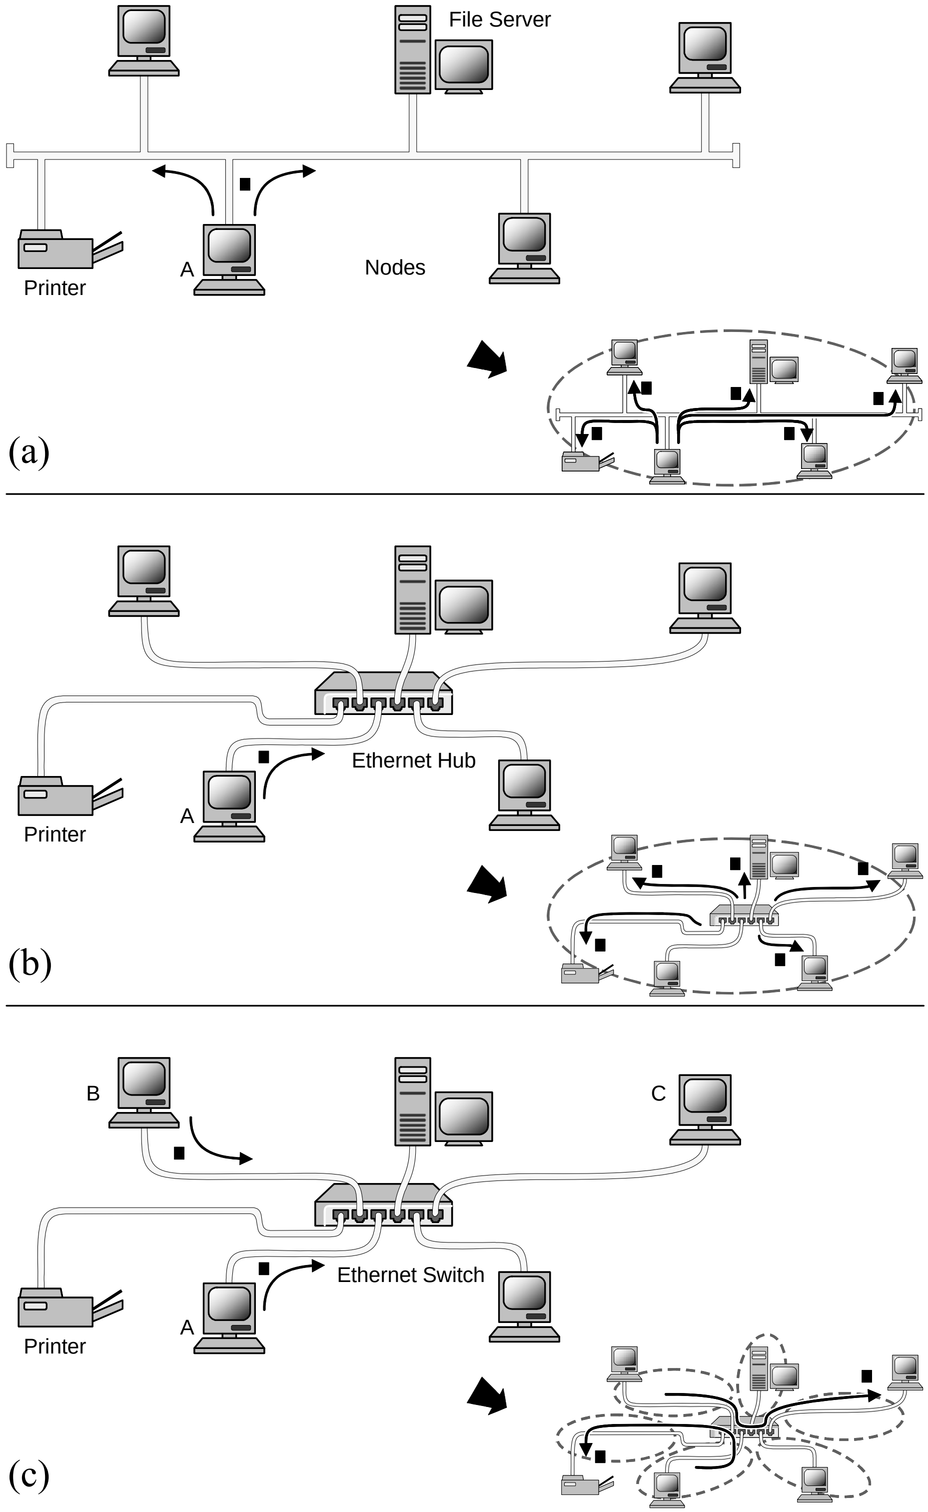

Figure 1-62: Comparing bus-based multidrop-cable Ethernet (a), hub-based Ethernet (b) and switch-based Ethernet (c). Dotted ovals indicate independent collision domains.

The star design in Figure 1-61 has many variations, depending on whether the central device operates at the physical layer (OSI Layer 1) or at the link layer (OSI Layer 2) and whether the

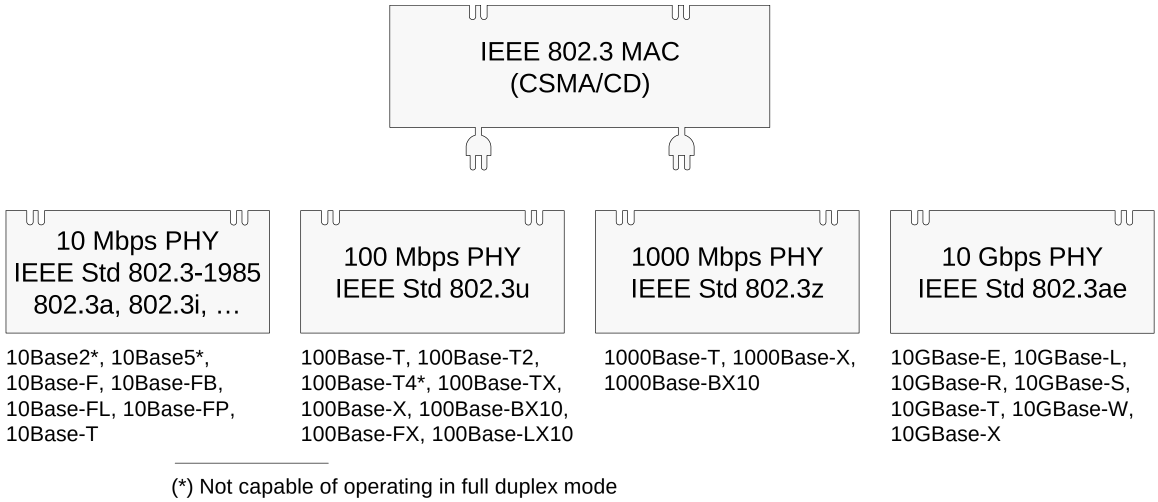

Figure 1-63: Ethernet standards family for IEEE Std 802.3-2008 (current).

links are half duplex or full duplex. The central device was historically first called **bridge**. In the simplest version, the bridge is called **hub** or **repeater** and it operates at the physical layer in a half-duplex mode. A hub does not understand anything beyond bits, i.e., does not recognize frames or knows about device addresses. It simply switches bits that come in one network interface (or, port) to *all* other interfaces. The whole network forms a *single collision domain*, so conceptually this design is equivalent to a bus-based design.

A more spohisticated bridge is known as a **switch**, but the term “bridge” is also often used synonymously. An Ethernet switch moves frames from input to output ports based on their Layer-2 destination addresses (described above as MAC-48 addresses). In other words, unlike a hub which switches bits to all interfaces, a switch switches a frame exclusively to a port determined by the frame’s destination address.

Figure 1-62 illustrates the difference between hubs and switches. Ethernet hubs (Figure 1-62(b)) are conceptually equivalent to the bus-based Ethernet (Figure 1-62(a)) because both designs form a single collision domain. Conversely, each network port of an Ethernet switch forms an *independent collision domain* (Figure 1-62(c)). With switch-based design, each cable has only two stations attached: on one end is a switch’s port and on the other end is a computer host. This is essentially a point-to-point link, but collisions are still possible between the endpoints and CSMA/CD must be employed. More detail on Ethernet switches is provided later in this section.

Figure 1-63 summarizes the current family of Ethernet protocols. The figure also indicates whether the physical sublayer has the ability to perform full-duplex link transmission and reception, which is described next.

Full-duplex Mode and Collision-free Ethernet

Traditionally, Ethernet MAC sublayer implements the CSMA/CD algorithm, which creates a half-duplex link. In half-duplex mode, media access method is the means by which two or more stations share a common transmission medium (broadcast). To transmit, a station waits (defers) for a quiet period on the medium (that is, no other station is transmitting) and then sends the intended message in bit-serial form. If, after initiating a transmission, the message collides with that of another station, then each transmitting station intentionally transmits for an additional predefined period to ensure propagation of the collision throughout the system. The station remains silent for a random amount of time (backoff) before attempting to transmit again.

Ethernet 802.3 standard provides for two modes of operation of the MAC sublayer:

(a) In **half-duplex mode**, stations contend for the use of the physical medium, using the CSMA/CD algorithms specified. This is the traditional CSMA/CD contention-based operation. Bidirectional communication is accomplished by sequential exchange of frames, rather than simultaneous transmission in both directions. Half-duplex operation is possible on all supported media; it is required on those media that are incapable of supporting simultaneous transmission and reception without interference, such as 10Base2 and 100Base-T4 (Figure 1-63).

(b) The **full-duplex mode** of operation allows simultaneous communication between a pair of stations using point-to-point media (dedicated channel). Full-duplex operation does not require that transmitters defer, nor do they monitor or react to receive activity ("collision detection"), as there is no contention for a shared medium in this mode. Full-duplex operation can be used when *all* of the following are true:

1) The physical medium is capable of supporting simultaneous transmission and reception without interference (Figure 1-63).

2) There are exactly two stations connected with a full duplex point-to-point link. Because there is no contention for use of a shared medium, the multiple access (i.e., CSMA/CD) algorithms are unnecessary.

3) Both stations on the LAN are capable of, and have been configured to use, full duplex operation.

The most common configuration envisioned for full-duplex operation consists of a central switch (or, bridge) with a dedicated LAN connecting each switch port to a single station. Ethernet hubs or repeaters are outside the scope of full duplex operation. By definition, an IEEE 802.3 LAN operating in full-duplex mode comprises exactly two stations, so full-duplex mode creates an **Ethernet point-to-point link**.

An Ethernet device operates in either half or full duplex mode at any one time. A device is configured for one specific mode of operation (e.g. 1000Base-X Full Duplex). Auto-Negotiation is performed as part of the initial set-up of the link, and allows the PHYs at each end to advertise their capabilities (speed, PHY type, half or full duplex) and to automatically select the operating mode for communication on the link. The term "CSMA/CD MAC" is used synonymously with "802.3 MAC," and may represent an instance of either a half duplex or full duplex mode device, although full-duplex devices do not implement the traditional CSMA/CD algorithm. In full-duplex mode, stations do not implement the CSMA/CD algorithms traditionally used to arbitrate access to shared-media LANs. Full-duplex operation constitutes a proper subset of the MAC functionality required for half-duplex operation.

The current Ethernet standard (IEEE Std 802.3-2008)

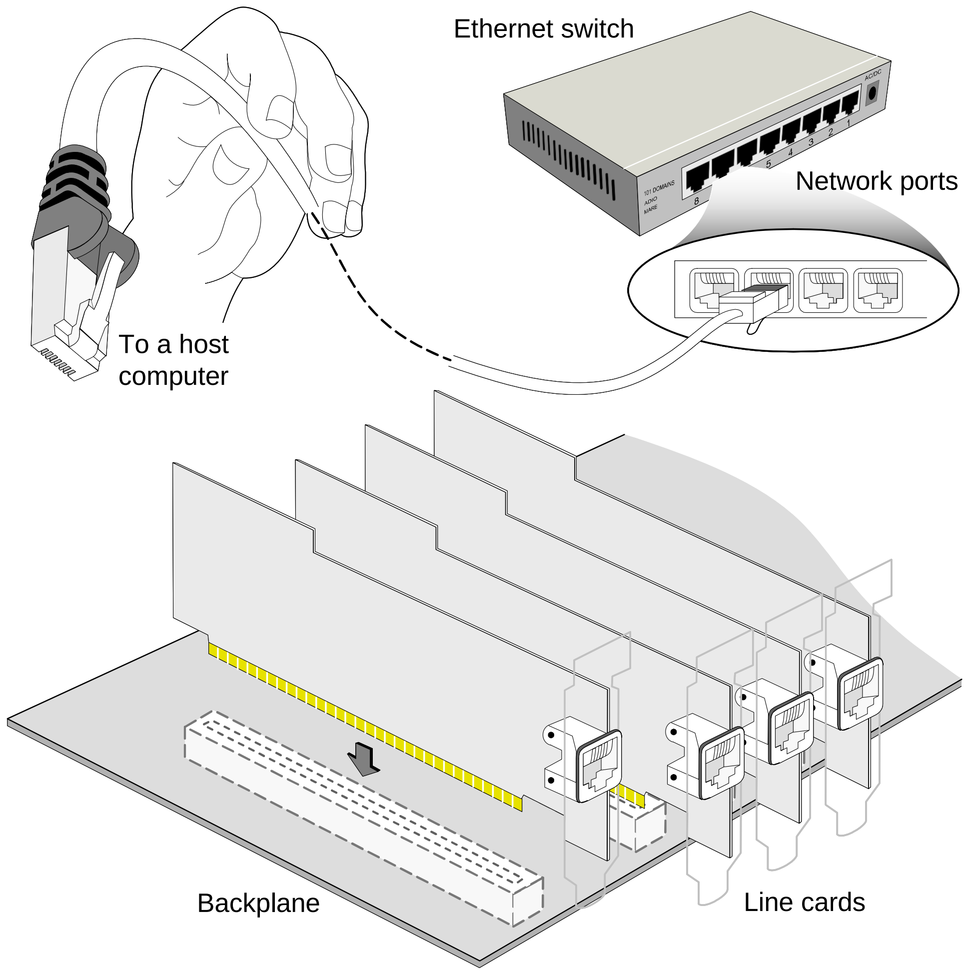

Figure 1-64: Ethernet switch architecture.

Ethernet Switches

An Ethernet switch consists of a high-speed backplane and a number of plug-in line cards, typically 4 to 32 (Figure 1-64). Each line card contains one or more (e.g., eight) network ports or connectors. A twisted pair cable leads from each connector to a host computer. When a computer sends a frame, the frame first reaches an associated line card, which checks whether the frame is destined to a station connected to the same card. If so, the frame is copied to the given port/connector on this line card. If not, the frame is sent over the backplane to the destination computer's line card. The backplane typically runs at data rates of many Gbps, using a proprietary protocol. More about switch design is available in Section 4.1.

A hub or repeater transmits a frame on an output port while it is being received on an input port. This is known as *cut-through switching*. Unlike a hub/repeater, a switch or bridge first receives the entire frame then stores it, waiting for the network attached to the frame's outgoing port to become idle. This is known as *store-and-forward switching*. With store-and-forward switching, it is possible for two stations on different ports of the switch to transmit simultaneously without a collision. We say that switch ports form independent collision domains (Figure 1-62).

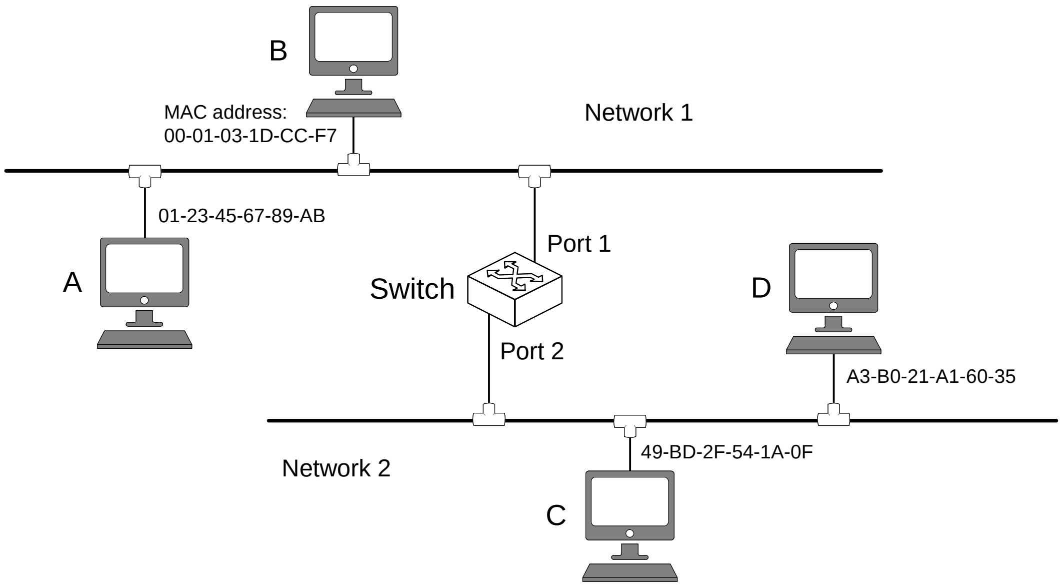

Figure 1-65: Example Ethernet networks connected by an Ethernet switch.

As already pointed out, switches switch packets based on their link-layer (or, MAC) addresses, i.e., switches operate at OSI Layer-2. They are also known as *LAN switches*. In Section 1.4 we also learned about another kind of switches: routers. Routers switch packets based their network-layer addresses, i.e., switches operate at OSI Layer-3. Routers are more complex because they need to run routing protocols to discover the topology of the entire internetwork consisting of many networks. It is also said that LAN switches are *transparent* to the computers in the network. Unlike routers, where nodes know of the next-hop routers, LAN nodes are unaware of intermediate switches and their forwarding role. When a computer sends a frame, the frame is addressed to another computer, rather than addressing the frame to a switch. The frame will pass through a switch when going from one LAN segment to another without the switch identifying itself as the device that transmitted the frame to the next segment. Therefore, switches are transparent to each other, as well. Routers are described in Section 4.1.

LAN switches perform two basic functions: frame forwarding and frame filtering. **Frame forwarding** helps move a frame toward its ultimate destination. A switch moves a frame from an input port to an output port based on frame's MAC address by looking up the **switching table**. The switching table is similar to a router's forwarding table. Consider a network in Figure 1-65. The switching table of the switch is shown in Table 1-5. The MAC addresses of stations *A*, *B*, and *D* are listed in the table. For example, if a frame arrives on Port-2 destined for MAC address 00-01-03-1D-CC-F7 (station *A*), the switch outputs the frame on Port-1. If a frame arrives on Port-1 destined for 49-BD-2F-54-1A-0F (station *C*), which is currently not listed in the switching table, the switch will output the frame to all other ports. In this example, Port-2 is the only other port.

Table 1-5: The switching table for the example LAN in Figure 1-65.

MAC address	Network port	Time last frame received
00-01-03-1D-CC-F7	1	10:39
01-23-45-67-89-AB	1	10:52
A3-B0-21-A1-60-35	2	10:17

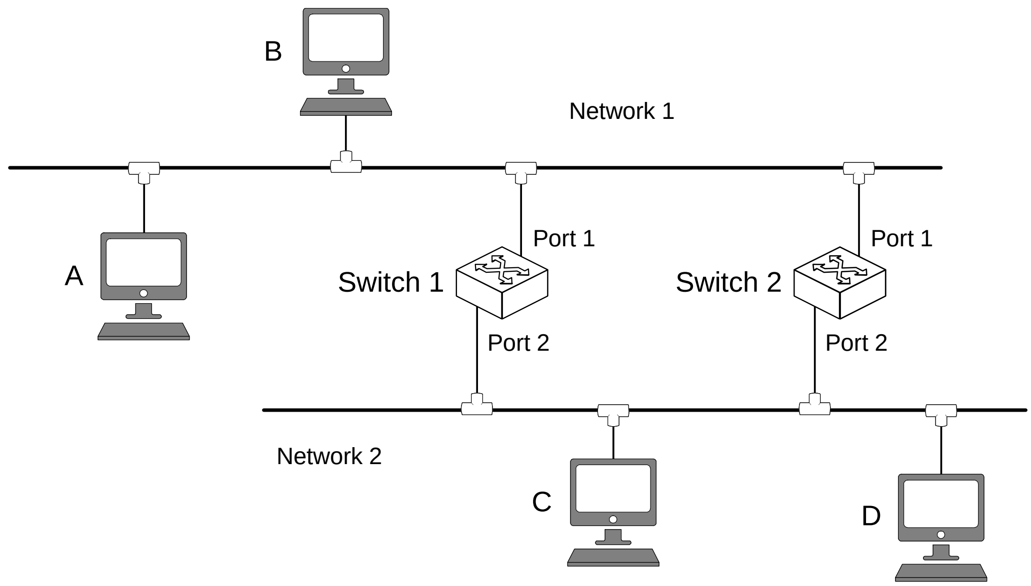

Figure 1-66: Example switched network with a loop formed by two switches.

Frame filtering relates to discarding frames that are headed in a direction where they do not need to go. For example, if in Figure 1-65 a frame arrives on Port-2 with the destination address 00-01-03-1D-CC-F7 (station *A*), then according to Table 1-5 this frame should be output on Port-1. On the other hand, assume that a frame arrives on Port-2 with the destination address A3-B0-21-A1-60-35. The switch realizes that it is a station on a network segment attached on Port-2 sending a frame to another station on the same segment. In our case, it is station *C* sending a frame to station *D*. There is no need to forward this frame because all other stations on the same segment already received this frame, so the switch filters this frame and discards it.

The switching table can be filled up manually, but this is a tedious and error-prone task for large number of stations. Instead, the switch performs **backward learning** of the switching table. Initially, the table is empty. When the switch receives a frame from a station for which it has no address in the table, the switch automatically creates a new entry. The entry records the MAC address in the frame's *Source address* field (Figure 1-59), the network port on which the frame arrived, and the time of the arrival. If every station on all attached networks sends a frame, then every station will eventually be recorded in the table. The parameter called **aging time** determines how long the table entries are valid. If the switch does not receive a frame with a given address as its source address, the entry will be deleted from the table. For example, in Table 1-5 a frame with source address A3-B0-21-A1-60-35 (station *D* Figure 1-59) arrived last time at 10:17 on Port-2. Suppose that the aging time for this switch is 50 minutes. If no frame arrives with source address A3-B0-21-A1-60-35 arrives until 11:07, the switch will remove the entry for station *D* from the table. In this way, if a computer is unplugged or moved around the building and plugged in again somewhere else, the network can operate without manual intervention.

Consider now the network in Figure 1-66. The two switches connect the two networks via two alternative paths, thus forming a *loop* (or, cycles) in the topology. This may happen by accidence, if the network administrator is not careful when upgrading the network, or it may be done purposefully to provide for alternate paths in case of switch failures (fault tolerance by redundancy). Let us assume that the switching tables of both switches are as in Table 1-5 and that

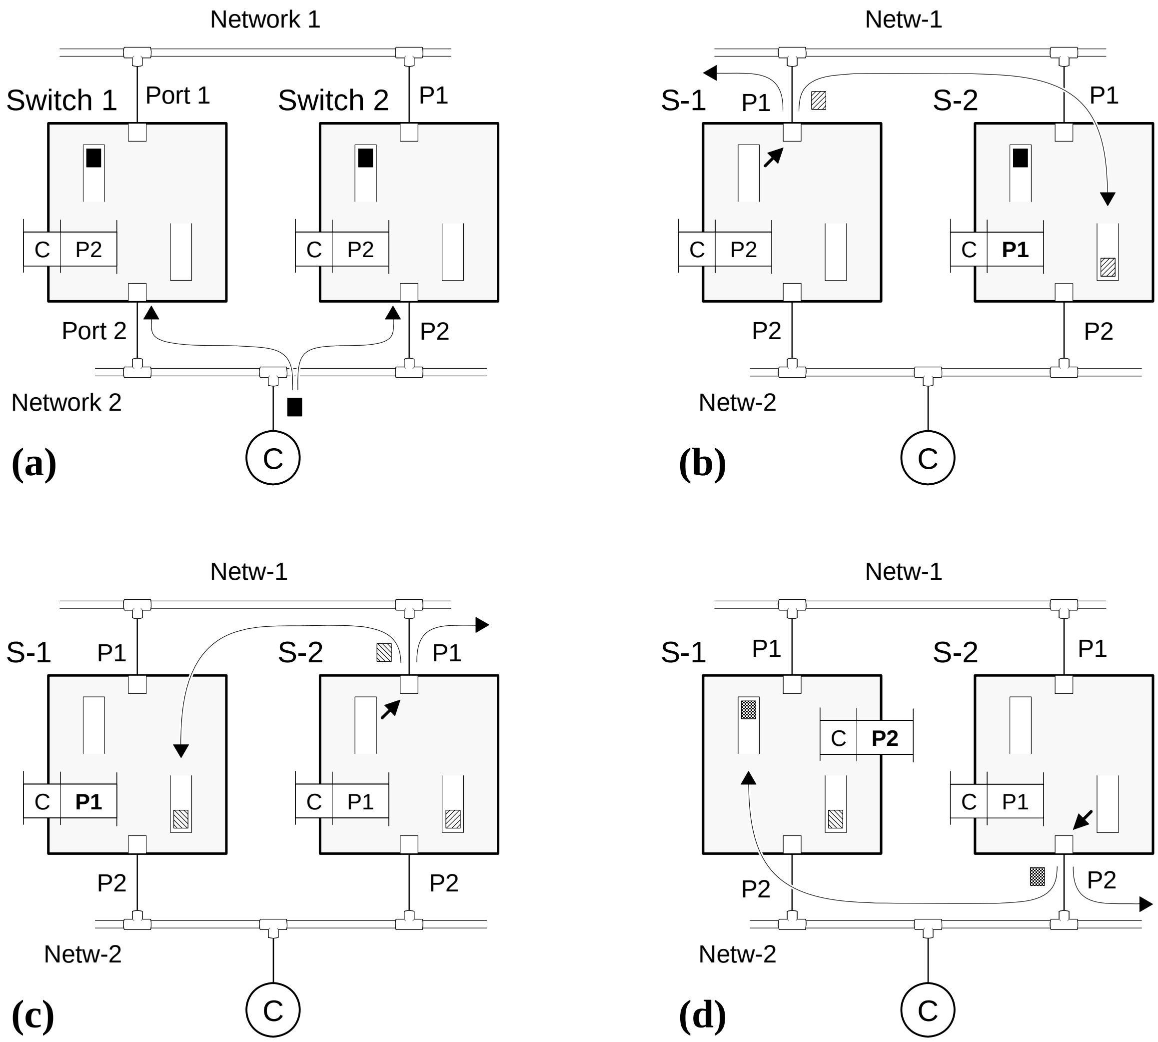

Figure 1-67: Packet proliferation in the network with a loop in Figure 1-66.

station *C* sends a frame to another station (the destination to which *C* sends is irrelevant for this example). The frame will arrive on both switches on Port-2, and each switch will record in its switching table that *C* arrived on Port-2 (resides on Network 2) and enqueue the frame for forwarding on its Port-1 (Figure 1-67(a)). Let us say that Switch 1 is the first to seize the access to the medium and succeed in relaying the frame to Network 1. Because switches are transparent to each other, the frame will appear on Port-1 of Switch 2 exactly as if transmitted by station *C*. Switch 2 will record in its table that *C* arrived on Port-1 (as if *C* now resides on Network 1!) and enqueue the frame for forwarding on its Port-2 (Figure 1-67(b)). Next, suppose that Switch 2 succeeds in transmitting its first received frame onto Network 1 (Figure 1-67(c)). Switch 1 will record that *C* moved to Port-1 and enqueue the frame on its Port-2. Figure 1-67(d) shows one more iteration, where Switch 2 transmits its second received frame onto Network 2, but this process continues to infinity. Notice also that during this process any frames from other stations heading to station *C* may be misdirected and eventually discarded.

The solution to this problem is to remove the loops, which produces a tree from a general graph. A **spanning tree** of a graph is a subgraph of this graph that connects (spans) all the nodes, but contains no cycles. That is, a spanning tree keeps all the nodes of the original graph, but removes some links. In terms of Ethernet networks, each LAN segment corresponds to a graph node, and each switch corresponds to a link in the graph. A spanning tree of a network can be derived automatically using the *spanning tree protocol* (STP), specified in the IEEE 802.1D standard.

Each switch in the network sends a configuration message on all of its attached networks, which includes the MAC address of the switch. The switches use the Spanning Tree Protocol to compute the spanning tree, which has these five steps:

1. Elect a root switch. The switches choose one switch as the root switch of the spanning tree. The choice is the switch with the smallest (lowest) identifier. Each switch has a unique identifier (its MAC address) and a configurable priority number; the switch ID contains both numbers. To compare two IDs, the priority is compared first. If two switches have equal priority, then their MAC addresses (48-bit binary numbers) are compared and the switch with the smaller address is chosen as the root switch. Of course, before configuration messages from all switches are received, some switches may have made incorrect choices due to insufficient information. The root switch always forwards frames out over all of its ports.

2. Compute the shortest path to the root. Each switch determines the cost of each possible path from itself to the root. From these paths, it selects one with the smallest cost (shortest path). The port connecting to that path becomes the *root port* of the switch. The cost of traversing a path is the sum of the costs of the LAN segments on the path. Different technologies have different default costs for LAN segments. A common approach is to assign to each segment the cost of 1 (i.e., one hop). All shortest paths form a spanning tree.

3. Determine any designated ports. All switches on a LAN segment collectively decide which one among them has the shortest path to the root. The elected switch becomes the *designated switch* that will be responsible for forwarding frames from this LAN segment toward the root switch. The port connecting the designated switch to the given LAN segment becomes a *designated port* of the switch. A switch may have no designated ports or may have more than one designated port (because each switch is connected to several LAN segments).

4. Disable all other ports. Every switch *blocks* all of its active ports that do not qualify as a root port or a designated port. In case there are ties, go to the next step.

5. Resolve the ties. It may happen that two or more ports on a single switch are attached to shortest paths to the root or two or more switches on the same LAN segment have equal least-cost paths to the root. Such ties are broken as follows:

5.a) Breaking ties for root ports. When multiple paths from a switch are shortest paths, the chosen path uses the neighbor switch with the lower identifier. The root port is thus the one connecting to the switch with the lowest identifier.

5.b) Breaking ties for designated ports. When more than one switch on a segment has a shortest path to the root, the switch with the smaller identifier is chosen to forward messages to the root. The port attaching that switch to the LAN segment is a designated port of that switch. A loser switch sets the port to the given LAN segment as being blocked.

5.c) The final tiebreaker. In some cases, there may still be a tie, as when two switches are connected by multiple cables. In this case, multiple ports on a single switch are candidates for root port. The path that passes through the port on the neighbor switch that has the lowest port priority is used.

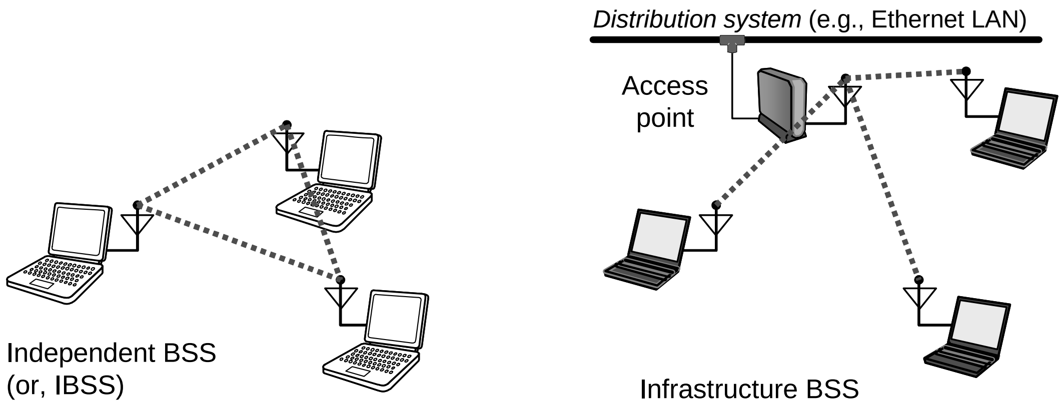

Figure 1-68: IEEE 802.11 (Wi-Fi) independent and infrastructure basic service sets (BSSs).

The switches run the STP protocol iteratively and exchange configuration messages containing this information:

(1) The identifier for the switch sending the message (includes the MAC address and a configurable priority number);

(2) The identifier for what the sending switch believes to be the root switch;

(3) The distance (measured in hops) from the sending switch to the root switch.

Initially, each switch thinks it is the root, and so it sends a configuration message out on each of its ports identifying itself as the root, with a distance to the root valued at 0. When a switch receives a message on a particular port, the switch checks if this message is better than the best configuration message previously recorded for this port. The message is considered “better” if:

- It identifies a root with a smaller identifier, or
- It identifies the same root but with a shorter distance (lower cost path), or
- It identifies the same root and distance, but the sending switch has a smaller identifier.

When a switch decides that it is not the root switch, it stops sending own configuration messages and only forwards messages from other switches. Similarly, when a switch decides that it is not the designated switch for a given LAN segment, it stops sending configuration messages over the port attached to this segment. When the system stabilizes, only the root switch will be generating configuration messages and all the other switches will be forwarding these messages only over the ports for which they are the designated switch.

1.5.3 Wi-Fi (IEEE 802.11)

Problems related to this section: Problem 1.44 → Problem 1.45

IEEE 802.11, also known as Wi-Fi, …

Architecture and Basics

The basic building block of an IEEE 802.11 network is the **basic service set** (BSS), which is simply a set of stations that communicate with one another. A BSS does not generally refer to a

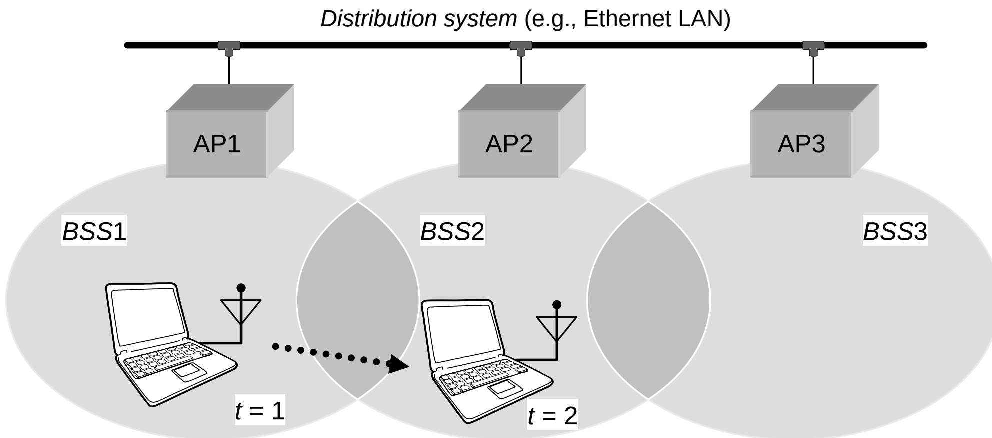

Figure 1-69: IEEE 802.11 (Wi-Fi) extended service set (ESS) allows connecting multiple access points to support long-range roaming.

particular area, due to the uncertainties of electromagnetic propagation. There are two types of BSS, as shown in Figure 1-68. When all of the stations in the BSS are mobile stations and there is no connection to a wired network, the BSS is called an **independent BSS** (or, IBSS). The IBSS is the entire network and only those stations communicating with each other in the IBSS are part of the LAN. This type of network is called an **ad hoc network** (see Chapter 6).

When all of the mobile stations in the BSS communicate with an **access point** (AP), the BSS is called an **infrastructure BSS** (never called an IBSS!). This configuration is also known as *wireless local area network* or W-LAN. The access point provides both the connection to the wired LAN (wireless-to-wired bridging), if any, and the local relay function for all stations in its BSS. Therefore, if one mobile station in the BSS must communicate with another mobile station, the packet is sent first to the AP and then from the AP to the other mobile station. This causes communications to consume more transmission capacity than in the case where the communications are directly between the source and the destination (as in the IBSS). However, in many cases the benefits provided by the AP outweigh the drawbacks. One of the benefits provided by the AP is that the AP can assist the mobile stations in saving battery power. The mobile stations can operate at a lower power, just to reach the AP, and not worry about how far away is the destination host. Also, the AP can buffer (temporarily store) the packets for a mobile station, if the station is currently in a power saving mode.

Extended service set (ESS) extends the range of mobility from a single infrastructure BSS Figure 1-68(b) to an arbitrary range by interconnecting a set of infrastructure BSSs (Figure 1-69). In ESS, multiple APs communicate among themselves to forward traffic from one BSS to another and to facilitate the roaming of mobile stations between the BSSs. This is conceptually similar to the cellular telephony network. The APs perform this communication via the *distribution system*, such as an Ethernet-based wireline network. The stations in an ESS see the wireless medium as a single link-layer connection. ESS is the highest-level abstraction supported by 802.11 networks. Roaming between different ESSs is not supported by IEEE 802.11 and must be supported by a higher-level protocol, e.g., Mobile IP (Section 8.3.4).

Wi-Fi supports dynamic data-rate adaptation to the current conditions of the wireless channel (Figure 1-70). The goal is to select the rate that minimizes the errors due to the channel noise. This behavior is not implemented in Ethernet, which operates over wire media, where channel

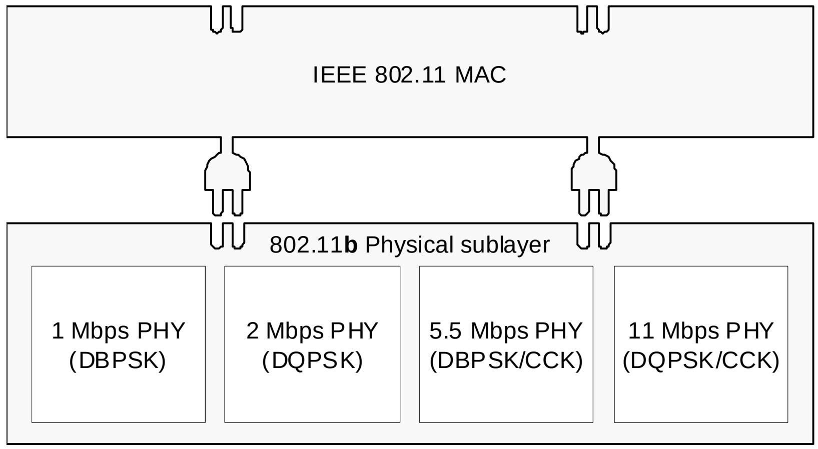

Figure 1-70: Wi-Fi supports dynamic data-rate adaptation at the physical sublayer. This figure shows the available rates for 802.11b.

conditions are stable and error rate is low. In Ethernet, the physical data rate is pre-configured and does not change at runtime (compare to Figure 1-63).

Figure 1-71 shows the 802.11 frame format. The general MAC-layer format (Figure 1-71(a)) is used for all data and control frames, but not all fields are used in all types of frames. There can be up to four address fields in an 802.11 frame. When all four fields are present, the address types include source, destination, transmitting station, and receiving station. The first two represent the end nodes and the last two may be intermediary nodes. 802.11 uses the same MAC-48 address format as Ethernet (Section 1.5.2). One of the fields could also be the BSS identifier, which is used in the probe request and response frames, used when mobile stations scan an area for existing 802.11 networks.

The *Duration*/*Connection-ID* field indicates the time (in microseconds) the channel will be reserved for successful transmission of a data frame. The stations that receive this frame, but are not intended receivers, use this information to defer their future transmissions until this transmission is completed. The deferral period is called **network allocation vector** (NAV), and we will see later in Figure 1-77 how it is used. In some control frames, this field contains a network association, or connection identifier.

The 802.11 physical-layer frame (Figure 1-71(b)) is known as PLCP protocol data unit (PPDU), where PLCP stands for "physical (PHY) layer convergence procedure." The version shown in Figure 1-71(b) is known as *Long PPDU format*. A **preamble** is a bit sequence that receivers watch for to lock onto the rest of the frame transmission. There are two different preamble and header formats defined for 802.11 physical-layer frames. The mandatory supported long preamble and header, shown in Figure 1-71(b), is interoperable with the basic 1 Mbps and 2 Mbps data transmission rates. There is also an optional short preamble and header (not illustrated here), known as *Short PPDU format*. This format is used at higher transmission rates to reduce the control overhead and improve the network performance. (More discussion is provided in Chapter 6.)

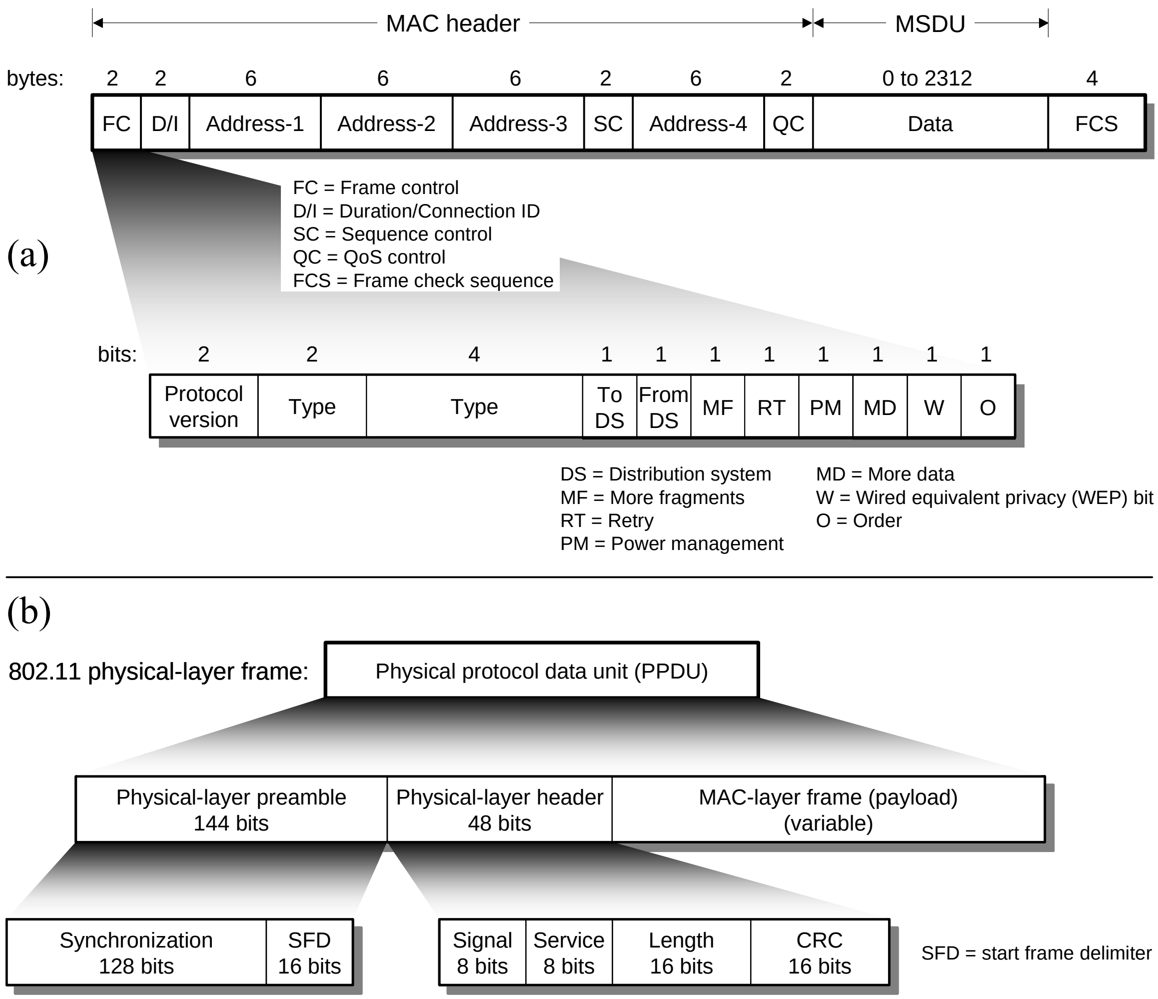

Figure 1-71: IEEE 802.11 (Wi-Fi) frame formats. (a) Link-layer (or, MAC-layer) frame format. (b) Physical-layer frame format (also known as Long PPDU format).

Medium Access Control (MAC) Protocol

The medium access control (MAC) protocol for IEEE 802.11 is a CSMA/CA protocol. As described earlier in Section 1.3.3, a CSMA/CA sender tries to avoid collision by introducing a variable amount of delay before starting with transmission. This is known as the *access deferral state*. The station sets a **contention timer** to a time interval randomly selected in the range [0, $CW-1$], and counts down to zero while sensing the carrier. If the carrier is idle when the countdown reaches zero, the station transmits.

Similar to an Ethernet adapter (Section 1.5.2), a Wi-Fi adapter needs time to decide that the channel is idle. Again, this period is known as *interframe space* (IFS). However, unlike Ethernet, the IFS delay is not fixed for all stations to 96-bit times. Wi-Fi has an additional use of the IFS delay, so that it can differentiate stations of different priority. Each station must delay its transmission according to the IFS period assigned to the station's priority class. A station with a higher priority is assigned a shorter interframe space and, conversely, lower priority stations are assigned longer IFSs. The idea behind different IFSs is to create different priority levels for

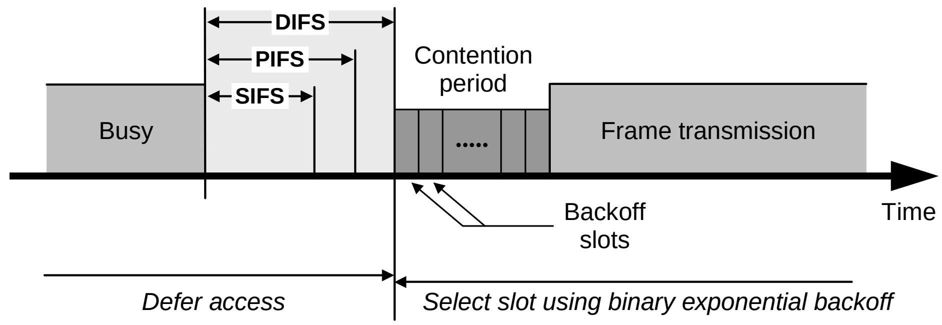

Figure 1-72: IEEE 802.11 interframe spacing relationships. Different length IFSs are used by different priority stations.

different types of traffic. Then, high-priority traffic can wait for shorter time after the medium has become idle. If there is any high-priority traffic, it grabs the medium before lower-priority frames have a chance to try.

Again, when a station wants to transmit data, it first senses whether the medium is busy. Two rules apply here:

1. If the medium has been idle for longer than an IFS corresponding to its priority level, transmission can begin immediately.
2. If the medium is busy, the station continuously senses the medium, waiting for it to become idle. When the medium becomes idle, the station first waits for its assigned IFS, and then enters the access deferral state. The station can transmit the packet if the medium is idle after the contention timer expires.

To assist with interoperability between different data rates, the interframe space is a fixed amount of time, independent of the physical layer bit rate. There are two basic intervals determined by the physical layer (PHY): the *short interframe space* (SIFS), which is equal to the parameter β, and the *slot time*, which is equal to $2\times\beta$. To be precise, the 802.11 slot time is the sum of the physical-layer Rx-Tx turnaround time[11], the *clear channel assessment* (CCA) interval, the air propagation delay on the medium, and the link-layer processing delay.

The four different types of IFSs defined in 802.11 are (see Figure 1-72):

SIFS: Short interframe space is used for the highest priority transmissions, such as control frames, or to separate transmissions belonging to a single dialog (e.g. Frame-fragment–ACK). This value is a fixed value per PHY and is calculated in such a way that the transmitting station will be able to switch back to receive mode and be capable of decoding the incoming packet. For example, for the 802.11 FH PHY this value is set to 28 microseconds.

PIFS: PCF (or priority) interframe space is used by the PCF during contention-free operation. The coordinator uses PIFS when issuing polls and the polled station may transmit after the

[11] *Rx-Tx turnaround time* is the maximum time (in μs) that the physical layer requires to change from receiving to transmitting the start of the first symbol. More information about the Rx-Tx turnaround time is available in: "IEEE 802.11 Wireless Access Method and Physical Specification," September 1993; doc: IEEE P802.11-93/147: http://www.ieee802.org/11/Documents/DocumentArchives/1993_docs/1193147.doc

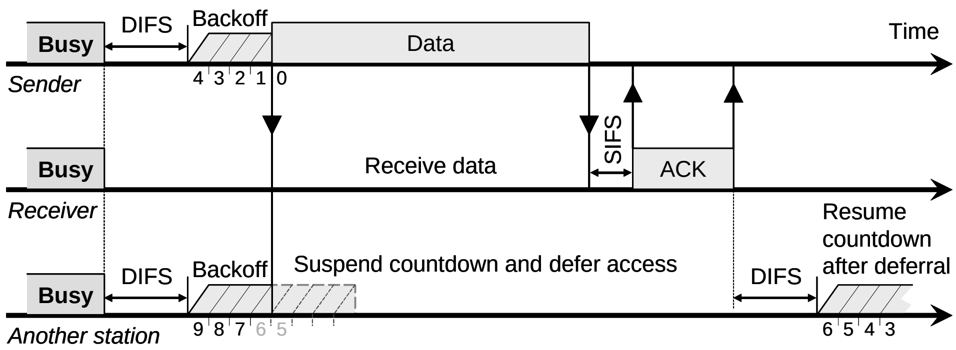

Figure 1-73: IEEE 802.11 basic transmission mode is a based on the stop-and-wait ARQ. Notice the backoff slot countdown during the contention period.

SIFS has elapsed and preempt any contention-based traffic. PIFS is equal to SIFS plus one slot time.

DIFS: DCF (or distributed) interframe space is the minimum medium idle time for asynchronous frames contending for access. Stations may have immediate access to the medium if it has been free for a period longer than the DIFS. DIFS is equal to SIFS plus two slot times.

EIFS: Extended interframe space (not illustrated in Figure 1-72) is much longer than any of the other interframe spaces. It is used by any station that has received a frame containing errors that it could not understand. This station cannot detect the duration information and set its NAV for the Virtual Carrier Sense (defined later). EIFS ensures that the station is prevented from colliding with a future packet belonging to the current dialog. In other words, EIFS allows the ongoing exchanges to complete correctly before this station is allowed to transmit.

The values of some important 802.11b system parameters are shown in Table 1-6. The values shown are for the 1Mbps channel bit rate and some of them are different for other bit rates.

Table 1-6: IEEE 802.11b system parameters. (PHY preamble serves for the receiver to distinguish silence from transmission periods and detect the beginning of a new packet.)

Parameter	Value for 1 Mbps channel bit rate
Slot time	20 μsec
SIFS	10 μsec
DIFS	50 μsec (DIFS = SIFS + 2 × Slot time)
EIFS	SIFS + PHY_preamble + PHY_header + ACK + DIFS = 364 μsec
CW_{min}	32 (minimum contention window size)
CW_{max}	1024 (maximum contention window size)
PHY_preamble	144 bits (144 μsec)
PHY_header	48 bits (48 μsec)
MAC data header	28 bytes = 224 bits
ACK	14 bytes + PHY_preamble + PHY_header = 304 bits (304 μsec)
RTS	20 bytes + PHY_preamble + PHY_header = 352 bits (352 μsec)
CTS	14 bytes + PHY_preamble + PHY_header = 304 bits (304 μsec)
MTU*	Adjustable, up to 2304 bytes for frame body before encryption

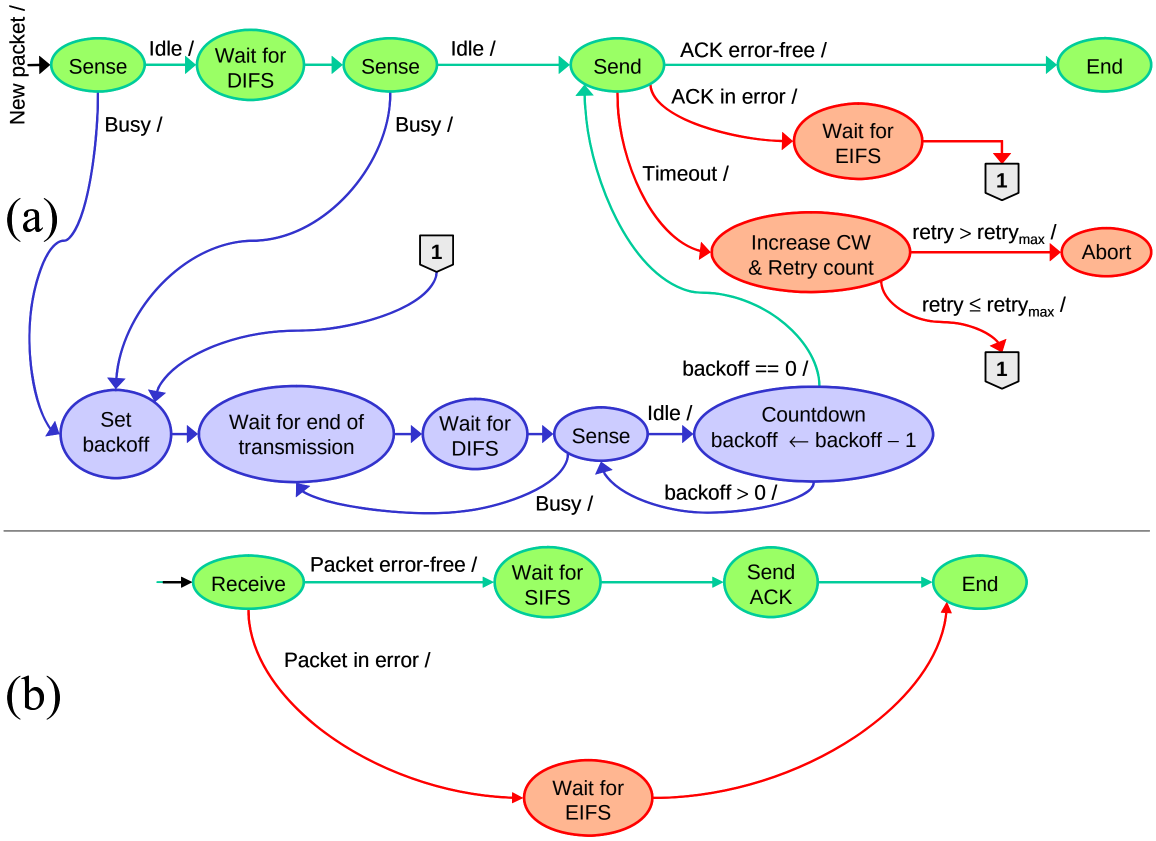

Figure 1-74: (a) Sender's state diagram of basic packet transmission for 802.11 MAC protocol. Compare to Figure 1-33. In "Set backoff," the backoff counter is set randomly to a number ∈ {0, …, *CW*–1}. (b) Receiver's state diagram for 802.11 MAC protocol.

(*) The Maximum Transmission Unit (MTU) size specifies the maximum size of a physical packet created by a transmitting device. The reader may also encounter the number 2312 (as in Figure 1-71(a)), which is the largest WEP encrypted frame payload (also known as MSDU, for *M*AC *S*ervice *D*ata *U*nit). Also, 2346 is the largest frame possible with WEP encryption and every MAC header field in use (including Address 4, see Figure 1-71(a)). In practice, MSDU size seldom exceeds 1508 bytes because of the need to bridge with Ethernet.

An example of a frame transmission from a sender to a receiver is shown in Figure 1-73. Notice that even the units of the atomic transmission (data and acknowledgement) are separated by SIFS, which is intended to give the transmitting station a short break so it will be able to switch back to receive mode and be capable of decoding the incoming (in this case ACK) packet.

The state diagrams for 802.11 senders and receivers are shown in Figure 1-74. Notice that sender's state diagram is based on the CSMA/CA protocol shown in Figure 1-33, with the key difference of introducing the interframe space.

Here is an example:

Example 1.6 Illustration of Timing Diagrams for IEEE 802.11

Consider a local area network (infrastructure BSS) using the IEEE 802.11 protocol shown in Figure 1-74. Show the timing diagrams for the following scenarios:

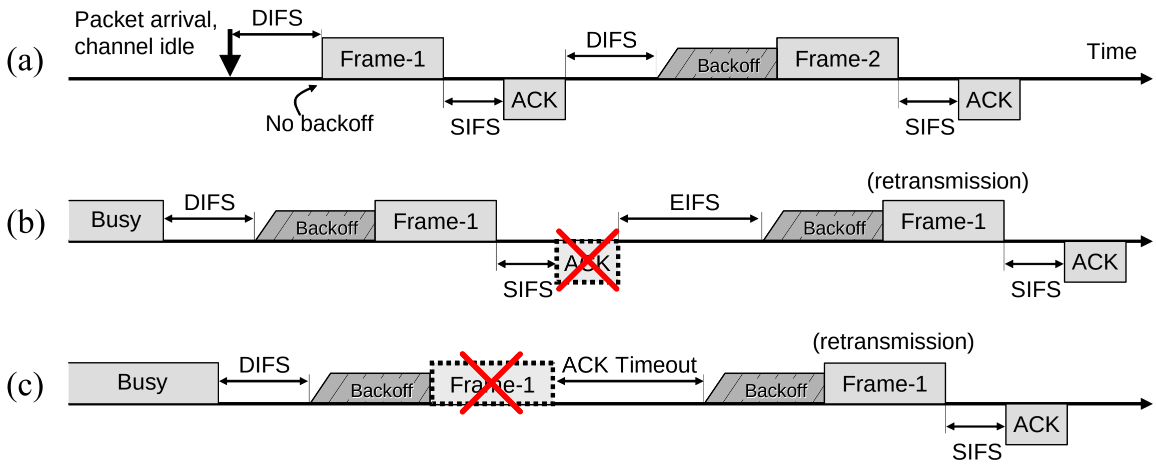

Figure 1-75: Timing diagrams. (a) Timing of successful frame transmissions under the DCF. (b) Frame retransmission due to ACK failure. (c) Frame retransmission due to an erroneous data frame reception.

(a) A single station has two frames ready for transmission on an idle channel.

(b) A single station has one frame ready for transmission on a busy channel. The acknowledgement for the frame is corrupted during the first transmission.

(c) A single station has one frame ready for transmission on a busy channel. The data frame is corrupted during the first transmission.

The solutions are shown in Figure 1-75. Sender's actions are shown above the time axis and receiver's actions are shown below the time axis. A crossed block represents a loss or erroneous reception of the corresponding frame.

The timing of successful frame transmissions is shown in Figure 1-75(a). If the channel is idle upon the packet arrival, the station transmits immediately, without backoff. However, it has to backoff for its own second transmission.

Figure 1-75(b) shows the case where an ACK frame is received in error, i.e., received with an incorrect frame check sequence (FCS). The transmitter re-contends for the medium to retransmit the frame after an EIFS interval. This is also indicated in the state diagram in Figure 1-74.

On the other hand, if no ACK frame is received within a timeout interval, due possibly to an erroneous reception at the receiver of the preceding data frame, as shown in Figure 1-75(c), the transmitter contends again for the medium to retransmit the frame after an ACK timeout. (Notice that the ACK timeout is much shorter than the EIFS interval; in fact, ACK_timeout = $t_{SIFS} + t_{ACK} + t_{slot}$. Check Table 1-6 for the values.)

Hidden Stations Problem

The hidden and exposed station problems are described earlier in Section 1.3.3. A common solution is to induce the receiver to transmit a brief "warning signal" so that other potential transmitters in its neighborhood are forced to defer their transmissions. IEEE 802.11 extends the basic access method (Figure 1-73) with two more frames: *request-to-send* (RTS) and *clear-to-send* (CTS) frames, which are very short frames exchanged before the data and ACK frames.

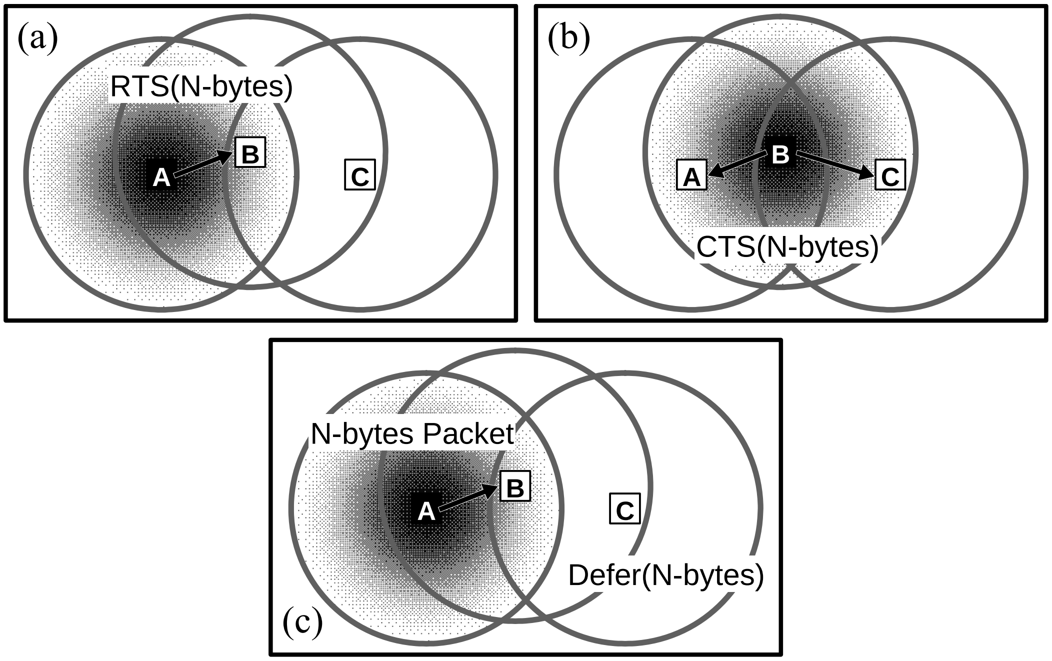

Figure 1-76: IEEE 802.11 protocol is augmented by RTS/CTS frames for hidden stations.

(Frame lengths, including RTS and CTS durations, are shown in Table 1-6.) The process is shown in Figure 1-76. The sender first sends the RTS frame to the receiver (Figure 1-76(a)). If the transmission succeeds, the receiver responds by outputting another short frame (CTS). The CTS frame is intended not only for the sender, but also for all other stations in the receiver's range (Figure 1-76(b)). All stations that receive a CTS frame know that this frame signals a transmission in progress and must avoid transmitting for the duration of the upcoming (large) data frame. Through this indirection, the sender performs "floor acquisition" so it can speak unobstructed because all other stations will remain silent for the duration of transmission (Figure 1-76(c)). Notice also that the frame length is indicated in each frame, which is the Duration D/I field of the frame header (Figure 1-71(a)).

The 4-way handshake of the RTS/CTS/DATA/ACK exchange of the 802.11 DCF protocol (Figure 1-77) requires that the roles of sender and receiver be interchanged several times between pairs of communicating nodes, so neighbors of both these nodes must remain silent *during the entire exchange*. This is achieved by relying on the virtual carrier sense mechanism of 802.11, i.e., by having the neighboring nodes set their *network allocation vector* (NAV) values from the Duration D/I field specified in either the RTS or CTS frames they overhear (Figure 1-71(a)). By using the NAV, the stations ensure that atomic operations are not interrupted. The NAV time duration is carried in the frame headers on the RTS, CTS, data and ACK frames. (Notice that the NAV vector is set only in the RTS/CTS access mode and *not* in the basic access mode shown in Figure 1-73, because RTS/CTS perform channel reservation for the subsequent data frame.)

The additional RTS/CTS exchange shortens the vulnerable period from the entire data frame in the basic method (Figure 1-73) down to the duration of the RTS/CTS exchange in the RTS/CTS method (Figure 1-77). If a "covered station" transmits simultaneously with the sender, they will collide within the RTS frame. If the hidden stations hear the CTS frame, they will *not* interfere with the subsequent data frame transmission. In either case, the sender will detect collision by the

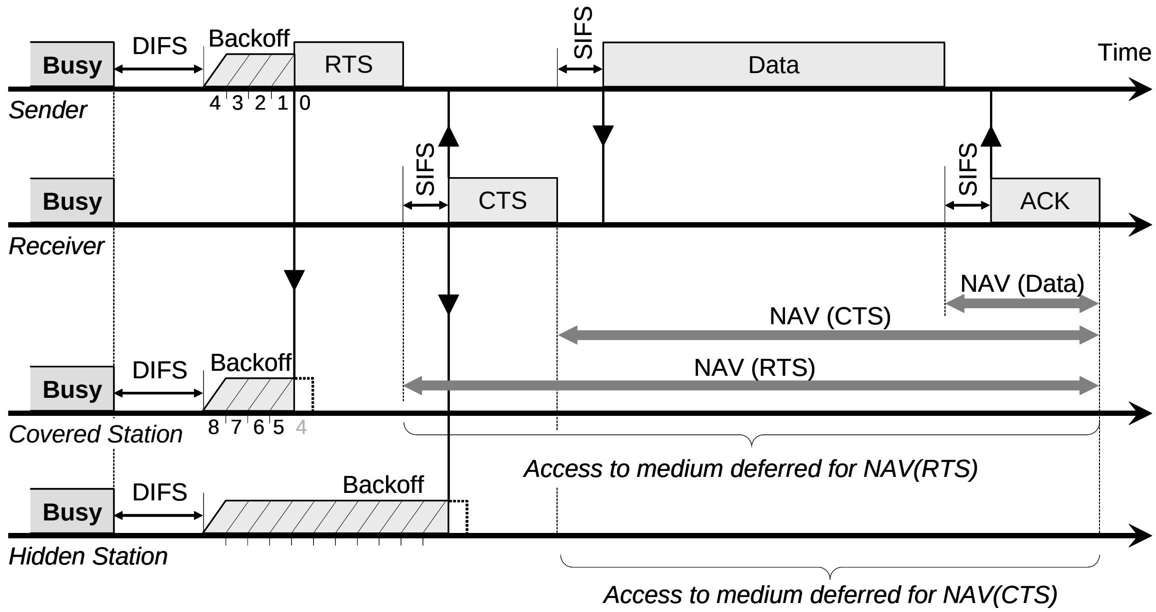

Figure 1-77: The 802.11 protocol atomic unit exchange in RTS/CTS transmission mode consists of four frames: RTS, CTS, Data, and ACK. (Compare to Figure 1-73.)

lack of the CTS frame. If a collision happens, it will last only a short period because RTS and CTS frames are very short, unlike data frames, which can be very long. This RTS/CTS exchange partially solves the hidden station problem but the exposed node problem remains unaddressed. The hidden station problem is solved only partially, because if a hidden station starts with transmission *simultaneously* with the CTS frame, the hidden station will not hear the CTS frame, the sender will receive the CTS frame correctly and start with the data frame transmission, and this will result in a collision at the receiver. (Of course, the probability of this event is very low.)

802.11 RTS/CTS protocol does not solve the exposed station problem (Figure 1-29(b)). Exposed stations could maintain their NAV vectors to keep track of ongoing transmissions. However, if an exposed station gets a packet to transmit while a transmission is in progress, it is allowed to transmit for the remainder of the NAV, before the sender needs to receive the ACK. Tailoring the frame to fit this interval and accompanied coordination is difficult and is not implemented as part of 802.11.

Recent extensions in the evolution of the 802.11 standard are described in Chapter 6.

1.6 Quality of Service Overview

This text reviews basic results about quality of service (QoS) in networked systems, particularly highlighting the wireless networks.

The recurring theme in this text is the ***delay and its statistical properties*** for a computing system. Delay (also referred to as *latency*) is modeled differently at different abstraction levels, but the

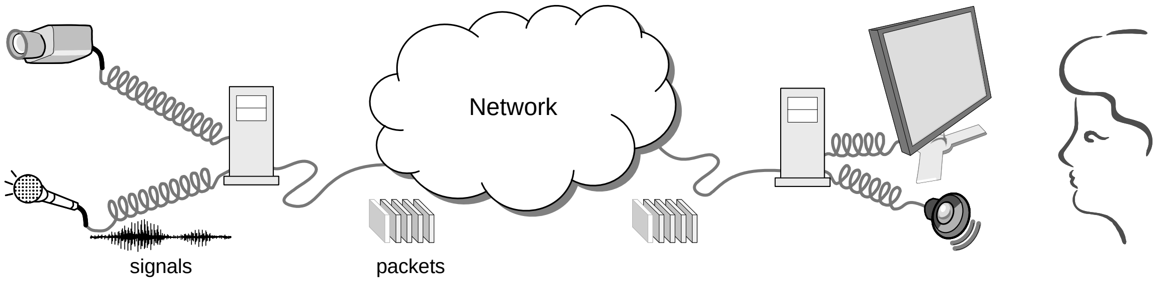

Figure 1-78: Conceptual model of multimedia information delivery over the network.

key issue remains the same: how to limit the delay so it meets task constraints. A complement of delay is *capacity*, also referred to as *bandwidth*, which was covered in the previous volume. Therein, we have seen that the system capacity is subject to physical (or economic) limitations. Constraints on delay, on the other hand, are imposed subjectively by the task of the information recipient—information often loses value for the recipient if not received within a certain deadline.

Processing and communication of information in a networked system are generally referred to as *servicing* of information. The dual constraints of capacity and delay naturally call for compromises on the quality of service. If the given capacity and delay specifications cannot provide the full service to the customer (in our case information), a compromise in service quality must be made and a sub-optimal service agreed to as a better alternative to unacceptable delays or no service at all. In other words, if the receiver can admit certain degree of information loss, then the latencies can be reduced to an acceptable range.

In order to achieve the optimal tradeoff, the players (source, intermediary, and destination) and their parameters must be considered as shown in Figure 1-79. We first define information qualities and then analyze how they get affected by the system processing.

Latency and information loss are tightly coupled and by adjusting one, we can control the other. Thus, both enter the quality-of-service specification. If all information must be received to meet the receiver's requirements, then the loss must be dealt with within the system, and the user is only aware of latencies.

Time is always an issue in information systems as is generally in life. However, there are different time constraints, such as soft, hard, as well as their statistical properties.

We are interested in assessing the servicing parameters of the intermediate and controlling them to achieve information delivery satisfactory for the receiver.

Because delay is inversely proportional to the packet loss, by adjusting one we can control the other. Some systems are "black box"—they cannot be controlled, e.g., Wi-Fi, where we cannot control the packet loss because the system parameters of the maximum number of retries determine the delay. In this case, we can control the input traffic to obtain the desired output.

Source is usually some kind of computing or sensory device, such as microphone, camera, etc. However, it may not be always possible to identify the actual traffic source. For example, it could be within the organizational boundaries, concealed for security or privacy reasons.

Figure 1-79 is drawn as if the source and destination are individual computers and (geographically) separated. The reality is not always so simple. Instead of computers, these may

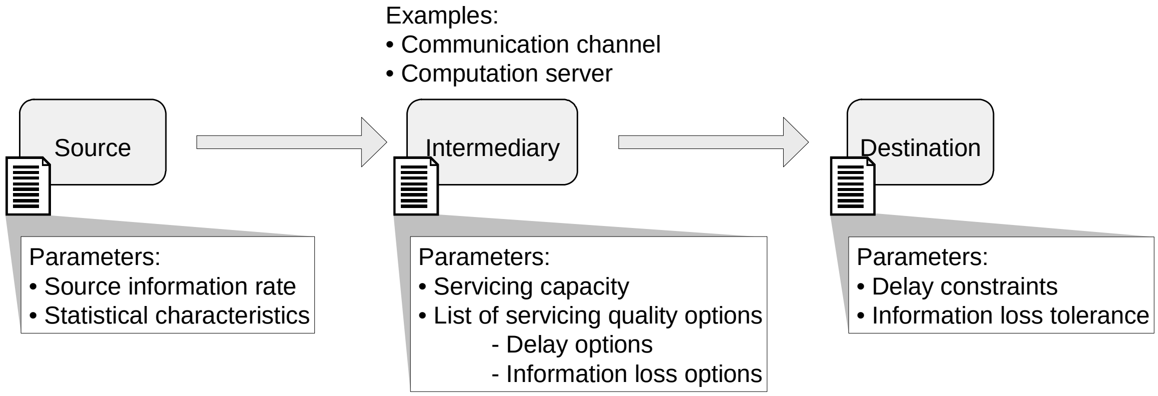

Figure 1-79: Key factors in quality of service assurance.

be people or organizations using multiple computers, or they could be networks of sensors and/or actuators.

Users exchange information through computing applications. A distributed application at one end accepts information and presents it at the other end. Therefore, it is common to talk about the characteristics of information transfer of different applications. These characteristics describe the traffic that the applications generate as well as the acceptable delays and information losses by the intermediaries (network) in delivering that traffic. We call **traffic** the aggregate bitstreams that can be observed at any cut-point in the system.

The information that applications generate can take many forms: text, audio, voice, graphics, pictures, animations, and videos. Moreover, the information transfer may be one-way, two-way, broadcast, or multipoint.

Traffic management is the set of policies and mechanisms that allow a network to satisfy efficiently a diverse range of service requests. The two fundamental aspects of traffic management, *diversity* in user requirements and *efficiency* in satisfying them, act at cross purposes, creating a tension that has led to a rich set of mechanisms. Some of these mechanisms include flow control and scheduling (Chapter 5).

QoS guarantees: hard and soft

Our primary concerns here are *delay* and *loss* requirements that applications impose on the network. We should keep in mind that other requirements, such as reliability and security may be important. When one or more links or intermediary nodes fail, the network may be unable to provide a connection between source and destination until those failures are repaired. *Reliability* refers to the frequency and duration of such failures. Some applications (e.g., control of electric power plants, hospital life support systems, critical banking operations) demand extremely reliable network operation. Typically, we want to be able to provide higher reliability between a few designated source-destination pairs. Higher reliability is achieved by providing multiple disjoint paths between the designated node pairs.

In this text we first concentrate on the parameters of the network players, traffic characteristics of information sources, information needs of information sinks, and delay and loss introduced by the intermediaries. Then we review the techniques designed to mitigate the delay and loss to meet the sinks information needs in the best possible way.

Performance Bounds

Network performance bounds can be expressed either *deterministically* or *statistically*. A deterministic bound holds for every packet sent on a connection. A statistic bound is a probabilistic bound on network performance. For example, a deterministic delay bound of 200 ms means that every packet sent on a connection experiences an end-to-end, from sender to receiver, delay smaller than 200 ms. On the other hand, a statistical bound of 200 ms with a parameter of 0.99 means that the probability that a packet is delayed by more than 200 ms is smaller than 0.01.

Quality of Service

Network operator can guarantee performance bounds for a connection only by reserving sufficient network resources, either on-the-fly, during the connection-establishment phase, or in advance.

There is more than one way to characterize quality-of-service (QoS). Generally, QoS is the ability of a network element (e.g., an application, a host or a router) to provide some level of assurance for consistent network data delivery. Some applications are more stringent about their QoS requirements than other applications, and for this reason (among others), we have two basic types of QoS available:

- **Resource reservation** (integrated services): network resources are apportioned according to an application's QoS request, and subject to bandwidth management policy.
- **Prioritization** (differentiated services): network traffic is classified and apportioned network resources according to bandwidth management policy criteria. To enable QoS, network elements give preferential treatment to classifications identified as having more demanding requirements.

These types of QoS can be applied to individual traffic "flows" or to flow aggregates, where a flow is identified by a source-destination pair. Hence, there are two other ways to characterize types of QoS:

- **Per Flow**: A "flow" is defined as an individual, unidirectional, data stream between two applications (sender and receiver), uniquely identified by a 5-tuple (transport protocol, source address, source port number, destination address, and destination port number).
- **Per Aggregate**: An aggregate is simply two or more flows. Typically, the flows will have something in common (e.g., any one or more of the 5-tuple parameters, a label or a priority number, or perhaps some authentication information).

1.6.1 QoS Outlook

QoS is becoming more important with the growth of real-time and multimedia applications. Unlike traditional network applications, such as email or Web browsing, where data transmission process is much more transparent, in real-time communications, such as phone calls, delay and loss problems are much more apparent to the users. IP networks are subject to much impairment, including:

* Packet loss due to network congestion or corruption of the data

* Variation in the amount of delay of packet delivery, which can result in poor voice quality

* Packets arriving out of sequence, which can result in discarded packets and cause more delay and disruption

There has been a great amount of research on QoS in wireline networks, but very little of it ended up being employed in actual products. Many researchers feel that there is higher chance that QoS techniques will be actually employed in wireless networks. Here are some of the arguments:

Wireline Network	**vs.**	**Wireless Network**
• Deals with *thousands* of traffic flows, thus not feasible to control		• Deals with *tens* of traffic flows (max about 50), thus it is feasible to control
• Am I a *bottleneck*?		• I am the *bottleneck*!
• Easy to add capacity		• Hard to add capacity
• Scheduling interval ~ 1 μs		• Scheduling interval ~ 1 ms, so a larger period is available to make decision

As will be seen in Chapter 3, service quality is always defined relative to human users of the network. As strange as it may sound, it is interesting to point out that the service providers may not always want to have the best network performance. In other words, the two types of end users (service providers vs. customers) may have conflicting objectives on network performance. For example, recent research of consumer behavior [Liu, 2008] has shown that task interruptions “clear the mind,” changing the way buyers make decisions. The buyers who were temporarily interrupted were more likely to shift their focus away from “bottom-up” details such as price. Instead, they concentrated anew on their overall goal of getting satisfaction from the purchased product, even if that meant paying more. The uninterrupted buyers remained more price-conscious. This seems to imply that those annoying pop-up advertisements on Web pages—or even a Web page that loads slowly—might enhance a sale!

1.6.2 Network Neutrality vs. Tiered Services

Network neutrality (or, net neutrality or Internet neutrality) is the principle that Internet service providers (ISPs) should not be allowed to block or degrade Internet traffic from their competitors in order to speed up their own. There is a great political debate going on at present related to this topic. On one hand, consumers’ rights groups and large Internet companies, such as Google and eBay, have tried to get US Congress to pass laws restricting ISPs from blocking or slowing Internet traffic. On the other hand, net neutrality opponents, such as major telecommunications companies Verizon and AT&T, argue that in order to keep maintaining and improving network

Search the Web for net neutrality

performance, ISPs need to have the power to use tiered networks to discriminate in how quickly they deliver Internet traffic. The fear is that net neutrality rules would relegate them to the status of “dumb pipes” that are unable to effectively make money on value-added services.

Today many ISPs enforce a usage cap to prevent “bandwidth hogs” from monopolizing Internet access. (Bandwidth hogs are usually heavy video users or users sharing files using peer-to-peer applications.) Service providers also enforce congestion management techniques to assure fair access during peak usage periods. Consumers currently do not have influence on these policies, and can only “vote” by exploiting a competitive market and switching to another ISP that has a larger usage cap or no usage cap at all. Consider, on the other hand, a business user who is willing to pay a higher rate that guarantees a high-definition and steady video stream that does not pause for buffering. Unfortunately, currently this is not possible, because regardless of the connection speeds available, Internet access is still a best effort service. Some industry analysts speak about a looming crisis as more Internet users send and receive bandwidth intensive content.

To discriminate against heavy video and file-sharing users, providers use what is known as deep packet inspection. *Deep packet inspection* is the act of an intermediary network node of examining the payload content of IP datagrams for some purpose. Normally, an intermediary node (router or switch) examines only link or network-layer packet headers but not the network-layer payload.

The Internet with its “best effort” service is not neutral in terms of its impact on applications that have different requirements. It is more beneficial for elastic applications that are latency-insensitive than for real-time applications that require low latency and low jitter, such as voice and real-time video. Some proposed regulations on Internet access networks define net neutrality as equal treatment among similar applications, rather than neutral transmissions regardless of applications.

See further discussion about net neutrality in Section 9.4.

1.7 Summary and Bibliographical Notes

This chapter covers some basic aspects of computer networks and wireless communications. Some topics are covered only briefly and many other important topics are left out. Practical implementations of the Internet protocols will be described in Chapter 8. To learn more about the basics of computer networking, the reader may also consult other networking books. Perhaps two of the most regarded introductory networking books currently are [Peterson & Davie 2007] and [Kurose & Ross 2010].

Section 1.1: Introduction

The end-to-end principle was formulated by Saltzer *et al*. [1984], who argued that reliable systems tend to require end-to-end processing to operate correctly, in addition to any processing in the intermediate system. They pointed out that most features in the lowest level of a communications system present costs for all higher-layer clients, even if those clients do not need the features, and are redundant if the clients have to reimplement the features on an end-to-end basis.

Keshav [1997: Chapter 5] argued from first principles that there are good reasons to require at least five layers and that no more are necessary. The layers that he identified are: physical, link, network, transport, and application layers. He argued that the functions provided by the session and presentation layers can be provided in the application with little extra work.

Early works on protocol implementation include [Clark, 1985; Hutchinson & Peterson, 1991; Thekkath, *et al.*, 1993]. Clark [1985] described the upcall architecture. Hutchinson and Peterson [1991] described a threads-based approach to protocol implementation. [Thekkath, *et al.*, 1993] is a pionerring work on user-level protocol implementation.

RFC-2679 [Almes, *et al.*, 1999(a)] defines a metric for one-way delay of packets across Internet paths. RFC-2681 [Almes, *et al.*, 1999(b)] defines a metric for round-trip delay of packets across Internet paths and follows closely the corresponding metric for one-way delay described in RFC-2679.

Section 1.2: Reliable Transmission via Redundancy

Section 1.3: Reliable Transmission by Retransmission

Automatic Repeat reQuest (ARQ) was invented by H. C. A. van Duuren during World Ward II to provide reliable transmission of characters over radio [van Duuren, 2001]. A classic early paper on ARQ and framing is [Gray, 1972]. RFC-3366 [Fairhurst & Wood, 2002] provides advice to the designers for employing link-layer ARQ techniques. This document also describes issues with supporting IP traffic over physical-layer channels where performance varies, and where link ARQ is likely to be used.

Broadcast media require medium access coordination. The earliest medium access control (MAC) protocol is called ALOHA. ALOHA was invented in the late 1960s by Norman Abramson and his colleagues at the University of Hawaii. Their goal was to use low-cost commercial radio equipment to connect computer users on Hawaiian Islands with a central time-sharing computer on the main campus.

Another MAC protocol is called Carrier Sense Multiple Access (CSMA). This means that before the sender sends a packet, it senses the medium to see if it is idle. If it is, the sender transmits the packet. A variant of CSMA called CSMA/CD (for: Collision Detection) continues to sense the carrier during the transmission to detect whether a collision will happen because some other sender connected on the same medium is transmitting at the same time. If a collision is detected, then the sender will wait for a random period of time before transmitting again.

Section 1.4: Routing and Addressing

IP version 4 along with the IPv4 datagram format was defined in RFC-791. Currently there is a great effort by network administrators to move to the next generation IP version 6 (IPv6), reviewed in Section 8.1.

Path MTU Discovery for IP version 4 is described in RFC-1191, and for IPv6 in RFC-1981.

Internet routing protocols are designed to rapidly detect failures of network elements (nodes or links) and route data traffic around them. In addition to routing around failures, sophisticated routing protocols take into account current traffic load and dynamically shed load away from congested paths to less-loaded paths.

The original ARPAnet distance vector algorithm used queue length as metric of the link cost. That worked well as long everyone had about the same line speed (56 Kbps was considered fast at the time). With the emergence of orders-of-magnitude higher bandwidths, queues were either empty or full (congestion). As a result, wild oscillations occurred and the metric was not functioning anymore. This and other reasons caused led to the adoption of Link State Routing as the new dynamic routing algorithm of the ARPAnet in 1979.

The first link-state routing concept was invented in 1978 by John M. McQuillan [McQuillan *et al.*, 1980] as a mechanism that would calculate routes more quickly when network conditions changed, and thus lead to more stable routing.

When IPv4 was first created, the Internet was rather small, and the model for allocating address blocks was based on a central coordinator: the *Internet Assigned Numbers Authority* (http://iana.org/). Everyone who wanted address blocks would go straight the central authority. As the Internet grew, this model became impractical. Today, IPv4's classless addressing scheme (CIDR) allows variable-length network IDs and hierarchical assignment of address blocks. Big Internet Service Providers (ISPs) get large blocks from the central authority, then subdivide them, and allocate them to their customers. In turn, each organization has the ability to subdivide further their address allocation to suit their internal requirements.

In Section 1.4.4, it was commented that CIDR optimizes the common case. *Optimizing the common case* is a widely adopted technique for system design. Check, for example, [Keshav, 1997, Section 6.3.5] for more details and examples. Keshav [1997] also describes several other widely adopted techniques for system design.

The IANA (http://iana.org/) is responsible for the global coordination of the DNS Root (Section 8.4), IP addressing, Autonomous System numbering (RFC-1930, RFC-4893), and other Internet protocol resources. It is operated by the *Internet Corporation for Assigned Names and Numbers*, better known as ICANN.

In Section 1.4.5, we saw how the global Internet with many independent administrative entities creates the need to reconcile economic forces with engineering solutions. A routing protocol within a single administrative domain (or, Autonomous System) just needs to move packets as efficiently as possible from the source to the destination. Unlike this, a routing protocol that spans multiple administrative domains (or, Autonomous Systems) must allow them to implement their economic preferences. In 1980s when NSFNet provided the backbone network (funded by the US National Science Foundation), and the whole Internet was organized in a single tree structure with the backbone at the root. The backbone routers exchanged routing advertisement over this tree topology using a routing protocol called Exterior Gateway Protocol (EGP), described in RFC-904. In the early 1990s, the Internet networking infrastructure opened up to competition in the US, and a number of ISPs of different sizes emerged. The evolution of the Internet from a singly-administered backbone to its current commercial structure made EGP obsolete, and it is now replaced by BGP (Section 8.2.3). Bordering routers (called speaker nodes) implement both intra-domain and inter-domain routing protocols. Inter-domain routing protocols are based on path vector routing. Path vector routing is discussed in RFC-1322 (http://tools.ietf.org/html/rfc1322)

[Berg, 2008] introduces the issues about peering and transit in the global Internet. [Johari & Tsitsiklis, 2004]. [He & Walrand, 2006] present a generic pricing model for Internet services jointly offered by a group of providers and propose a fair revenue-sharing policy based on the weighted proportional fairness criterion. [Shrimali & Kumar, 2006] develop game-theoretic models for Internet Service Provider peering when ISPs charge each other for carrying traffic and study the incentives under which rational ISPs will participate in this game. [Srinivas & Srikant, 2006] study economics of network pricing with multiple ISPs.

Section 1.5: Link-Layer Protocols and Technologies

IEEE 802.2 is the IEEE 802 standard defining Logical Link Control (LLC). The standard is available online here: http://standards.ieee.org/getieee802/802.2.html. Both Ethernet (802.3) and Wi-Fi (802.11) have different physical sublayers and MAC sublayers but converge on the same LLC sublayer (i.e., 802.2), so they have the same interface to the network layer.

The IETF has defined a serial line protocol, the **Point-to-Point Protocol (PPP)**, in RFC-1661 [Simpson, 1994] (see also RFC-1662 and RFC-1663). RFC-1700 and RFC-3232 define the 16-bit protocol codes used by PPP in the Protocol field (Figure 1-56). PPP uses HDLC (bit-synchronous or byte synchronous) framing. High-Level Data Link Control (HDLC) is a link-layer protocol developed by the International Organization for Standardization (http://www.iso.org/). The current standard for HDLC is ISO-13239, which replaces previous standards. The specification of PPP adaptation for IPv4 is RFC-1332, and the IPv6 adaptation specification is RFC-5072. Current use of the PPP protocol is in traditional serial lines, authentication exchanges in DSL networks using PPP over Ethernet (PPPoE) [RFC-2516], and in the Digital Signaling Hierarchy (generally referred to as Packet-on-SONET/SDH) using PPP over SONET/SDH [RFC-2615].

Ethernet was originally developed at Xerox's Palo Alto Research Center (PARC) in 1973 to 1975 by Robert Metcalfe and David Boggs. The name "Ethernet" refers to the cable (the ether) and it originates from the *luminiferous ether*, through which electromagnetic radiation was once thought to propagate. Network hosts on an Ethernet network use a MAC protocol based on CSMA/CD to coordinate their access to the broadcast medium. As a historic curiosity, in March 1974, Robert Z. Bachrach wrote a memo to Metcalfe and Boggs and their management, stating that "technically and conceptually there is nothing new in your proposal" and that "analysis would show that your system would be a failure." However, this simple technology pretty much blew away any sophisticated technologies that competed with it over more than thirty years. (Check also Mr. Bachrach's response here:
http://www.reddit.com/comments/1xz13/in_1974_xerox_parc_engineers_invented_ethernet).

Digital Equipment Corporation (DEC), Intel, and Xerox published the Ethernet Version 1.0 standard in 1978 for a 10-Mbps version of Ethernet, called the **DIX standard**. In September

1980, the IEEE 802.3 working group released a draft standard **802.3** of the 10-Mbps version of Ethernet, with some minor changes from the DIX standard. In 1982, DEC, Intel, and Xerox published **Ethernet Version 2.0** or **Ethernet II**. Meanwhile, the IEEE draft standard was approved in 1983 and was subsequently published as an official standard in 1985 (ANSI/IEEE Std 802.3-1985). Although there are some minor differences between the two technologies, the terms **Ethernet** and **802.3** are generally used synonymously. Since then, a number of supplements to the standard have been defined to take advantage of improvements in the technologies and to support additional communication media and higher data rate capabilities, plus several new optional medium access control features. The latest version of the 802.3 standard is available online at: http://standards.ieee.org/getieee802/802.3.html.

Ethernet's collision detection does severely limit practical throughput on loaded, unswitched networks. As a result, almost no one uses unswitched networks anymore, and full duplex Gigabit Ethernet does not support unswitched operation. Basically, today's modern, high speed Ethernet connections are synchronized connections that can use the full bandwidth of the wire without worrying about collisions.

The **IEEE Std 802.11** is a wireless local area network specification. [Crow, *et al*., 1997] discusses IEEE 802.11 wireless local area networks. In fact, 802.11 is a family of evolving wireless LAN standards. The 802.11n specification is described in this book in Section 6.3.1. The latest version of the 802.11 standard is available online at:
http://standards.ieee.org/getieee802/802.11.html.

Raj Jain, "Books on Quality of Service over IP,"
Online at: http://www.cse.ohio-state.edu/~jain/refs/ipq_book.htm

Useful information about QoS can be found here:

Leonardo Balliache, "Practical QoS," Online at: http://www.opalsoft.net/qos/index.html

Problems

Note: *Look for problem solutions on the back of this book, starting on page* Error!

Problem 1.1

Suppose you wish to transmit a long message from a source to a destination over a network path that crosses two routers. Assume that all communications are error free and no acknowledgements are used. The propagation delay and the bit rate of the communication lines are the same. Ignore all delays other than transmission and propagation delays.

(a) Given that the message consists of N packets, each L bits long, how long will it take to transmit the entire message?

(b) If, instead, the message is sent as $2\times N$ packets, each $L/2$ bits long, how long will it take to transmit the entire message?

(c) Are the durations in (a) and (b) different? Will anything change if we use smaller or larger packets? Explain why yes or why no.

Problem 1.2

Suppose host A has four packets to send to host B using Stop-and-Wait protocol. If the packets are unnumbered (i.e., the packet header does not contain the sequence number), draw a time-sequence diagram to show what packets arrive at the receiver and what ACKs are sent back to host A if the ACK for packet 2 is lost.

Problem 1.3

Suppose two hosts, sender A and receiver B, communicate using Stop-and-Wait ARQ method. Subsequent packets from A are alternately numbered with 0 or 1, which is known as the *alternating-bit protocol*.

(a) Show that the receiver B will never be confused about the packet duplicates *if and only if* there is a single path between the sender and the receiver.

(b) In case there are multiple, alternative paths between the sender and the receiver and subsequent packets are numbered with 0 or 1, show step-by-step an example scenario where receiver B is unable to distinguish between an original packet and its duplicates.

Problem 1.4

Assume that the network configuration shown in the figure below runs the Stop-and-Wait protocol. The signal propagation speed for both links is 2×10^8 m/s. The length of the link from the sender to the router is 100 m and from the router to the receiver is 10 km. Determine the sender utilization.

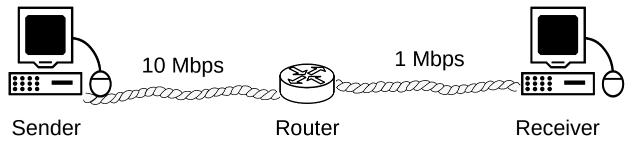

Problem 1.5

Suppose two hosts are using Go-back-2 ARQ. Draw the time-sequence diagram for the transmission of seven packets if packet 4 was received in error.

Problem 1.6

Consider a system using the Go-back-*N* protocol over a fiber link with the following parameters: 10 km length, 1 Gbps transmission rate, and 512 bytes packet length. (Propagation speed for fiber $\approx 2 \times 10^8$ m/s and assume error-free and full-duplex communication, i.e., link can transmit simultaneously in both directions. Also, assume that the acknowledgment packet size is negligible.) What value of *N* yields the maximum utilization of the sender?

Problem 1.7

Consider a system of two hosts using a sliding-window protocol sending data simultaneously in both directions. Assume that the maximum frame sizes (MTUs) for both directions are equal and the acknowledgements may be piggybacked on data packets, i.e., an acknowledgement is carried in the header of a data frame instead of sending a separate frame for acknowledgment only.

(a) In case of a full-duplex link, what is the minimum value for the retransmission timer that should be selected for this system?

(b) Can a different values of retransmission timer be selected if the link is half-duplex?

Problem 1.8

Suppose three hosts are connected as shown in the figure. Host *A* sends packets to host *C* and host *B* serves merely as a relay. However, as indicated in the figure, they use different ARQ's for reliable communication (Go-back-*N* vs. Selective Repeat). Notice that *B* is *not* a router; it is a regular host running both receiver (to receive packets from *A*) and sender (to forward *A*'s packets to *C*) applications. *B*'s receiver immediately relays in-order packets to *B*'s sender.

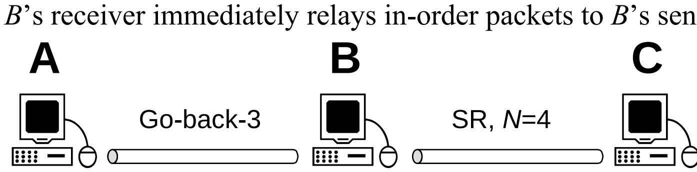

Draw side-by-side the timing diagrams for A→B and B→C transmissions up to the time where the first seven packets from *A* show up on *C*. Assume that the 2nd and 5th packets arrive in error to host *B* on their first transmission, and the 5th packet arrives in error to host *C* on its first transmission.

Discuss the merits of sending ACKs end-to-end, from destination C to source A, as opposed to sending ACKs independently for each individual link.

Problem 1.9

Consider the network configuration as in Problem 1.8. However, this time around assume that the protocols on the links are reverted, as indicated in the figure, so the first pair uses Selective Repeat and the second uses Go-back-*N*, respectively.

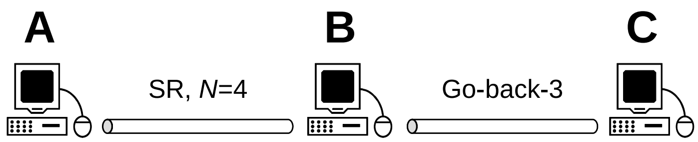

Draw again side-by-side the timing diagrams for A→B and B→C transmissions assuming the same error pattern. That is, the 2nd and 5th packets arrive in error to host *B* on their first transmission, and the 5th packet arrives in error to host *C* on its first transmission.

Problem 1.10

Assume the following system characteristics (see the figure below):

The link transmission speed is 1 Mbps; the physical distance between the hosts is 300 m; the link is a copper wire with signal propagation speed of 2×10^8 m/s. The data packets to be sent by the hosts are 2 Kbytes each, and the acknowledgement packets (to be sent separately from data packets) are 10 bytes long. Each host has 100 packets to send to the other one. Assume that the transmitters are somehow synchronized, so they *never* attempt to transmit simultaneously from both endpoints of the link.

Each sender has a *window* of 5 packets. If at any time a sender reaches the limit of 5 packets outstanding, it stops sending and waits for an acknowledgement. Because there is no packet loss (as stated below), the timeout timer value is irrelevant. This is similar to a Go-back-*N* protocol, with the following difference.

The hosts do *not* send the acknowledgements immediately upon a successful packet reception. Rather, the acknowledgements are sent periodically, as follows. At the end of an 82 ms period, the host examines whether any packets were successfully received during that period. If one or more packets were received, a single (cumulative) acknowledgement packet is sent to acknowledge all the packets received in this period. Otherwise, no acknowledgement is sent.

Consider the two scenarios depicted in the figures (a) and (b) below. The router in (b) is 150 m away from either host, i.e., it is located in the middle. If the hosts in each configuration start sending packets at the same time, which configuration will complete the exchange sooner? Show the process.

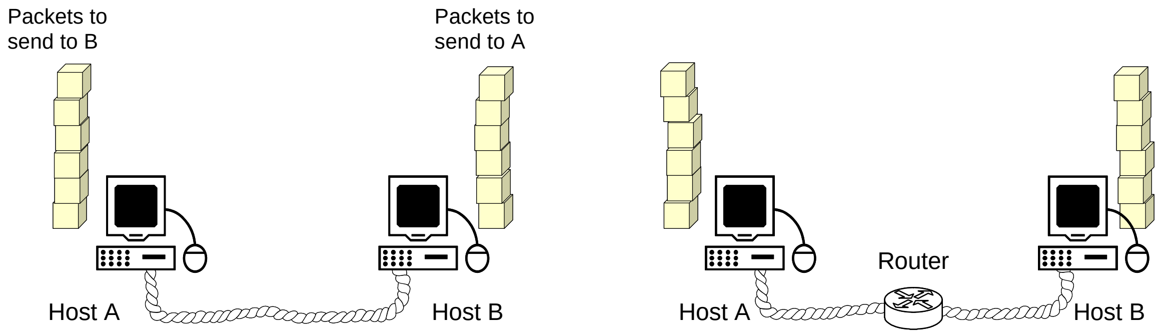

(a) (b)

Assume no loss or errors on the communication links. The router buffer size is unlimited for all practical purposes and the processing delay at the router approximately equals zero. Notice that the router can simultaneously send and receive packets on different links.

Problem 1.11

Consider two hosts directly connected and communicating using Go-back-*N* ARQ in the presence of channel errors. Assume that data packets are of the same size, the transmission delay t_x per packet, one-way propagation delay t_p, and the probability of error for data packets equals p_e. Assume that ACK packets are effectively zero bytes and always transmitted error free.

(a) Find the expected delay per packet transmission. Assume that the duration of the timeout t_{out} is large enough so that the source receives ACK before the timer times out, when both a packet and its ACK are transmitted error free.

(b) Assume that the sender operates at the maximum utilization and determine the expected delay per packet transmission.

Note: This problem considers only the expected delay from the start of the first attempt at a packet's transmission until its successful transmission. It does not consider the *waiting delay*, which is the time the packet arrives at the sender until the first attempt at the packet's transmission. The waiting delay will be considered later in Section 4.4.

Problem 1.12

Given a 64Kbps link with 1KB packets and RTT of 0.872 seconds:

(a) What is the maximum possible throughput, in packets per second (pps), on this link if a Stop-and-Wait ARQ scheme is employed and all transmissions are error-free?

(b) Again assuming that S&W ARQ is used, what is the expected throughput (pps) if the probability of error-free transmission is p=0.95?

(c) If instead a Go-back-*N* (GBN) sliding window ARQ protocol is deployed, what is the average throughput (pps) assuming error-free transmission and fully utilized sender?

(d) For the GBN ARQ case, derive a lower bound estimate of the expected throughput (pps) given the probability of error-free transmission p=0.95.

Problem 1.13

Assume a slotted ALOHA system with 10 stations, a channel with transmission rate of 1500 bps, and the slot size of 83.33 ms. What is the maximum throughput achievable per station if packets are arriving according to a Poisson process?

Problem 1.14

Consider a slotted ALOHA system with m stations and unitary slot length. Derive the following probabilities:

(a) A new packet succeeds on the first transmission attempt

(b) A new packet suffers exactly K collisions and then a success

Problem 1.15

Consider a slotted ALOHA network with m mobile stations and packet arrivals modeled as a Poisson process with rate λ. Solve the following:

(a) Assuming that this system operates with maximum efficiency, what are the fractions of slots that, on average, go unused (idle), slots that are used for successful transmission, and slots that experience packet collisions?

(b) Describe under what scenarios the system would operate with a less-than-maximum efficiency. Under such scenarios, what are the fractions of idle, successful, and collision slots?

(c) Given a steady arrival rate λ, would each non-maximum-efficiency operating point remain stable? Explain why yes or why no.

Hint: Carefully examine Figure 1-28 and other figures related to ALOHA.

Problem 1.16

Problem 1.17

Suppose two stations are using nonpersistent CSMA with a modified version of the binary exponential backoff algorithm, as follows. In the modified algorithm, each station will always wait 0 or 1 time slots with equal probability, regardless of how many collisions have occurred.

(a) What is the probability that contention ends (i.e., one of the stations successfully transmits) on the first round of retransmissions?

(b) What is the probability that contention ends on the second round of retransmissions (i.e., success occurs after one retransmission ends in collision)?

(c) What is the probability that contention ends on the third round of retransmissions?

(d) In general, how does this algorithm compare against the nonpersistent CSMA with the normal binary exponential backoff algorithm in terms of performance under different types of load?

Problem 1.18

A network using random access protocols has three stations on a bus with source-to-destination propagation delay τ. Station A is located at one end of the bus, and stations B and C are together at the other end of the bus. Frames arrive at the three stations are ready to be transmitted at stations A, B, and C at the respective times $t_A = 0$, $t_B = \tau/2$, and $t_C = 3\tau/2$. Frames require transmission times of 4τ. In appropriate timing diagrams with time as the horizontal axis, show the transmission activity of each of the three stations for (a) Pure ALOHA; (b) Non-persistent CSMA; (c) CSMA/CD.
Note: In case of collisions, show only the first transmission attempt, *not* retransmissions.

Problem 1.19

Problem 1.20

Consider a local area network using the CSMA/CA protocol shown in Figure 1-33. Assume that three stations have frames ready for transmission and they are waiting for the end of a previous transmission. The stations choose the following backoff values for their first frame: STA1 = 5, STA2 = 9, and STA3=2. For their second frame backoff values, they choose: STA1 = 7, STA2 = 1, and STA3=4. For the third frame, the backoff values are STA1 = 3, STA2 = 8, and STA3=1. Show the timing diagram for the first 5 frames. Assume that all frames are of the same length.

Problem 1.21

Consider a CSMA/CA protocol that has a backoff window size equal 2 slots. If a station transmits successfully, it remains in state 1. If the transmission results in collision, the station randomly chooses its backoff state from the set {0, 1}. If it chooses 1, it counts down to 0 and transmits. (See Figure 1-33 for details of the algorithm.) What is the probability that the station will be in a particular backoff state?

Hint: Figure 1-80 shows an analogy with a playground slide. A kid is climbing the stairs and upon reaching Platform-1 decides whether to enter the slide or to proceed climbing to Platform-2 by tossing a fair coin. If the kid enters the slide on Platform-1, he slides down directly through Tube-1. If the kid enters the slide on Platform-2, he slides through Tube-2 first, and then continues down through Tube-1. Think of this problem as the problem of determining the probabilities that the kid will be found in Tube-1 or in Tube-2. For the sake of simplicity, we will distort the reality and assume that climbing the stairs takes no time, so the kid is always found in one of the tubes.

Problem 1.22

Consider a CSMA/CA protocol with two backoff stages. If a station previously was idle, or it just completed a successful transmission, or it experienced two collisions in a row (i.e., it exceeded the retransmission limit, equal 2), then it is in the first backoff stage. In the first stage, when a new packet arrives, the station randomly chooses its backoff state from the set {0, 1}, counts down to 0, and then transmits. If the transmission is successful, the station remains in the first backoff stage and waits for another packet to send. If the transmission results in collision, the

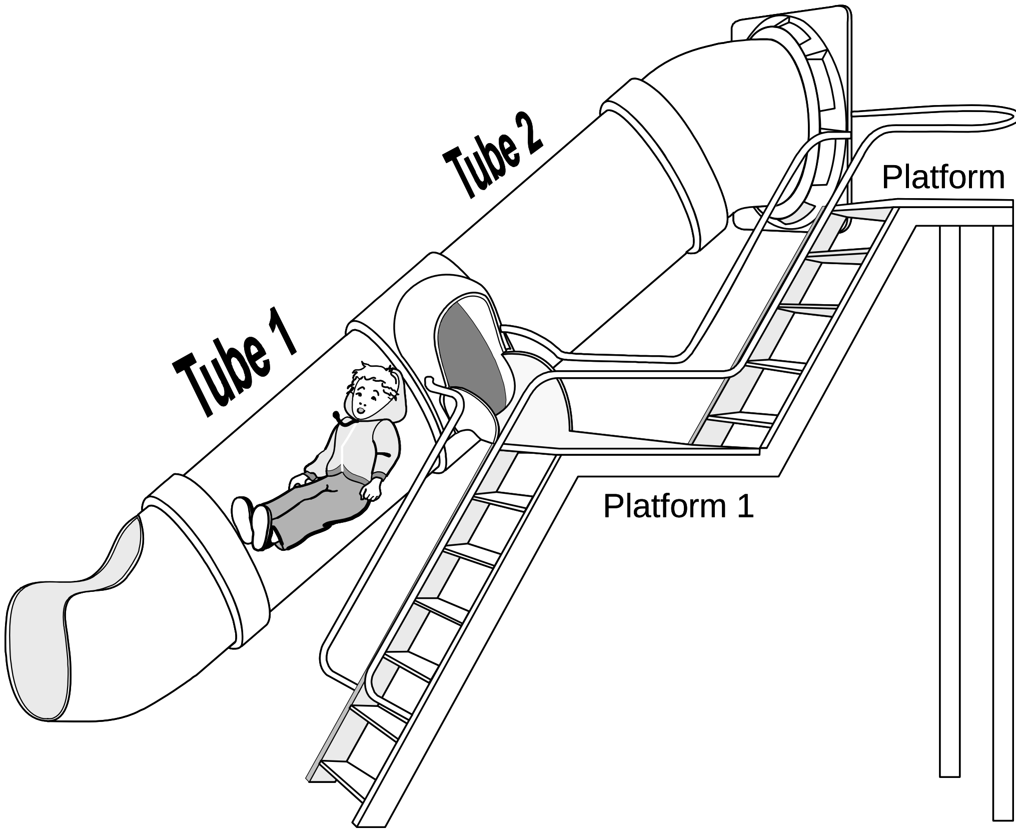

Figure 1-80: A playground-slide analogy to help solve Problem 1.21.

station jumps to the second backoff stage. In the second stage, the station randomly chooses its backoff state from the set {0, 1, 2, 3}, counts down to 0, and then retransmits the previously collided packet. If the transmission is successful, the station goes to the first backoff stage and waits for another packet to send. If the transmission results in collision, the station discards the packet (because it reached the retransmission limit, equal 2), and then jumps to the first backoff stage and waits for another packet to send.

Continuing with the playground-slide analogy of Problem 1.21, we now imagine an amusement park with two slides, as shown in Figure 1-81. The kid starts in the circled marked “START.” It first climbs Slide-1 and chooses whether to enter it at Platform-11 or Platform-12 with equal probability, i.e., 0.5. Upon sliding down and exiting from Slide-1, the kid comes to two gates. Gate-1 leads back to the starting point. Gate-2 leads to Slide-2, which consists of four tubes. The kid decides with equal probability to enter the tube on one of the four platforms. That is, on Platform-21, the kid enters the tube with probability 0.25 or continues climbing with probability 0.75. On Platform-22, the kid enters the tube with probability 1/3 or continues climbing with probability 2/3. On Platform-23, the kid enters the tube with probability 0.5 or continues climbing with probability 0.5. Upon sliding down and exiting from Slide-1, the kid always goes back to the starting point.

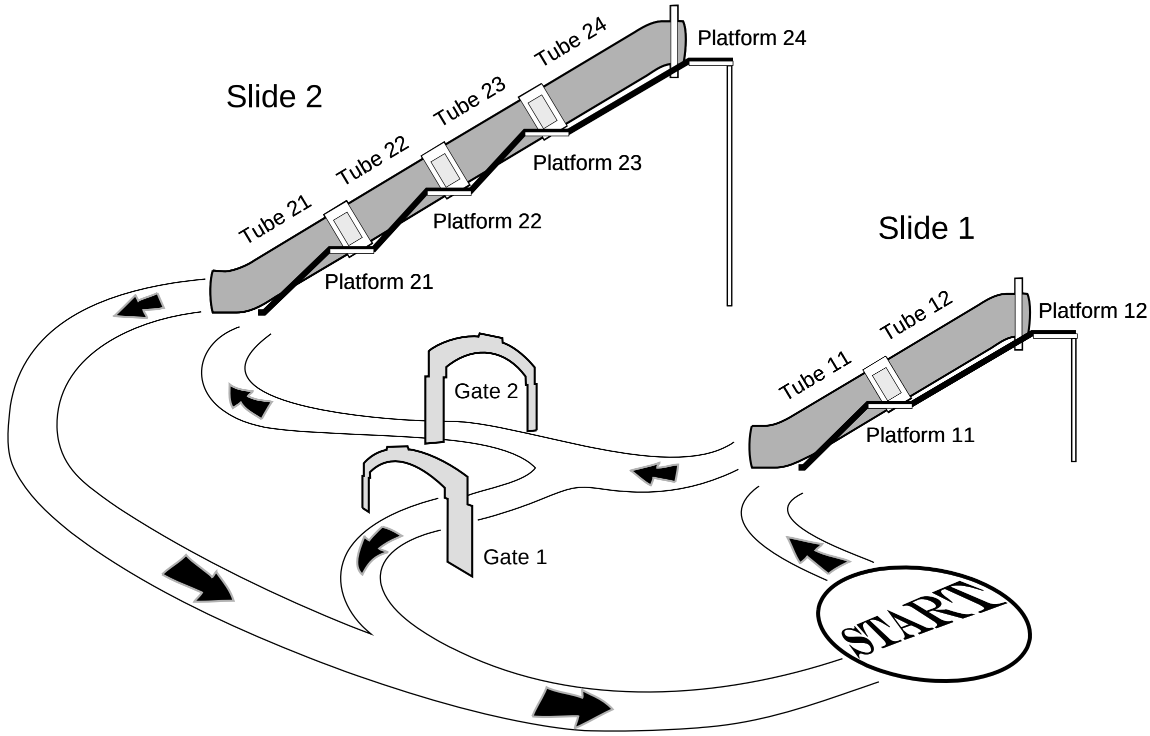

Figure 1-81: An analogy of a playground with two slides to help solve Problem 1.22.

Problem 1.23

Consider three wireless stations using the CSMA/CA protocol at the channel bit rate of 1 Mbps. The stations are positioned as shown in the figure. Stations *A* and *C* are hidden from each other and both have data to transmit to station *B*. Each station uses a timeout time for acknowledgements equal to 334 μsec. The initial backoff window range is 32 slots and the backoff slot duration equals 20 μsec. Assume that both *A* and *C* each have a packet of 44 bytes to send to *B*.

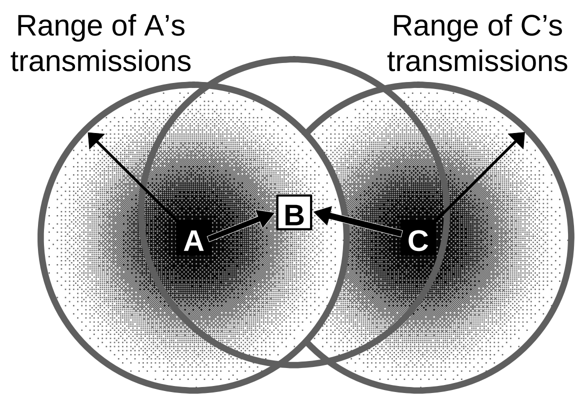

Suppose that stations *A* and *C* just heard station *B* send an acknowledgement for a preceding transmission. Let us denote the time when the acknowledgement transmission finished as $t = 0$. Do the following:

(a) Assuming that station *A* selects the backoff countdown $b_{A1} = 12$ slots, determine the vulnerable period for reception of the packet for *A* at receiver *B*.

(b) Assuming that simultaneously station C selects the backoff countdown b_{C1} = 5 slots, show the exact timing diagram for any packet transmissions from all three stations, until either a successful transmission is acknowledged or a collision is detected.

(c) After a completion of the previous transmission (ended in step (b) either with success or collision), assume that stations A and C again select their backoff timers (b_{A2} and b_{C2}, respectively) and try to transmit a 44-bytes packet each. Assume that A will start its second transmission at t_{A2}. Write the inequality for the range of values of the starting time for C's second transmission (t_{C2}) in terms of the packet transmission delay (t_x), acknowledgement timeout time (t_{ACK}) and the backoff periods selected by A and C for their first transmission (b_{A1} and b_{C1}, respectively).

Problem 1.24

Kathleen is emailing a long letter to Joe. The letter size is 16 Kbytes. Assume that TCP is used and the connection crosses 3 links as shown in the figure below. Assume link layer header is 40 bytes for all links, IP header is 20 bytes, and TCP header is 20 bytes.

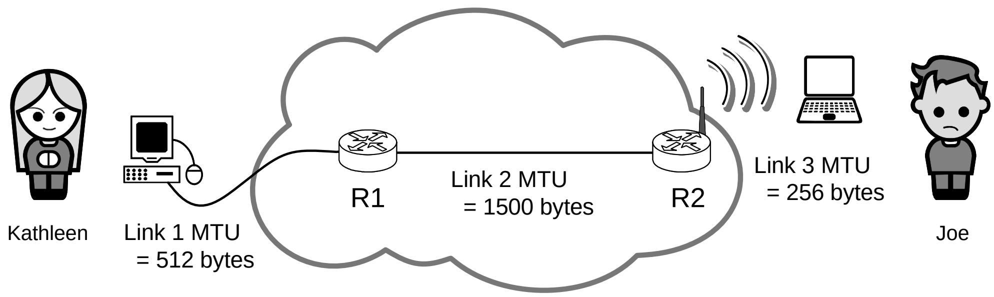

(a) How many packets/datagrams are generated in Kathleen's computer on the IP level? Show the derivation of your result.

(b) How many fragments Joe receives on the IP level? Show the derivation.

(c) Show the first 4 and the last 5 IP fragments Joe receives and specify the values of all relevant parameters (data length in bytes, ID, offset, flag) in each fragment header. Fragment's ordinal number should be specified. Assume initial ID = 672.

(d) What will happen if the very last fragment is lost on Link 3? How many IP datagrams will be retransmitted by Kathleen's computer? How many retransmitted fragments will Joe receive? Specify the values of all relevant parameters (data length, ID, offset, flag) in each fragment header.

Problem 1.25

Consider the network in the figure below, using link-state routing (the cost of all links is 1):

Suppose the following happens in sequence:

(a) The *BF* link fails

(b) New node *H* is connected to *G*

(c) New node *D* is connected to *C*

(d) New node *E* is connected to *B*

(e) A new link *DA* is added

(f) The failed *BF* link is restored

Show what link-state advertisements (LSAs) will flood back and forth for each step above. Assume that (*i*) the initial LSA sequence number at all nodes is 1, (*ii*) no packets time out, (*iii*) each node increments the sequence number in their LSA by 1 for each step, and (*iv*) both ends of a link use the same sequence number in their LSA for that link, greater than any sequence number either used before.

[You may simplify your answer for steps (b)-(f) by showing only the LSAs which change (not only the sequence number) from the previous step.]

Problem 1.26

Problem 1.27

Consider the network in the figure below and assume that the distance vector algorithm is used for routing. Show the distance vectors *after* the routing tables on all nodes are stabilized. Now assume that the link $\overline{AC}$ with weight equal to 1 is broken. Show the distance vectors on all nodes for up to five subsequent exchanges of distance vectors or until the routing tables become stabilized, whichever comes first.

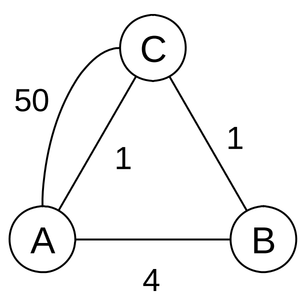

Problem 1.28

Consider the following network, using distance-vector routing:

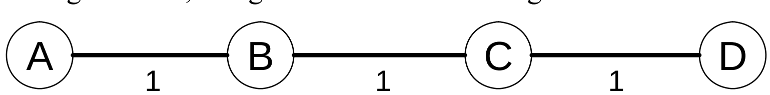

Suppose that, after the network stabilizes, link *C*–*D* goes down. Show the routing tables on the nodes *A*, *B*, and *C*, for the subsequent five exchanges of the distance vectors. How do you expect the tables to evolve for the future steps? State explicitly all the possible cases and explain your answer.

Problem 1.29

Consider the network shown in the figure, using the distance vector routing algorithm. Assume that all routers exchange their distance vectors periodically every 60 seconds regardless of any changes. If a router discovers a link failure, it broadcasts its updated distance vector within 1 second of discovering the failure.

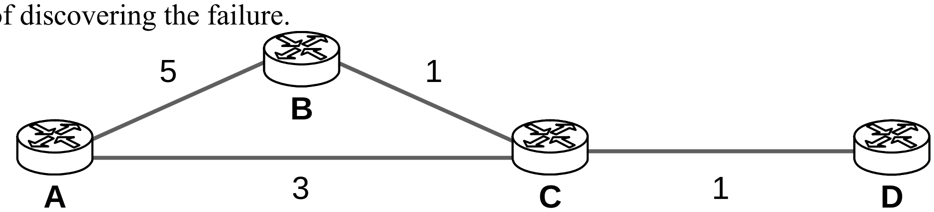

(a) Start with the initial state where the nodes know only the costs to their neighbors, and show how the routing tables at all nodes will reach the stable state.
(b) Use the results from part (a) and show the forwarding table at node *A*. [Note: use the notation *AC* to denote the output port in node *A* on the link to node *C*.]
(c) Suppose the link *CD* fails. Give a sequence of routing table updates that leads to a routing loop between *A*, *B*, and *C*.
(d) Would a routing loop form in (c) if all nodes use the split-horizon routing technique? Would it make a difference if they use split-horizon with poisoned reverse? Explain your answer.

Problem 1.30

Problem 1.31

You are hired as a network administrator for the network of sub-networks shown in the figure. Assume that the network will use the CIDR addressing scheme.

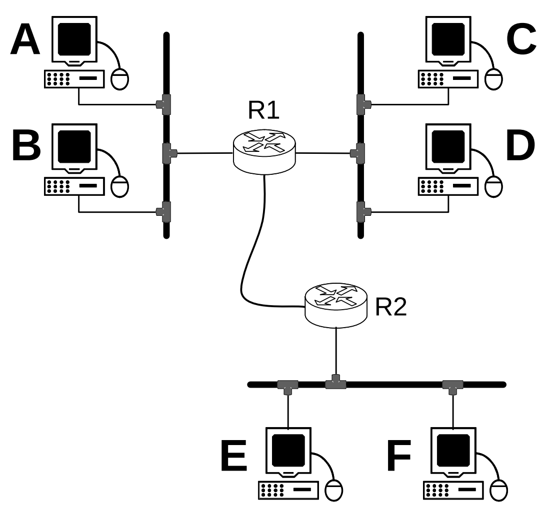

(a) Assign meaningfully the IP addresses to all hosts on the network. Allocate the minimum possible block of addresses for your network, assuming that no new hosts will be added to the current configuration.
(b) Show how routing/forwarding tables at the routers should look *after* the network stabilizes (do not show the process).

Problem 1.32

The following is a forwarding table of a router *X* using CIDR. Note that the last three entries cover every address and thus serve in lieu of a default route.

Subnet Mask	Next Hop
223.92.32.0 / 20	A
223.81.196.0 / 12	B
223.112.0.0 / 12	C
223.120.0.0 / 14	D
128.0.0.0 / 1	E
64.0.0.0 / 2	F
32.0.0.0 / 3	G

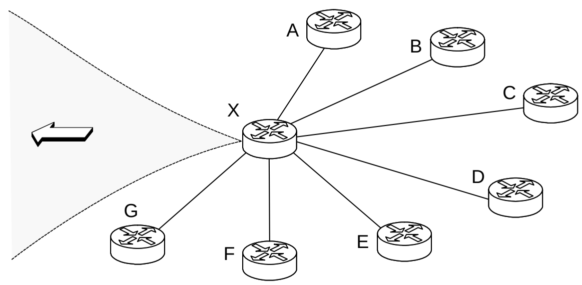

State to what next hop the packets with the following destination IP addresses will be delivered:

(a) 195.145.34.2
(b) 223.95.19.135
(c) 223.95.34.9
(d) 63.67.145.18
(e) 223.123.59.47
(f) 223.125.49.47

(Keep in mind that the default matches should be reported only if no other match is found.)

Problem 1.33

Suppose a router receives a set of packets and forwards them as follows:

(a) Packet with destination IP address **128.6.4.2**, forwarded to the next hop **A**
(b) Packet with destination IP address **128.6.236.16**, forwarded to the next hop **B**
(c) Packet with destination IP address **128.6.29.131**, forwarded to the next hop **C**
(d) Packet with destination IP address **128.6.228.43**, forwarded to the next hop **D**

Reconstruct only the part of the router's forwarding table that you suspect is used for the above packet forwarding. Use the CIDR notation and select the ***shortest network prefixes*** that will produce unambiguous forwarding:

Network Prefix	Subnet Mask	Next Hop
……………………………	………………..	…
……………………………	………………..	…
……………………………	………………..	…
……………………………	………………..	…

Problem 1.34

Problem 1.35

Consider the internetwork of autonomous systems shown in Figure 1-49. Assume that all stub ASs already advertised the prefixes of their networks so by now the internetwork topology is known to all other ASs. (We assume that non-stub ASs do not advertise any destinations within themselves—they advertise their presence only if they lay on a path to a stub-AS destination.) Given the business interests of different ASs, and as illustrated in Figure 1-50 and Figure 1-51, customers of different ISPs will not learn about all links between all ASs in the internetwork. For example, customers of ISP φ will see the internetwork topology as shown in the figure below:

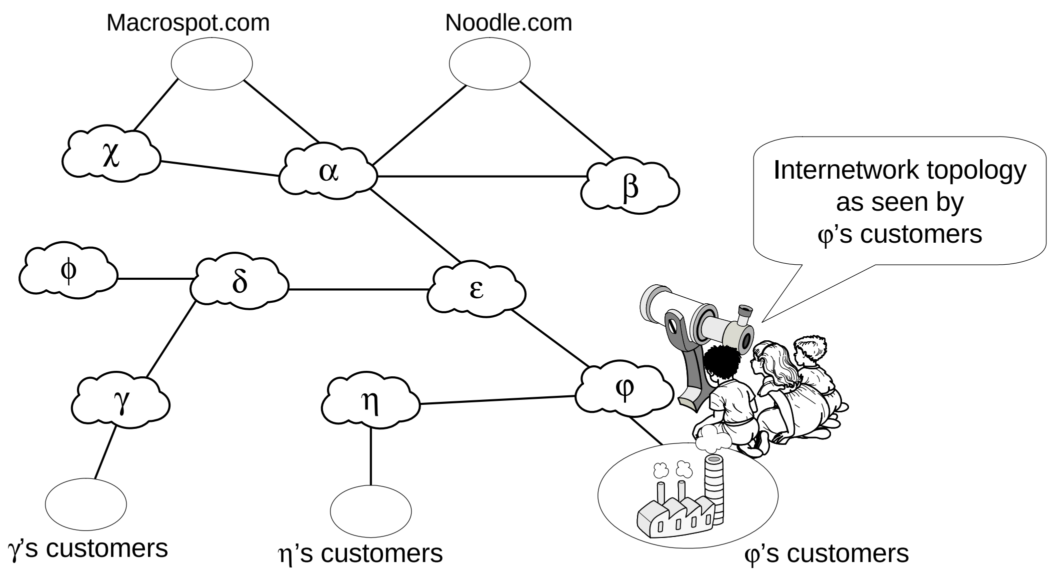

As explained in the description of Figure 1-50, ASη has no interest in carrying ASφ's transit traffic, so it will not advertise that it has a path to ASδ. Therefore, ASφ and its customers will not learn about this path. For the same reason the connection between ASδ and ASα will not be visible. Notice that ASφ will learn about alternative paths from/to Macrospot.com or Noodle.com through ASs χ and β, respectively, because multihomed ASs will advertise their network prefixes to all directly connected ISPs.

Your task is to show the respective views of the internetwork topology for the customers of ASs γ and η, and for the corporations Macrospot.com or Noodle.com. For each of these ASs, draw the internetwork topology as they will see it and explain why some connections or ASs will not be visible, if there are such.

Problem 1.36

Consider the internetwork of autonomous systems shown in the figure below. Autonomous systems α and β are peers and they both buy transit from AS γ. Assume that all links cost 1 (one hop) and every AS uses hot-potato routing to forward packets (destined for other ASs) towards speaker nodes.

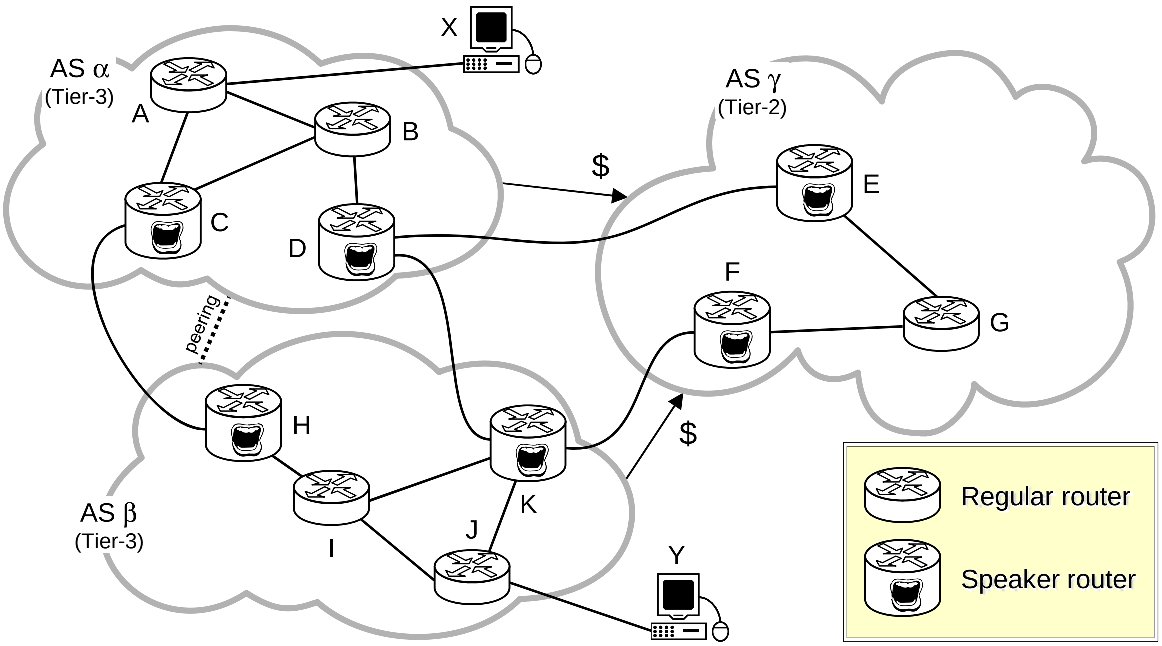

(d) How many paths (in terms of autonomous systems, not individual routers) to customers of AS β are available for use to reach customers of AS α?

(e) What path will traffic from host X to host Y normally take? List the autonomous systems and within each autonomous system list the individual routers (hop-by-hop). How about traffic from host Y to host X?

(f) Is it possible for *all* traffic from host X to host Y and vice versa to take the same path? Explain why yes or why no, and if yes, how this can be achieved.

(g) Is it possible for traffic from host X to host Y and vice versa to take the same that includes AS γ? Explain why yes or why no, and if yes, how this can be achieved.

Problem 1.37

Problem 1.38

Problem 1.39

Problem 1.40

PPP uses the Link Control Protocol (LCP) to establish, maintain, and terminate the physical connection. During the link establishment phase, the configuration parameters are negotiated by an exchange of LCP frames. Before information can be sent on a link, each of the two computers that make up the connection must test the link and agree on a set of parameters under which the link will operate. If the negotiation converges, the link is established and either the authentication is performed or the network layer protocol can start sending data. If the endpoints fail to negotiate a common configuration for the link, it is closed immediately. LCP is also responsible for maintaining and terminating the connection.

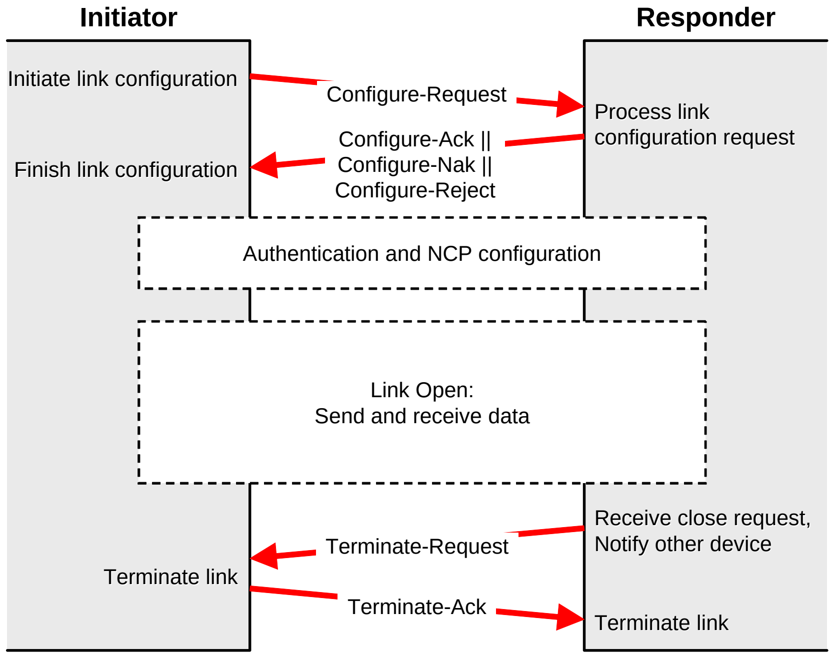

The *Configure-Request* message is sent to request a link establishment and it contains the various options requested. This request is responded with a *Configure-Ack* ("acknowledge") if every requested option is acceptable. A *Configure-Nak* ("negative acknowledge") is sent if all the requested options are recognizable but some of their requested values are not acceptable. This message includes a copy of each configuration option that the Responder found unacceptable and it suggests an acceptable negotiation.

A *Configure-Reject* is sent if any of the requested options were either unrecognizable or represent unacceptable ways of using the link or are not subject to negotiation. This message includes the objectionable options. Configure-Request frames are transmitted periodically until either a Configure-Ack is received, or the number of frames sent exceeds the maximum allowed value.

A simplified format of an LCP frame is as follows (Figure 1-58):

Code	Identifier	Data

The *Code* field identifies the type of LCP frame, such as *Configure-Request*, *Configure-Ack*, *Terminate-Request*, etc. The *Identifier* field carries an identifier that is used to match associated requests and replies. When a frame is received with an invalid Identifier field, the frame is silently discarded without affecting the protocol execution.

- For a Configure-Request frame, the Identifier field *must* be changed whenever the contents of the Data field changes (Data carries the link configuration options), and whenever a valid reply has been received for a previous request. For retransmissions, the Identifier *may* remain unchanged.
- For a configuration response frame (Configure-Ack, Configure-Nak, or Configure-Reject), the Identifier field is an exact copy of the Identifier field of the Configure-Request that caused this response frame.
- For a Terminate-Request frame, the Identifier field *must* be changed whenever the content of the Data field changes, and whenever a valid reply has been received for a previous request. For retransmissions, the Identifier *may* remain unchanged.

- For a Terminate-Ack frame, on reception, the Identifier field of the Terminate-Request is copied into the Identifier field of the Terminate-Ack frame.

During the link establishment phase, only LCP frames should be transmitted in the PPP frames. Any non-LCP frames received during this phase must be silently discarded.

The LCP link termination frames are: *Terminate-Request*, which represents the start of the link termination phase; and, *Terminate-Ack*, which acknowledges the receipt of a recognizable Terminate-Request frame, and accepts the termination request. Under ideal conditions, the link termination phase is signaled end-to-end using LCP link termination frames. However, the link termination phase also can be caused by a loss of carrier or an immediate shutdown by the system administrator.

Explain the following:

(a) Why are unique identifiers needed in LCP frames? Why a configuration response frame, such as Configure-Ack, without the Identifier field is insufficient? Illustrate your argument by drawing a time-diagram for a scenario where LCP link configuration would fail if the Identifier field did not exist.

(b) Why is the “two-way handshake” using Terminate-Ack frame needed in the link termination phase? Illustrate your argument by drawing a time-diagram of a scenario where LCP link termination would fail if Terminate-Ack were not sent.

Problem 1.41

Problem 1.42

Consider the following Ethernet network where all the switches employ the spanning tree protocol (STP) to remove the loops in the network topology. The numbers in parentheses represent the switches’ identifiers.

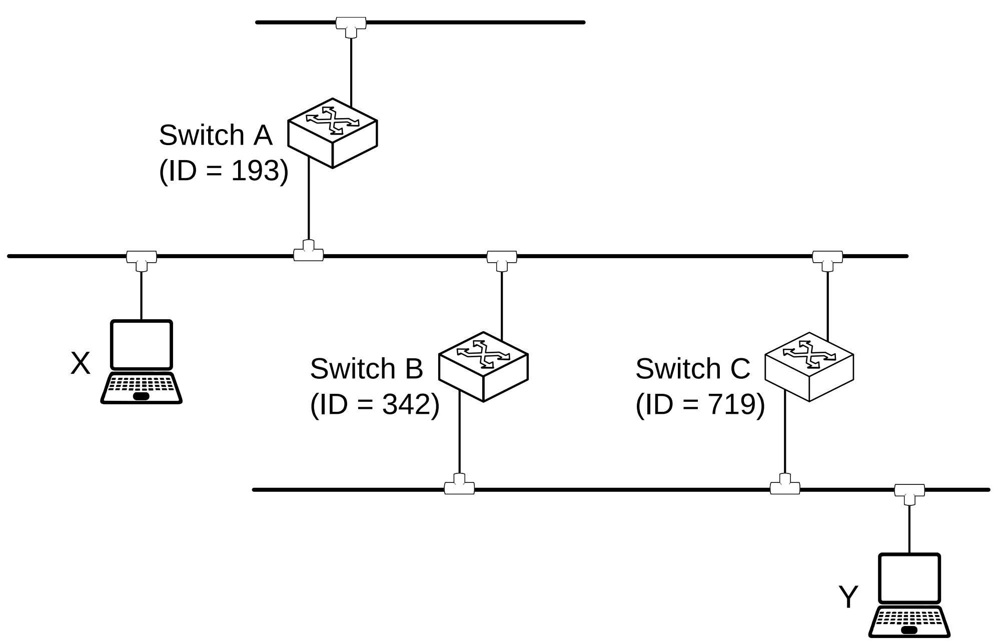

Start when the network is powered up and stop when the network stabilizes (i.e., only the root switch remains generating configuration frames). Assume that all switches send configuration messages in synchrony with each other (although in reality generally this is not the case). Do the following:

(a) List all the configuration messages sent by all switches until the network stabilizes. Recall that a configuration message carries ⟨source-ID, root-ID, root-distance⟩.
(b) For each switch, indicate which ports will be selected as "root," "designated," or "blocked" by the spanning tree protocol.
(c) How many iterations will take for the network to stabilize?
(d) After the network stabilizes, draw the path that a frame sent by station *X* will traverse to reach station *Y*.

Problem 1.43

Problem 1.44

Consider a local area network using the CSMA/CA protocol shown in Figure 1-33. Assume that three stations have frames ready for transmission and they are waiting for the end of a previous transmission. The stations choose the following backoff values for their first frame: STA1 = 5, STA2 = 9, and STA3=2. For their second frame backoff values, they choose: STA1 = 7, STA2 = 1, and STA3=4. For the third frame, the backoff values are STA1 = 3, STA2 = 8, and STA3=3. Show the timing diagram for the first 5 frames. Assume that all frames are of the same length.

Note: Compare the solution with that of Problem 1.20.

Problem 1.45

Consider an independent BSS (IBSS) with two mobile STAs, *A* and *B*, where each station has a single packet to send to the other one. Draw the precise time diagram for each station from the start to the completion of the packet transmission. For each station, select different packet arrival time and a reasonable number of backoff slots to count down before the station commences its transmission so that no collisions occur during the entire session. (Check Table 1-6 for contention window ranges.) Assume that both stations are using the basic transmission mode and only the first data frame transmitted is received in error (due to channel noise, not collision).

Problem 1.46

Chapter 2
Transmission Control Protocol (TCP)

Contents

2.1 Introduction

Transmission Control Protocol (TCP) is usually not associated with quality of service; but one could argue that TCP offers QoS in terms of assured delivery and efficient use of bandwidth, although it provides no delay guarantees. TCP is, after all, mainly about efficiency and adaptation: how to deliver data utilizing the maximum available (but fair) share of a dynamically changing network capacity so to reduce the delay. That is why our main focus here is only one aspect of TCP—congestion avoidance and control. The interested reader should consult additional sources for other aspects of TCP, e.g., [Stevens 1994; Peterson & Davie 2007; Kurose & Ross 2010]. We start quality-of-service review with TCP because it does not assume any knowledge of or any cooperation from the network. The network is essentially seen as a black box.

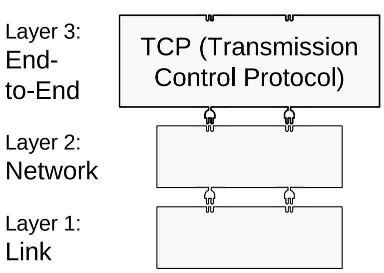

In Chapter 1 we have seen that pipelined ARQ protocols, such as Go-back-*N*, increase the utilization of network resources by allowing multiple packets to be simultaneously in transit (or, in flight) from sender to receiver. The "flight size" is controlled by a parameter called window size which must be set according to the available network resources. Remember that network is responsible for data from the moment it accepts them at sender's end until they are delivered at receiver's end. The network is storing the data for the "flight duration" and for this it must reserve resources, avoiding the possibility of becoming overbooked. In case of two end hosts connected by a single link, the optimal window size is easy to determine and remains static for the duration of session. However, this task is much more complex in a general multi-hop network.

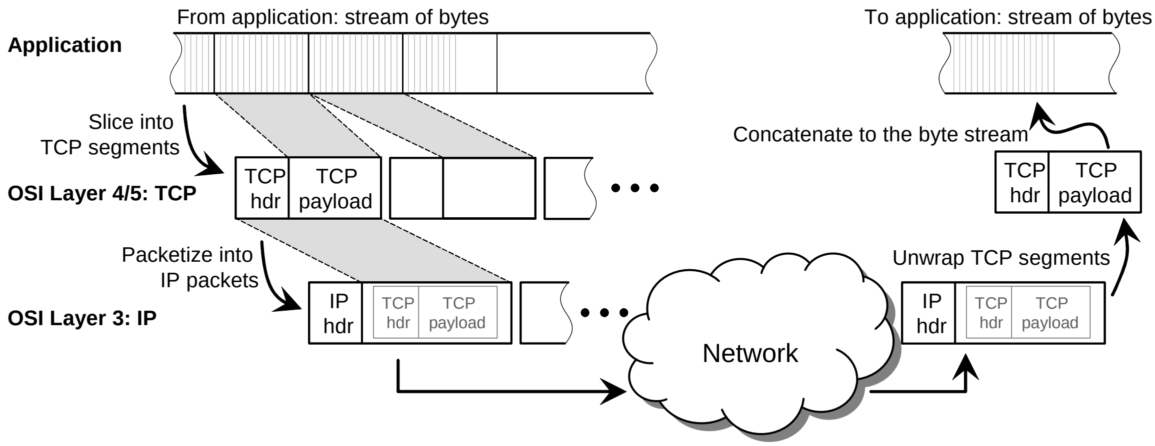

Figure 2-1: TCP accepts a stream of bytes as input from the application, slices it into segments, and passes to the IP layer as IP packets.

2.1.1 Reliable Byte Stream Service

TCP provides a *byte stream service*, which means that a stream of 8-bit bytes is exchanged across the TCP connection between the two applications. TCP does not automatically insert any delimiters of the data records. An application using TCP might "write" to it several times only to have the data compacted into a common segment and delivered as such to its peer. For example, if the application on one end writes 20 bytes, followed by a write of 30 bytes, followed by a write of 10 bytes, the application at the other end of the connection cannot tell what size the individual writes were. The other end may read the 60 bytes in two reads of 30 bytes at a time. One end puts a stream of bytes into TCP and the same, identical stream of bytes appears at the other end.

It is common to use the term "segment" for TCP packets. The TCP segments are encapsulated into IP packets and sent to the destination (Figure 2-1).

Like any other data packet, the TCP segment consists of the header and the data payload (Figure 2-2). The header consists of a 20-byte mandatory part, plus a variable-size options field. Most of regular TCP segments found on the Internet will have fixed 20-byte header and the options field is rarely used. The description of the header fields is as follows.

Source port number and **destination port number:** These numbers identify the sending and receiving applications on their respective hosts. A network application is rarely the sole "inhabitant" of a host computer; usually, the host runs multiple applications (processes), such as a web browser, email client, multimedia player, etc. Similar to an apartment building, where an apartment number is needed in addition to the street address to identify the recipient uniquely, the applications communicating over TCP are uniquely identified by their hosts' IP addresses and the applications' port numbers.

Sequence number: The 32-bit sequence number field identifies the position of the *first* data byte of this segment in the sender's byte stream during data transfer (when SYN bit is not set). Because TCP provides a byte-stream service, each byte of data has a sequence number.

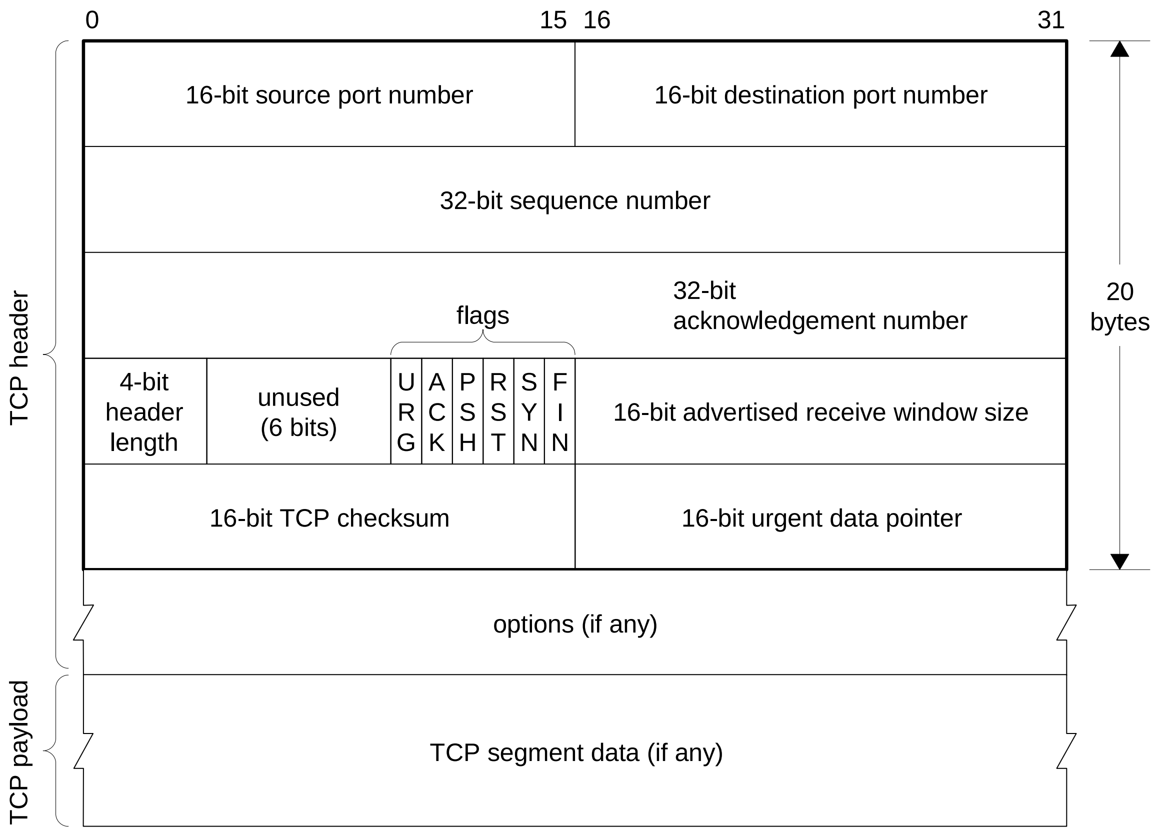

Figure 2-2: TCP segment format.

Acknowledgment number: The 32-bit acknowledgement number field identifies the sequence number of the next data byte that the receiver expects to receive. This field is valid only if the ACK bit is set; otherwise, it should be ignored by the recipient of this segment.

Header length: This field specifies the length of the TCP header in 32-bit words. This field is also known as the **Offset** field, because it informs the segment receiver where the data begins relative to the start of the segment. Regular header length is 20 bytes, so the default (and minimum allowed) value of this field equals 5. In case the options field is used, the value can be up to $4^2 - 1 = 15$, which means that the options field may contain up to $(15 - 5) \times 4 = 40$ bytes.

Unused: This field is reserved for future use and must be set to 0.

Flags: There are six bits allocated for the flags field, as follows:

URG: If this bit is set, the urgent data pointer field of the header is valid (described later).

ACK: When this bit is set, the acknowledgement number field of the header is valid and the recipient of this segment should pay attention to the acknowledgement number.

PSH: If this bit is set, it requires the TCP receiver to pass the received data to the receiving application immediately. Normally, this bit is not set and the TCP receiver may choose to buffer the received segment until it accumulates more data in the receive buffer.

RST: When set, this bit requires the TCP receiver to abort the connection because of some abnormal condition. For example, the segment's sender may have received a segment it did not expect to receive and wants to abort the connection.

SYN: This bit requests a connection (discussed later).

FIN: When set, this bit informs the TCP receiver that the sender does not have any more data to send. The sender can still receive data from the receiver until it receives a segment with the FIN bit set from the other direction.

Receive window size: This field specifies the number of bytes the sender is currently willing to accept. This field can be used to control the flow of data and congestion, as described later in Sections 2.1.3 and 2.2, respectively.

Checksum: This field helps in detecting errors in the received segments.

Urgent data pointer: When the URG bit is set, the value of this field should be added to the value of the sequence number field to obtain the location of the last byte of the "urgent data." The first byte of the urgent data is never explicitly defined. Because the TCP receiver passes data to the application in sequence, any data in the receive buffer up to the byte pointed by the urgent-data pointer may be considered urgent.

Options: The options field may be used to provide functions other than those covered by the regular header. This field may be up to 40 bytes long and if its length is not a multiple of 32 bits, extra padding bits should be added. The options field is used by the TCP sender and receiver at the connection establishment time, to exchange some special parameters, as described later.

The pseudo code in Listing 2-1 summarizes the TCP sender side protocol. In reality, both TCP sender and TCP receiver are implemented within the same TCP protocol module. Notice also that the method `send()` is part of sender's code, whereas the method `handle()` is part of receiver's code. However, to keep the discussion manageable, I decided to focus only on sender's side. See also Figure 2-3 for the explanation of the buffering parameters and Figure 2-8 for TCP sender's state diagram.

Listing 2-1: Pseudo code for RTO timer management in a TCP sender.

```
public class TCPSender {

    // window size that controls the maximum number of outstanding segments
    //     equation (2.3a) explains how EffectiveWindow is calculated
    private long effectiveWindow;

    // maximum segment size (MSS)
    private long MSS;

    // sequence number of the last byte sent thus far, initialized randomly
    private long lastByteSent;

    // sequence number of the last byte for which the acknowledgement
    //     from the receiver arrived thus far, initialized randomly
    private long lastByteAcked;
```

```
    // list of unacknowledged segments that may need to be retransmitted
    private ArrayList unacknowledgedSegments = new ArrayList();

    // network layer protocol that provides services to TCP protocol (normally, IP protocol)
    private ProtocolNetworkLayer networkLayerProtocol;

    // constructor
    public TCPSender(ProtocolNetworkLayer networkLayerProtocol) {
        this.networkLayerProtocol = networkLayerProtocol;

        lastByteSent = initial sequence number;
        lastByteAcked = initial sequence number;
    }

    // reliable byte stream service offered to an upper layer (or application)
    //     takes as input a long message ('data' input parameter) and transmits in segments
    public void send(
        byte[] data, String destinationAddr,
        ProtocolLayer_iUP upperProtocol
    ) throws Exception {
        // slice the application into segments of size MSS and send one-by-one
        for (i = 0; i < (data.length % MSS); i++) {

          // if the sender already used up the limit of outstanding packets, then wait
          if (effectiveWindow - unacknowledgedSegments.size() > 0){
           suspend this thread;
           wait until some ACKs are received;
        }

        // create a new TCP segment with sequence number equal to LastByteSent;
        // if (data.length < (i+1)*MSS), i.e., the remaining data slice is smaller than one MSS
        //     then use padding or Nagle's algorithm (described later)
        current_data_pointer = data + i*MSS;
        TCPSegment outgoingSegment =
            new TCPSegment(
              current_data_pointer, destinationAddr, upperProtocol
            );

        if (RTO timer not already running) {  start the timer;  }

        unacknowledgedSegments.add(outgoingSegment);
        lastByteSent += outgoingSegment.getLength();

        // hand the packet down the stack to IP for transmission
        // Note: outgoingSegment must be serialized to a byte-array as in Listing 1-1
        networkLayerProtocol.send(              // (omitted for clarity)
             outgoingSegment, destinationAddr, this
        );
    }

    // upcall method (called from the IP layer), when an acknowledgment is received
    public void handle(byte[] acknowledgement) {

        // acknowledgement carries the sequence number of
```

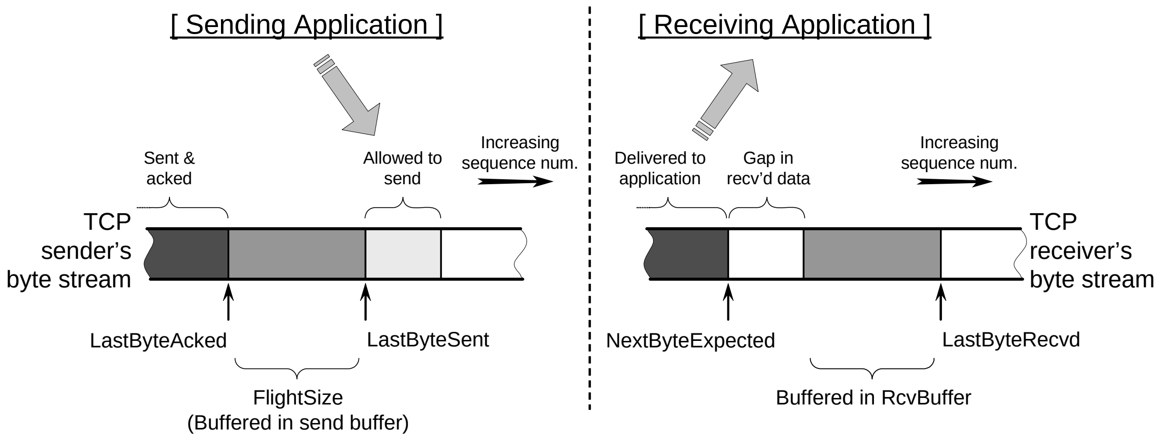

Figure 2-3: Parameters for TCP send and receive buffers.

```
32a         //       the next byte expected by the TCP receiver
33          if (acknowledgement.nextByteExpected > lastByteAcked) {
34              remove the acknowledged segment from
34a                 the unacknowledgedSegments array;
35              lastByteAcked = acknowledgement.nextByteExpected;
36              if (lastByteAcked < lastByteSent) {
37                  re-start the RTO timer;
38              } // i.e., there are segments not yet acknowledged
39          }
40      }

41      // this method is called when the RTO timer expires (timeout)
41a     // this event signifies segment loss, and the oldest unacknowledged segment is retransmitted
42      public void RTOtimeout() {
43          retrieve the segment with sequence_number == LastByteAcked
43a             from unacknowledgedSegments array and retransmit it;

43          double the TimeoutInterval;
44          start the timer;
45      }
46  }
```

The code description is as follows: … to be described … The reader should compare Listing 2-1 to Listing 1-1 in Section 1.1.4 for a generic protocol module.

The method `handle()`, which starts on Line 31, normally handles bidirectional traffic and processes TCP segments that are received from the other end of the TCP connection. Recall also that TCP acknowledgements are piggybacked on TCP segments (Figure 2-2).

Listing 2-1 provides only a basic skeleton of a TCP sender. The details will be completed in the following sections as we learn more about different issues.

2.1.2 Retransmission Timer

Problems related to this section: Problem 2.1 → ??, Problem 2.13, and Problem 2.15

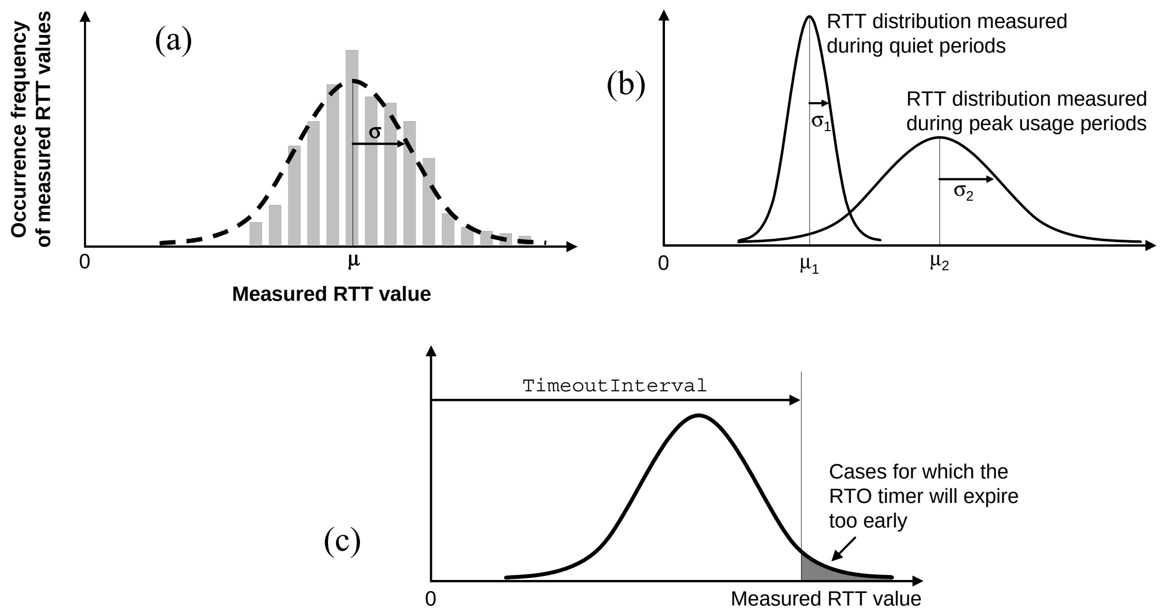

Figure 2-4: Distribution of measured round-trip times for a given sender-receiver pair.

An important parameter for reliable transport over multihop networks is retransmission timer. This timer triggers the retransmission of packets that are presumed lost. Obviously, it is very important to set the right value for the timer. For, if the timeout time is too short, the packets will be unnecessarily retransmitted thus wasting the network bandwidth. And, if the timeout time is too long, the sender will unnecessarily wait when it should have already retransmitted thus underutilizing and perhaps wasting the network bandwidth.

It is relatively easy to set the timeout timer for single-hop networks because the propagation time remains effectively constant. However, in multihop networks queuing delays at intermediate routers and propagation delays over alternate paths introduce significant uncertainties.

TCP has a special algorithm for dynamically updating the retransmission timeout (RTO) value. The details are available in RFC-2988 [Paxson & Allman, 2000], and here is a summary. The RTO timer value, denoted as `TimeoutInterval`, is initially set as 3 seconds. When the retransmission timer expires (presumably because of a lost packet), the earliest unacknowledged data segment is retransmitted and the next timeout interval is set to twice the previous value:

$$\texttt{TimeoutInterval}(t) = 2 \times \texttt{TimeoutInterval}(t{-}1) \tag{2.1}$$

This property of doubling the RTO on each timeout is known as **exponential backoff**.[12] If a segment's acknowledgement is received before the retransmission timer expires, the TCP sender measures the round-trip time (RTT) for this segment, denoted as `SampleRTT`. TCP only measures `SampleRTT` for segments that have been transmitted *once* and *not* for segments that have been retransmitted.

[12] Recall the discussion from Section 1.3.3 above. In the TCP case, the sender assumes that concurrent TCP senders are contending for the network resources (router buffers), thereby causing congestion and packet loss. To reduce the congestion, the sender doubles the retransmission delay by doubling its RTO.

Suppose you want to determine the statistical characteristics of round-trip time values for a given sender-receiver pair. As illustrated in Figure 2-4(a), the histogram obtained by such measurement can be approximated by a normal distribution[13] $N(\mu, \sigma)$ with the mean value μ and the standard deviation σ. If the measurement were performed during different times of the day, the obtained distributions may look quite different, Figure 2-4(b). Therefore, the timeout interval should be *dynamically adapted* to the observed network condition. The remaining decision is about setting the `TimeoutInterval`. As illustrated in Figure 2-4(c), there will always be some cases for which any finite `TimeoutInterval` is too short and the acknowledgement will arrive after the timeout already expired. Setting `TimeoutInterval` = $\mu + 4\sigma$ will cover nearly 100 % of all the cases. Therefore, for the subsequent data segments, the `TimeoutInterval` is set according to the following equation:

$$\mathtt{TimeoutInterval}(t) = \mathtt{EstimatedRTT}(t) + 4 \cdot \mathtt{DevRTT}(t) \qquad (2.2)$$

where `EstimatedRTT`(t) is the currently estimated mean value of RTT:

$$\mathtt{EstimatedRTT}(t) = (1-\alpha) \cdot \mathtt{EstimatedRTT}(t-1) + \alpha \cdot \mathtt{SampleRTT}(t)$$

The initial value is set as `EstimatedRTT`(0) = `SampleRTT`(0) for the first RTT measurement. This approach to computing the running average of a variable is called Exponential Weighted Moving Average (EWMA). Similarly, the current standard deviation of RTT, `DevRTT`(t), is estimated as:

$$\mathtt{DevRTT}(t) = (1-\beta) \cdot \mathtt{DevRTT}(t-1) + \beta \cdot \left|\mathtt{SampleRTT}(t) - \mathtt{EstimatedRTT}(t-1)\right|$$

The initial value is set as $\mathtt{DevRTT}(0) = \frac{1}{2}\mathtt{SampleRTT}(0)$ for the first RTT measurement. The recommended values of the control parameters α and β are $\alpha = 0.125$ and $\beta = 0.25$. These values were determined empirically.

In theory, it is simplest to maintain individual retransmission timer for each outstanding packet. In practice, timer management involves considerable complexity, so most protocol implementations maintain single timer per sender. RFC-2988 recommends maintaining *single* retransmission timer per TCP sender, even if there are multiple transmitted-but-not-yet-acknowledged segments. Of course, individual implementers may decide otherwise, but in this text, we follow the single-timer recommendation for TCP.

TCP sends segments in *bursts* (or, groups of segments), every burst containing the number of segments limited by the current window size. Recall from Section 1.3.2 that in all sliding window protocols, the sender is allowed to have only up to the window-size outstanding amount of data (yet to be acknowledged). The same holds for the TCP sender. Once the window-size worth of segments is sent, the sender stops and waits for acknowledgements to arrive. For every arriving ACK, the sender is allowed to send certain number of additional segments, as governed by the rules described later. The retransmission timer management is included in the pseudo code in Listing 2-1 (Section 2.1.1). The following summary extracts and details the key points of the retransmission timer management from Listing 2-1:

---

[13] In reality, multimodal RTT distributions (i.e., with several peaks) are observed. The interested reader can find relevant links at this book's website—follow the link "Related Online Resources," then look under the topic of "Network Measurement."

Search the Web for RTT distribution measurement

```
   In method TCPSender.send(), Listing 2-1 // called by application layer above
    in Line 24:
       if (RTO timer not already running) {
           set the RTO timer to the current value
               as calculated in methods handle() and RTOtimeout();
           start the timer;
       }

In method TCPSender.handle(), Listing 2-1 // called by IP layer when ACK arrives
    in Lines 36 - 38:
        if ((lastByteAcked < lastByteSent) {
            calculate the new value of the RTO timer using Eq. (2.2);
            re-start the RTO timer;
        }

In method TCPSender.RTOtimeout(), Listing 2-1 // called when RTO timer timeout
    in Line 43:
           double the TimeoutInterval; // see Eq. (2.1)
           start the timer;
```

An important peculiarity to notice about TCP is as follows. When a window-size worth of segments is sent, the timer is set for the first one, assuming that the timer is not already running (Line 24 in Listing 2-1). For every acknowledged segment of the burst, the timer is restarted for its subsequent segment in Line 37 in Listing 2-1. Thus, the actual timeout time for the segments towards the end of a burst can run quite longer than for those near the beginning of the burst. An example will be seen later in Section 2.2 in the solution of Example 2.1.

2.1.3 Flow Control

TCP receiver accepts out-of-order segments, but they are buffered and not delivered to the application above the TCP layer before the gaps are filled. For this, the receiver allocates memory space of the size `RcvBuffer`, which is typically set to 4096 bytes, although older versions of TCP set it to 2048 bytes. The *receive buffer* is used to store in-order segments as well, because the application may be busy with other tasks and does not fetch the incoming data immediately. For the sake of simplicity, in the following discussion we will assume that in-order segments are immediately fetched by the application, unless stated otherwise.

To avoid having its receive buffer overrun, the receiver continuously advertises the remaining buffer space to the sender using a field in the TCP header; we call this variable `RcvWindow`. It is dynamically changing to reflect the current occupancy state of the receiver's buffer. The sender should never have more than the current `RcvWindow` amount of data outstanding. This process is called *flow control*. Figure 2-5 illustrates the difference between the flow control as opposed to congestion control, which is described later in Section 2.2.

Figure 2-6 shows how an actual TCP session might look like. The notation 0:512(512) means transmitted data bytes 1 through but *not* included 512, which is a total of 512 bytes. The first action is to establish the session, which is done by the first three segments, which represent the **three-way handshake** procedure. Here I briefly summarize the three-way handshake. The

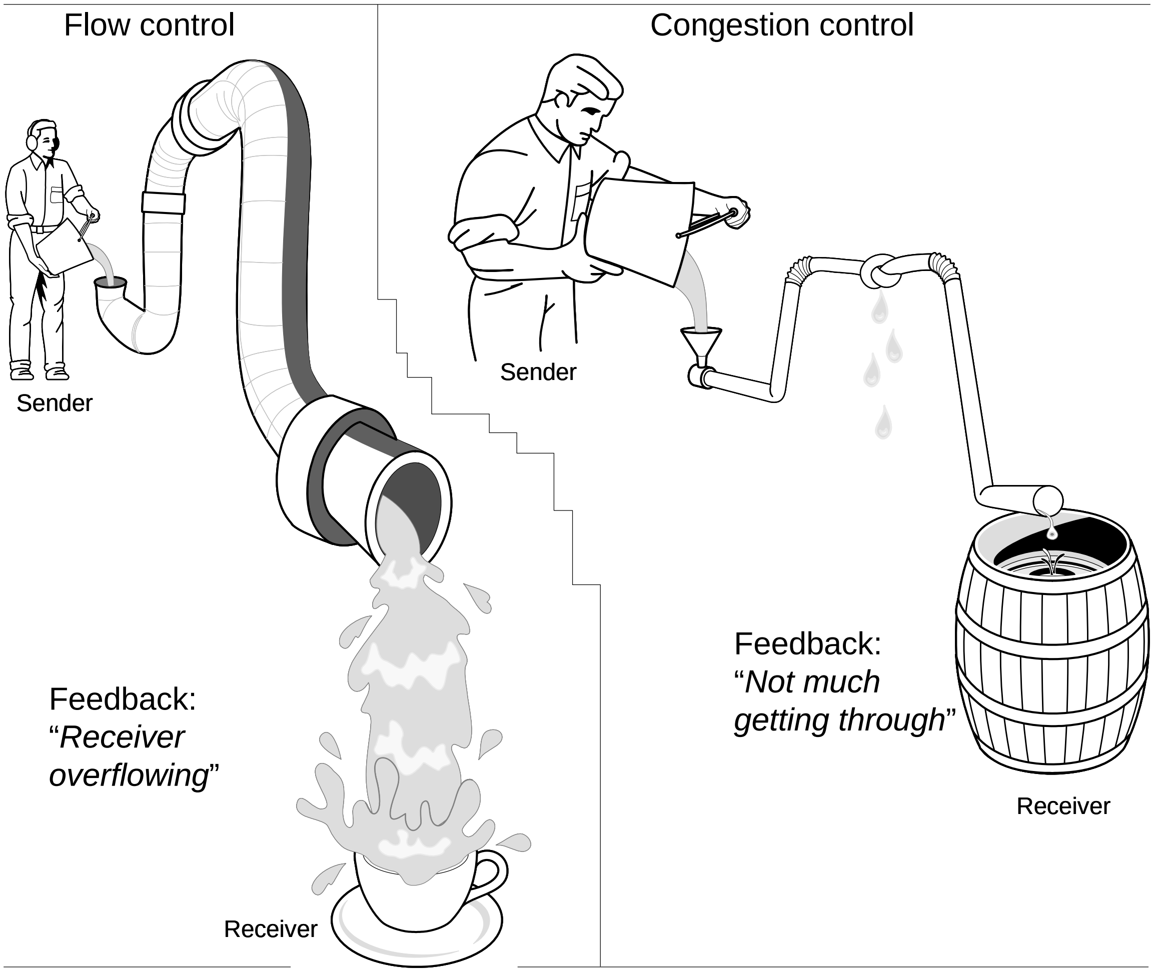

Figure 2-5: Flow control compared to congestion control.

interested reader should consult another source for more details, e.g., [Stevens 1994; Peterson & Davie 2007; Kurose & Ross 2010].

The first three segments are special in that they do not contain data (i.e., they have only a header), and the SYN flag in the header is set (Figure 2-2). In this example, the client offers `RcvWindow` = 2048 bytes, and the server offers `RcvWindow` = 4096 bytes. In our case, the client happens to be the "sender," but server or both client and server can simultaneously be senders and receivers. They also exchange the size of the future segments, `MSS` (to be described later, Table 2-1), and settle on the smaller one of 1024 and 512 bytes, i.e., 512 bytes. During the connection-establishment phase, the client and the server will transition through different states, such as LISTEN, SYN_SENT, and ESTABLISHED (marked in Figure 2-6 on the right-hand side). As stated earlier, both sender and receiver instantiate their sequence numbers randomly. In Figure 2-6, the sender selects 122750000, while the receiver selects 2363371521. Hence, the sender sends a zero-bytes segment

```
122750000:122750000(0)
```

The receiver acknowledges it and sends its own initial sequence number by sending

```
2363371521:2363371521(0); ack 122750000
```

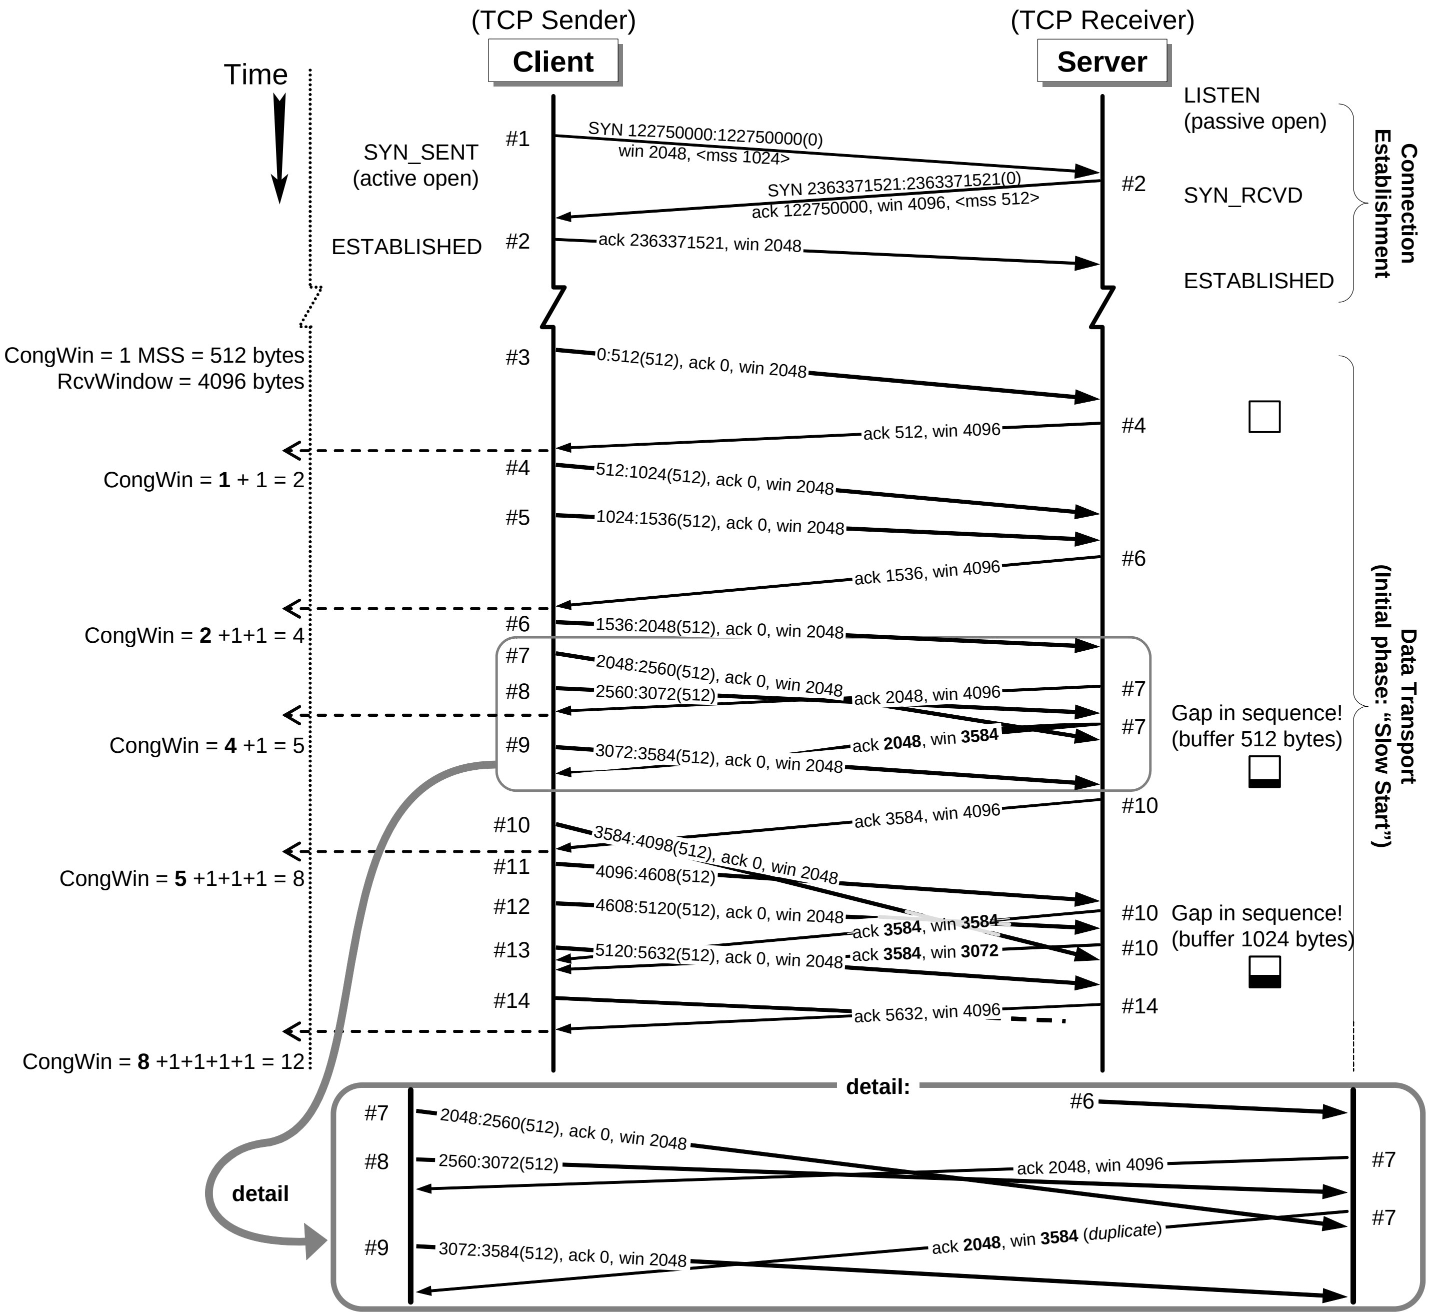

Figure 2-6: Initial part of the time line of an example TCP session. Time increases down the page. See text for details. (The CongWin parameter on the left side of the figure will be described later in Section 2.2.)

i.e., it sends zero data bytes and acknowledges zero data bytes received from the sender (the ACK flag is set). To avoid further cluttering the diagram, I am using these sequence numbers only for the first three segments. For the remaining segments, I simply assume that the sender starts with the sequence number equal to zero. In Figure 2-6, the server sends no data to the client, so the sender keeps acknowledging the first segment from the receiver and acknowledgements from

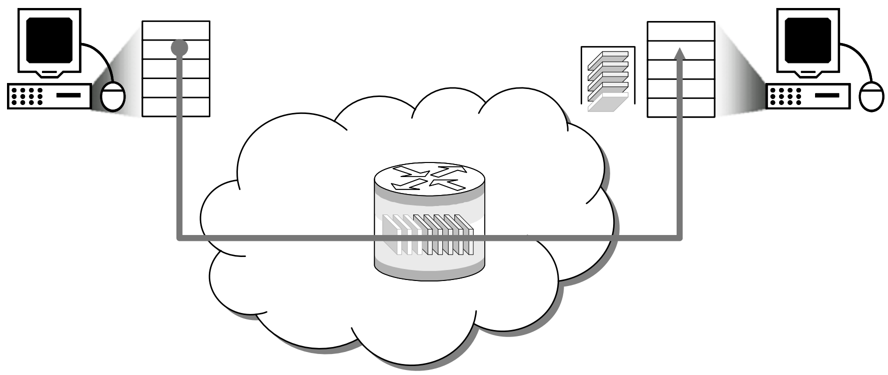

Figure 2-7: Simple congestion-control scenario for TCP.

sender to receiver carry the value 0 for the sequence number (recall that the actual value of the server's sequence number is 2363371521).

After establishing the connection, the sender starts sending packets. Figure 2-6 illustrates how TCP incrementally increases the number of outstanding segments. This procedure is called **slow start**, and it will be described later. TCP assigns byte sequence numbers, but for simplicity we usually show packet sequence numbers. Notice that the receiver is *not* obliged to acknowledge *individually* every single in-order segment—it can use *cumulative ACKs* to acknowledge several of them up to the most recent contiguously received data. Conversely, the receiver *must* immediately generate (duplicate) ACK—*dupACK*—for every out-of-order segment, because dupACKs help the sender detect segment loss (as described in Section 2.2).

Notice that receiver might send dupACKs even for successfully transmitted segments because of random re-ordering of segments in the network. This is the case with segment #7 (detail at the bottom of Figure 2-6), which arrives after segment #8. Thus, if a segment is delayed further than three or more of its successors, the duplicate ACKs will trigger the sender to re-transmit the delayed segment, and the receiver may eventually receive a duplicate of such a segment.

2.2 Congestion Control

TCP maneuvers to avoid congestion in the first place, and controls the damage if congestion occurs. The key characteristic of TCP is that all the intelligence for congestion avoidance and control is in the end hosts—no help is expected from the intermediary hosts.

A key problem addressed by the TCP protocol is to determine the optimal window size *dynamically* in the presence of uncertainties and dynamic variations of available network resources.

Early versions of TCP would start a connection with the sender injecting multiple segments into the network, up to the window size advertised by the receiver. The problems would arise due to intermediate router(s), which must queue the packets before forwarding them. If that router runs out of memory space, large number of packets would be lost and had to be retransmitted.

Jacobson [1988] showed how this naïve approach could reduce the throughput of a TCP connection drastically.

The problem is illustrated in Figure 2-7, where the whole network is abstracted as a single *bottleneck router*. It is easy for the receiver to know about its own available buffer space and advertise the right window size to the sender (denoted by `RcvWindow`). The problem is with the intermediate router(s), which serve data flows between many sources and receivers. Bookkeeping and policing of fair use of router's resources is a difficult task, because router must forward the packets as quickly as possible, and it is practically impossible to dynamically determine the "right window size" of the router's memory allocated for each flow and advertise it back to the sender.

TCP approaches this problem by putting the entire burden of determining the right window size of the bottleneck router onto the end hosts. Essentially, the sender dynamically probes the network and adjusts the amount of data in flight to match the bottleneck resource. The algorithm used by TCP sender can be summarized as follows:

---

1. Start with a small size of the sender window
2. Send a burst (size of the current sender window) of packets into the network
3. Wait for feedback about success rate (acknowledgements from the receiver end)
4. When feedback obtained:
    a. If the success rate is greater than zero, *increase* the sender window size and go to Step 2
    b. If loss is detected, *decrease* the sender window size and go to Step 2

This simplified procedure will be elaborated as we present the details in the following text. It is important to notice that TCP *controls* congestion in the sense that it first needs to cause congestion, next to observe it though the feedback, and then to react by reducing the input. This cycle is repeated in a never-ending loop. Section 2.4 describes other variants of TCP that try to avoid congestion, instead of causing it (and then controlling it).

Table 2-1 shows the most important parameters (all the parameters are maintained in integer units of bytes). Buffering parameters are shown in Figure 2-3. Figure 2-8 and Figure 2-9 summarize the algorithms run at the sender and receiver. These are digested from RFC 2581 and RFC 2001 and the reader should check the details on TCP congestion control in [Allman *et al*. 1999; Stevens 1997]. [Stevens 1994] provides a detailed overview with traces of actual runs.

Table 2-1. TCP congestion control parameters (measured in integer number of bytes). Also see Figure 2-3.

Variable	Definition
`MSS`	The size of the largest segment that the sender can transmit. This value can be based on the maximum transmission unit (MTU) of the network, the path MTU discovery algorithm, or other factors. The size does *not* include the TCP/IP headers and options. [Note that RFC 2581 distinguishes sender maximum segment size (`SMSS`) and receiver maximum segment size (`RMSS`).]

RcvWindow	The size of the most recently advertised receiver window.
CongWindow	Sender's current estimate of the available buffer space in the bottleneck router.
LastByteAcked	The highest sequence number currently acknowledged.
LastByteSent	The sequence number of the last byte the sender sent.
FlightSize	The amount of data that the sender has sent, but not yet had acknowledged.
EffectiveWindow	The maximum amount of data that the sender is currently allowed to send. At any given time, the sender must not send data with a sequence number higher than the sum of the highest acknowledged sequence number and the minimum of CongWindow and RcvWindow.
SSThresh	The slow start threshold used by the sender to decide whether to employ the slow-start or congestion-avoidance algorithm to control data transmission.

Notice that the sender must assure at all times that:

$$\texttt{LastByteSent} \le \texttt{LastByteAcked} + \min\{\texttt{CongWindow}, \texttt{RcvWindow}\}$$

Therefore, the amount of unacknowledged data (denoted as FlightSize) should not exceed this value at any time:

$$\texttt{FlightSize} = \texttt{LastByteSent} - \texttt{LastByteAcked} \le \min\{\texttt{CongWindow}, \texttt{RcvWindow}\}$$

At any moment during a TCP session, the maximum amount of data the TCP sender is allowed to send is (marked as "*allowed to send*" in Figure 2-3):

$$\texttt{EffectiveWindow} = \min\{\texttt{CongWindow}, \texttt{RcvWindow}\} - \texttt{FlightSize} \qquad (2.3a)$$

Here we assume that the sender can only send MSS-size segments; the sender holds with transmission until it collects at least an MSS worth of data. This is not always true, and the application can request speedy transmission, thus generating small packets, so called *tinygrams*. The application does this using the TCP_NODELAY socket option, which sets PSH flag (Figure 2-2). This is particularly the case with interactive applications, such as telnet or secure shell. Nagle's algorithm [Nagle 1984] constrains the sender to have unacknowledged at most one segment smaller than one MSS. For simplicity, we assume that the effective window is always rounded down to integer number of MSS-size segments:

$$\texttt{EffectiveWindow} = \lfloor \min\{\texttt{CongWindow}, \texttt{RcvWindow}\} - \texttt{FlightSize} \rfloor \qquad (2.3b)$$

Figure 2-6 illustrates the TCP slow start phase. In **slow start**, CongWindow starts at one segment and gets incremented by one segment every time an ACK is received. As it can be seen, this opens the congestion window exponentially: send one segment, then two, four, eight and so on.

The only "feedback" TCP receives from the network is by having packets lost in transport. TCP considers that these are solely lost to congestion, which, of course, is not necessarily true—packets may be lost to channel noise or even to a broken link. A design is good as long as its assumptions hold and TCP works fine over wired networks, because other types of loss are uncommon therein. However, in wireless networks, this underlying assumption breaks and it causes a great problem as will be seen later in Section 2.5.

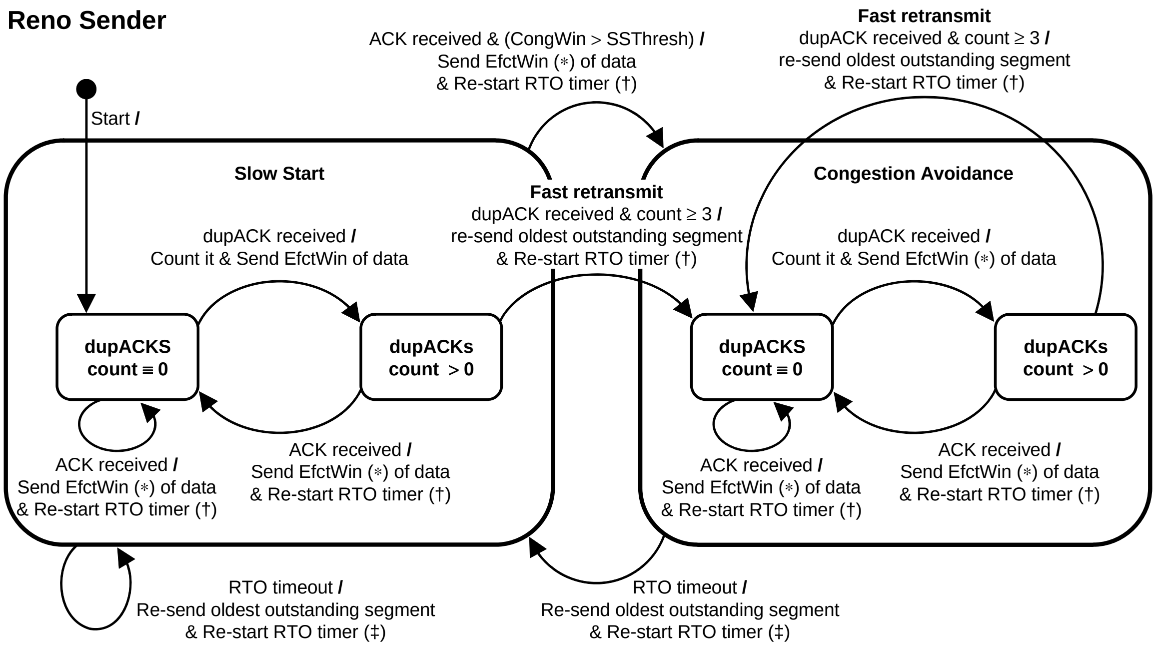

Figure 2-8: TCP Reno sender state diagram. (∗) Effective window depends on CongWin, which is computed differently in slow-start vs. congestion-avoidance. (†) RTO timer is restarted if LastByteAcked < LastByteSent. (‡) RTO size doubles, SSThresh = CongWin/2.

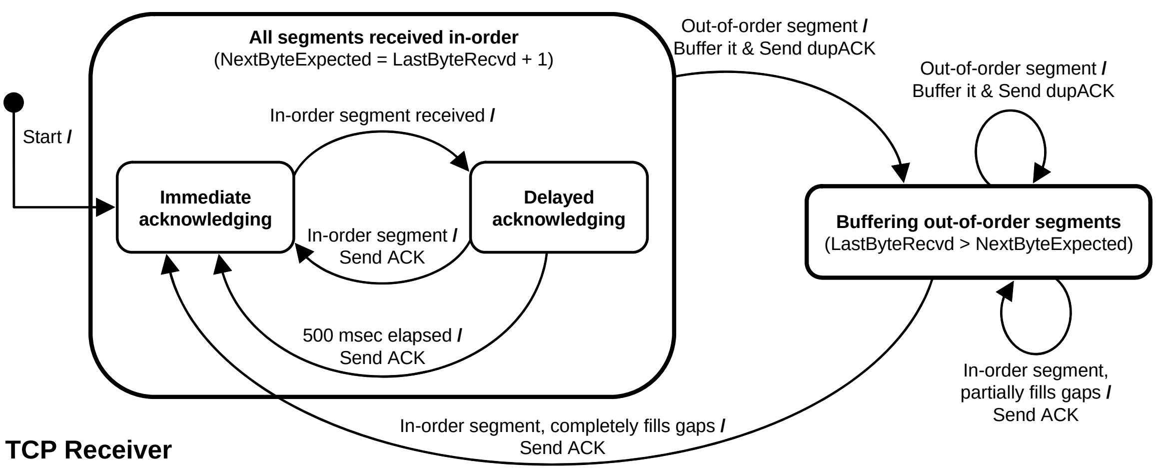

Figure 2-9: TCP receiver state diagram.

As far as TCP is concerned, it does not matter when a packet *loss happened* (somewhere in the network, on a router); what matters is when the *loss is detected* (at the TCP sender). Packet loss happens in the network and the network is not expected to notify the TCP endpoints about the loss—the endpoints have to detect loss *indirectly* and deal with it on their own. Packet loss is of little concern to TCP receiver, except that it buffers out-of-order segments and waits for the gap in sequence to be filled. TCP sender is the one mostly concerned about the loss and the one that takes actions in response to detected loss. TCP sender detects loss via two types of events (whichever occurs first):

1. Timeout timer expiration

2. Reception of three[14] duplicate ACKs (*four* identical ACKs without the arrival of any other intervening packets)

Upon detecting the loss, TCP sender takes action to *avoid* further loss by reducing the amount of data injected into the network. (TCP also performs *fast retransmission* of what appears to be the lost segment, without waiting for a RTO timer to expire.) There are many versions of TCP, each having different reaction to loss. The two most popular ones are **TCP Tahoe** and **TCP Reno**, of which TCP Reno is more recent and currently prevalent in the Internet. Table 2-2 shows how they detect and handle segment loss.

Table 2-2: How different TCP senders detect and deal with segment loss.

Event	TCP Version	TCP Sender's Action
Timeout	Tahoe Reno	Set `CongWindow` = 1×`MSS`
≥ 3×dup ACKs	Tahoe	Set `CongWindow` = 1×`MSS`
	Reno	Set `CongWindow` = max {⌊½ `FlightSize`⌋, 2×`MSS`} + 3×`MSS`

As seen in Table 2-2, different versions react differently to three dupACKs: the more recent version of TCP, i.e., TCP Reno, reduces the congestion window size to a lesser degree than the older version, i.e., TCP Tahoe. The reason is that researchers realized that three dupACKs signalize lower degree of congestion than RTO timeout. If the RTO timer expires, this may signal a "severe congestion" where nothing is getting through the network. Conversely, three dupACKs imply that three packets got through, although out of order, so this signals a "mild congestion."

The initial value of the slow start threshold `SSThresh` is commonly set to 65535 bytes = 64 KB. When a TCP sender detects segment loss using the retransmission timer, the value of `SSThresh` *must* be set to no more than the value given as:

$$\texttt{SSThresh} = \max\{\lfloor \tfrac{1}{2}\,\texttt{FlightSize} \rfloor, 2\times\texttt{MSS}\} \tag{2.4}$$

where `FlightSize` is the amount of outstanding data in the network (for which the sender has not yet received an acknowledgement). The floor operation $\lfloor\cdot\rfloor$ rounds the first term down to the next multiple of `MSS`. Notice that some networking books and even TCP implementations state that, after a loss is detected, the slow start threshold is set as `SSThresh` = ½ `CongWindow`, which according to RFC-2581 is *incorrect*.[15]

---

[14] The reason for *three* dupACKs is as follows. Because TCP does not know whether a lost segment or just a reordering of segments causes a dupACK, it waits for a small number of dupACKs to be received. It is assumed that if there is just a reordering of the segments, there will be only one or two dupACKs before the reordered segment is processed, which will then generate a fresh ACK. Such is the case with segments #7 and #10 in Figure 2-6. If three or more dupACKs are received in a row, it is a strong indication that a segment has been lost.

[15] The formula `SSThresh` = ½ `CongWindow` is an older version for setting the slow-start threshold, which appears in RFC-2001 as well as in [Stevens 1994]. I surmise that it was regularly used in TCP Tahoe, but should not be used with TCP Reno.

Congestion can occur when packets arrive on a big pipe (a fast LAN) and are sent out a smaller pipe (a slower WAN). Congestion can also occur when multiple input streams arrive at a router whose output capacity (transmission speed) is less than the sum of the input capacities. Here is an example:

Example 2.1 Congestion Due to Mismatched Pipes with Limited Router Resources

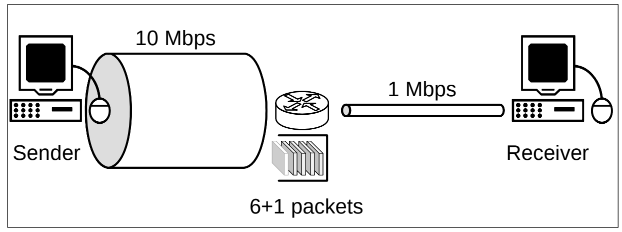

Consider an FTP application that transmits a huge file (e.g., 20 MBytes) from host A to B over the two-hop path shown in the figure. The link between the router and the receiver is called the "bottleneck" link because it is much slower than any other link on the sender-receiver path. Assume that the router can always allocate the buffer size of only six packets for our session and in addition have one of our packets currently being transmitted. Packets are only dropped when the buffer fills up. We will assume that there is no congestion or queuing on the path taken by ACKs.

Assume `MSS` = 1KB and a constant `TimeoutInterval` = 3×RTT = 3×1 sec. Draw the graphs for the values of `CongWindow` (in KBytes) over time (in RTTs) for the first 20 RTTs if the sender's TCP congestion control uses the following:

(a) TCP Tahoe: Additive increase / multiplicative decrease and slow start and fast retransmit.

(b) TCP Reno: All the mechanisms in (a) plus fast recovery.

Assume a large `RcvWindow` (e.g., 64 KB) and error-free transmission on all the links. Assume also that duplicate ACKs do not trigger growth of the `CongWindow` (i.e., only regular ACKs increase the `CongWindow` size). Finally, to simplify the graphs, assume that all ACK arrivals occur exactly at unit increments of RTT and that the associated `CongWindow` update occurs exactly at that time, too.

The solutions for (a) and (b) are shown in Figure 2-10 through Figure 2-15. The discussion of the solutions is in the following text. Notice that, unlike Figure 2-6, the transmission rounds are "clocked" and neatly aligned to the units of RTT. This idealization is only for the sake of illustration and the real world would look more like Figure 2-6. [Note that this idealization would stand in a scenario with propagation delays much longer than transmission delays.]

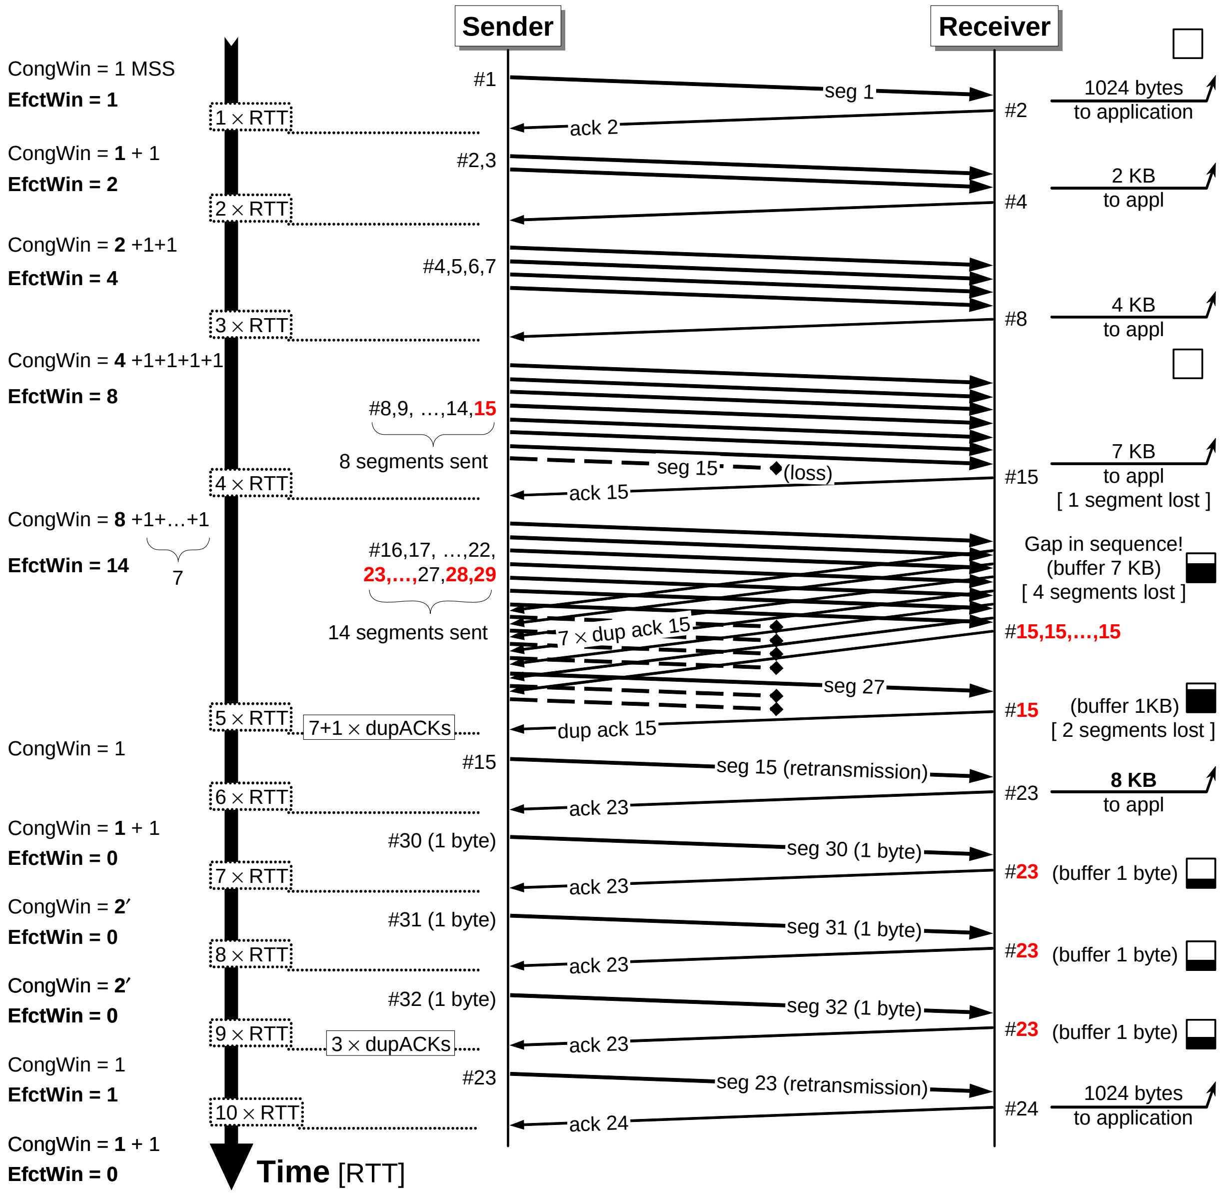

Figure 2-10: TCP Tahoe—partial timeline of segment and acknowledgement exchanges for Example 2.1. Shown on the sender's side are ordinal numbers of the sent segments and on the receiver's side are those of the ACKs (which indicate the next expected segment).

Let us first consider what happens at the router, as illustrated in Figure 2-11. The reader should recall the illustration in Figure 1-17, which shows that packets are first completely received at the link layer before they are passed up the protocol stack (to IP and on to TCP). The link speeds are mismatched by a factor of 10 : 1, so the router will transmit only a single packet on the second link while the sender already transmitted ten packets on the first link. Normally, this would only cause delays at the router, but with limited router resources there is also a loss of packets. This is detailed in Figure 2-11, where the three packets in excess of the router buffer capacity are discarded (numbered #23, #24, and #25). Thereafter, until the queue slowly drains, the router has

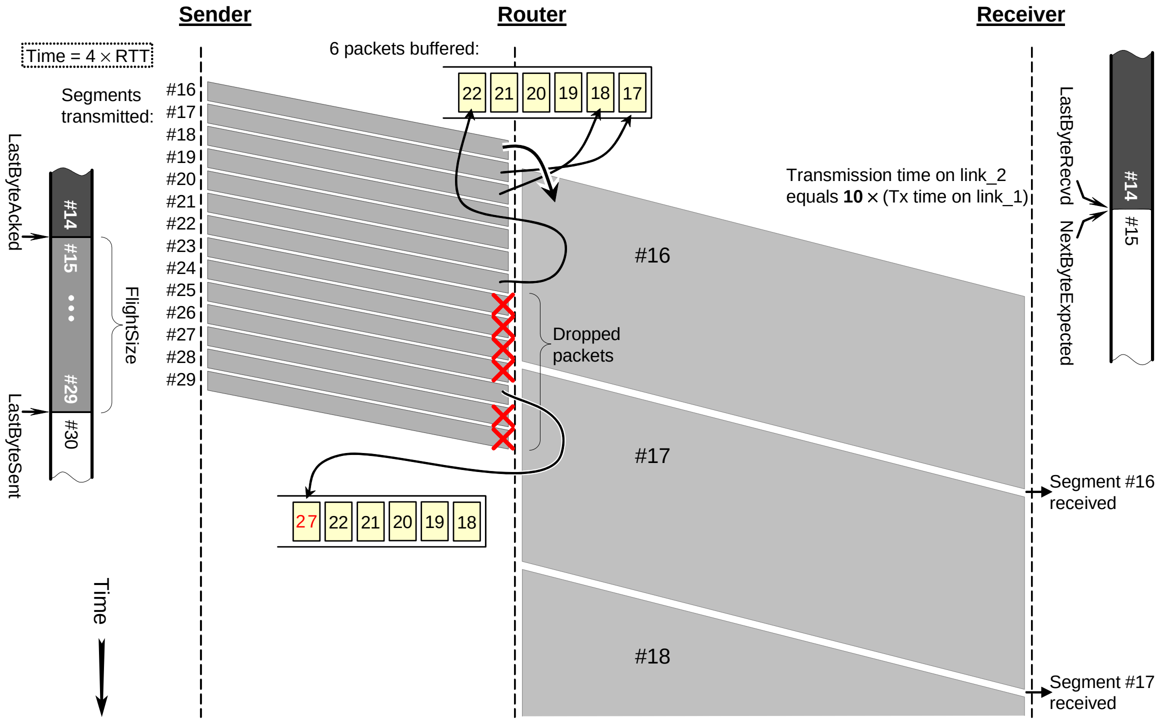

Figure 2-11: Detail from Figure 2-10 starting at time = 4×RTT. Mismatched transmission speeds result in packet loss at the bottleneck router.

one buffer slot available for every ten new packets that arrive. More details about how routers forward packets are available in Section 4.1.

It is instructive to observe how the retransmission timer is managed (Figure 2-12). Up to time = 4×RTT, the timer is always reset for the next burst of segments. However, at time = 4×RTT the timer is set for the 15th segment, which was sent in the same burst as the 8th segment, and not for the 16th segment because the acknowledgement for the 15th segment is still missing. The reader is encouraged to inspect the timer management for all other segments in Figure 2-12.

2.2.1 TCP Tahoe

Problems related to this section: Problem 2.2 → Problem 2.7 and Problem 2.9 → Problem 2.12

TCP sender begins with a congestion window equal to one segment and incorporates the slow start algorithm. In ***slow start*** the sender follows a simple rule: *For every acknowledged segment, increment the congestion window size by one* `MSS` (unless the current congestion window size exceeds the `SSThresh` threshold, as described later in this section). This procedure continues until a segment loss is detected. Of course, a duplicate acknowledgement does not contribute to increasing the congestion window size.

When the sender receives a dupACK, it does nothing but count it. If this counter reaches three or more dupACKs, the sender decides, by inference, that a loss occurred. In response, it adjusts the congestion window size and the slow-start threshold (`SSThresh`), and re-sends the oldest unacknowledged segment. (The dupACK counter also should be reset to zero.) As shown in Figure 2-12, the sender detects the loss first time at the fifth transmission round, i.e., at 5×RTT,

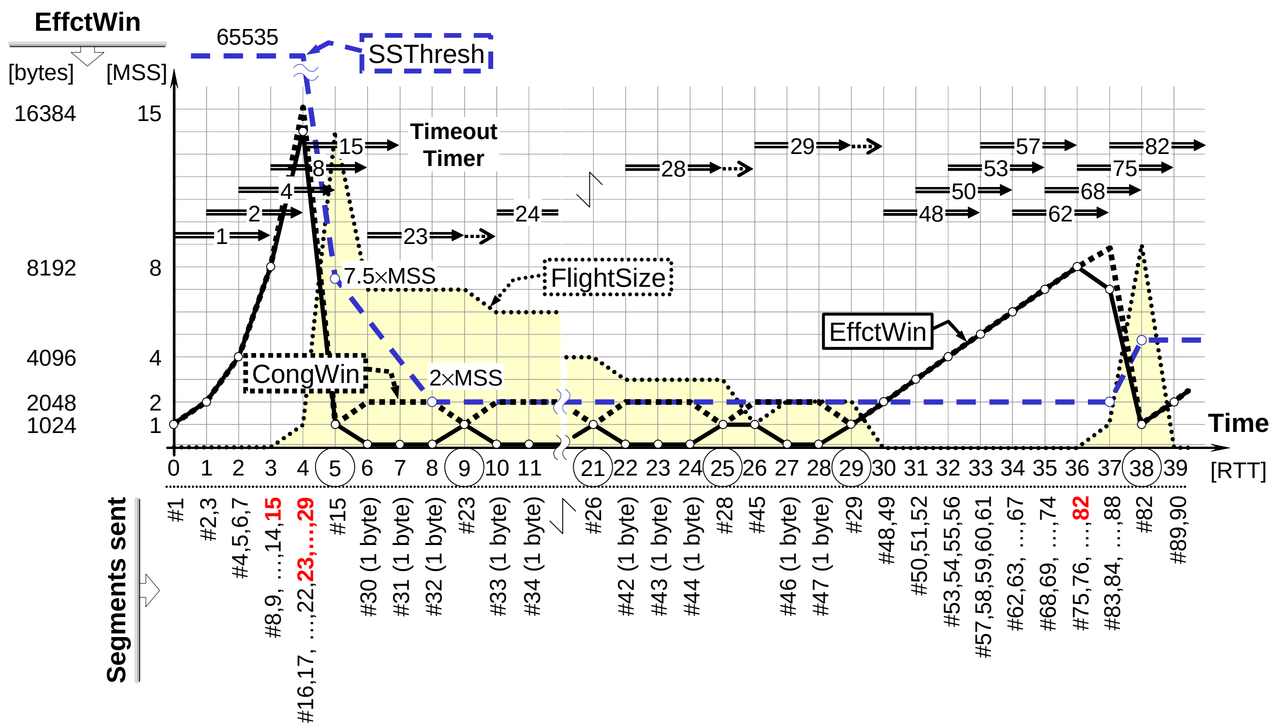

Figure 2-12: TCP Tahoe sender—the evolution of the effective and congestion window sizes for Example 2.1. The sizes are given on vertical axis (left) both in bytes and `MSS` units.

by receiving eight duplicate ACKs. The congestion window size at this instance is equal to 15360 bytes or 15×`MSS`. After detecting a segment loss, the sender sharply reduces the congestion window size in accordance with TCP's ***multiplicative decrease*** behavior. As explained earlier (Table 2-2), a Tahoe sender resets `CongWin` to one `MSS` and reduces `SSThresh` as given by Eq. (2.4). Just before the moment the sender received eight dupACKs `FlightSize` equaled 15, so the new value of `SSThresh` = 7.5×`MSS` is set.

Notice that in TCP Tahoe any additional dupACKs in excess of three do not matter—no new packet can be transmitted while additional dupACKs after the first three are received. As will be seen later, TCP Reno sender differs from TCP Tahoe sender in that it starts *fast recovery* based on the additional dupACKs received after the first three.

Upon completion of multiplicative decrease, TCP carries out **fast retransmit** to quickly retransmit the segment that is suspected lost, without waiting for the RTO timer timeout. Notice that Figure 2-12 at time = 5×RTT shows `EffectiveWindow` = 1×`MSS`. Obviously, this is not in accordance with Eq. (2.3b), because currently `CongWin` equals 1×`MSS` and `FlightSize` equals 15×`MSS`. This simply means that the sender in fast retransmit ignores the `EffectiveWindow` size and simply retransmits the segment that is suspected lost. The times when three (or more dupACKs are received and fast retransmit is employed are highlighted with circle in Figure 2-12.

Only after receiving a regular, non-duplicate ACK (most likely the ACK for the fast retransmitted packet), the sender enters a new slow start cycle. After the 15th segment is retransmitted at time = 6×RTT, the receiver's acknowledgement requests the 23rd segment thus cumulatively acknowledging all the previous segments. The sender does not re-send #23 immediately because

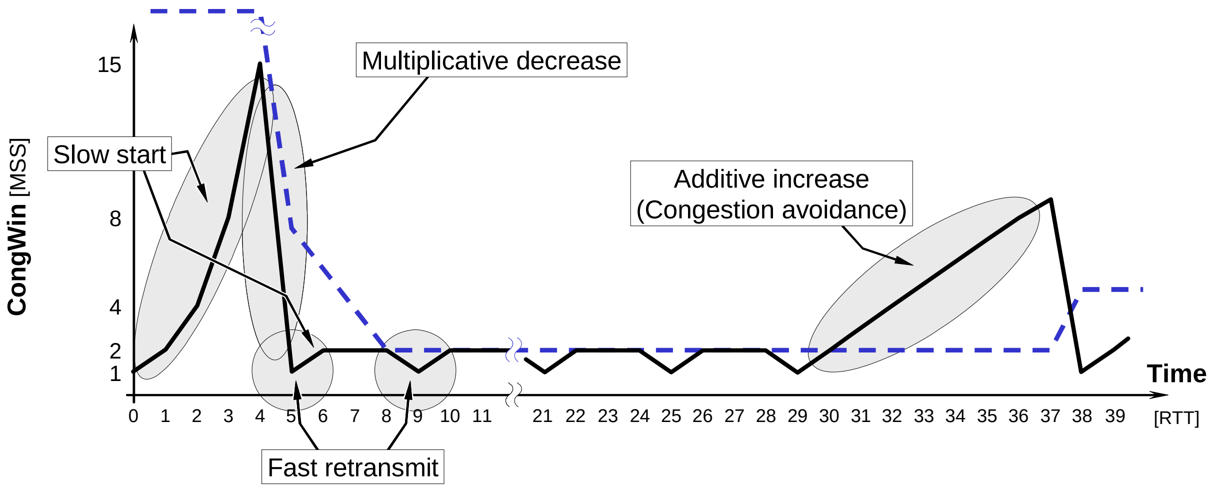

Figure 2-13: TCP Tahoe sender—highlighted are the key mechanisms for congestion avoidance and control; compare to Figure 2-12.

it still has no indication of loss. Although at time = 7×RTT the congestion window doubles to 2×`MSS` (because the sender is currently back in the *slow start* phase), there is so much data in flight that `EffectiveWindow` = 0 and the sender is shut down. Notice also that for repetitive slow starts, only ACKs for the segments sent after the loss was detected count. Cumulative ACKs for segments before the loss was detected do not count towards increasing `CongWin`. That is why, although at 6×RTT the acknowledgement for #23 cumulatively acknowledges packets 15–22, `CongWin` grows only by 1×`MSS` although the sender is in slow start (because there are too many outstanding segments).

However, even if `EffectiveWindow` = 0, TCP sender must send a 1-byte segment as indicated in Figure 2-10 and Figure 2-12. This usually happens when the receiver end of the connection advertises a window of `RcvWindow` = 0, and there is a *persist timer* (also called the zero-window-probe timer) associated with sending these segments. The tiny, 1-byte segment is treated by the receiver the same as any other segment. The sender keeps sending these tiny segments until the effective window becomes non-zero or a loss is detected.

In our example, three duplicate ACKs are received by time = 9×RTT at which point the 23rd segment is retransmitted. (Although `TimeoutInterval` = 3×RTT, we assume that ACKs are processed first, and the RTO timer is simply restarted for the just-retransmitted segment, without being declared as expired.) This continues until time = 29×RTT at which point the congestion window exceeds `SSThresh` and congestion avoidance takes off. The sender is in the **congestion avoidance** (also known as *additive increase*) phase when the current congestion window size is greater than the slow start threshold (`SSThresh`). During congestion avoidance, each time an ACK is received, the congestion window is increased as[16]:

$$\text{CongWin}(t) = \text{CongWin}(t-1) + \text{MSS} \times \frac{\text{MSS}}{\text{CongWin}(t-1)} \quad \text{[bytes]} \tag{2.5}$$

---

[16] The formula remains the same for cumulative acknowledgements, which acknowledge more than a single segment, but the reader should check further discussion in [Stevens 1994].

where `CongWin`(t–1) is the congestion window size before receiving the current ACK. The parameter t is *not* necessarily an integer multiple of round-trip time. Rather, t is just a time step that occurs whenever a new ACK is received and this can occur several times in a single RTT, i.e., a transmission round. It is important to notice that the resulting `CongWin` is *not* rounded down to the next integer value of `MSS` as in other equations. The congestion window can increase by at most one segment each round-trip time (regardless of how many ACKs are received in that RTT). This results in a linear increase.

Figure 2-13 summarizes the key congestion avoidance and control mechanisms. Notice that the second slow-start phase, starting at 5×RTT, is immediately aborted due to the excessive amount of unacknowledged data. Thereafter, the TCP sender enters a prolonged phase of dampened activity until all the lost segments are retransmitted through "fast retransmits."

It is interesting to notice that TCP Tahoe in this example needs 39×RTT in order to successfully transfer 71 segments (not counting 17 one-byte segments to keep the connection alive, which makes a total of 88 segments). Conversely, should the bottleneck bandwidth been known and constant, Go-back-7 ARQ would need 11×RTT to transfer 77 segments (assuming error-free transmission). In this example, bottleneck resource uncertainty and its dynamics introduce delay greater than three times the minimum possible one.

2.2.2 TCP Reno

Problems related to this section: Problem 2.8 → Problem 2.13

TCP Tahoe and Reno senders differ in their reaction to three duplicate ACKs. As seen earlier, Tahoe enters *slow start*; conversely, Reno enters *fast recovery*. This is illustrated in Figure 2-14, derived from Example 2.1.

After the fast retransmit algorithm sends what appears to be the missing segment, the **fast recovery** algorithm governs the transmission of new data until a *non-duplicate ACK* arrives. It is recommended [Stevens 1994; Stevens 1997; Allman *et al.* 1999] that `CongWindow` be incremented by one `MSS` for each additional duplicate ACK received over and above the first three dupACKs. This artificially inflates the congestion window in order to reflect the additional segment that has left the network. Because three dupACKs are received by the sender, this means that three segments have left the network and arrived successfully, but out-of-order, at the receiver. The fast recovery ends when either a retransmission timeout occurs or an ACK arrives that acknowledges all of the data up to and including the data that was outstanding when the fast recovery procedure began. After fast recovery is finished, the sender enters congestion avoidance.

As mentioned earlier in the discussion of Table 2-2, the reason for performing fast recovery rather than slow start is that the receipt of the dupACKs not only indicates that a segment has been lost, but also that segments are most likely leaving the network (although a massive segment duplication by the network can invalidate this conclusion). In other words, as the receiver can only generate a duplicate ACK when an error-free segment has arrived, that segment has left the network and is in the receive buffer, so we know it is no longer consuming network resources. Furthermore, because the ACK "clock" [Jac88] is preserved, the TCP sender can continue to transmit new segments (although transmission must continue using a reduced `CongWindow`).

TCP Reno sender retransmits the lost segment and sets congestion window to:

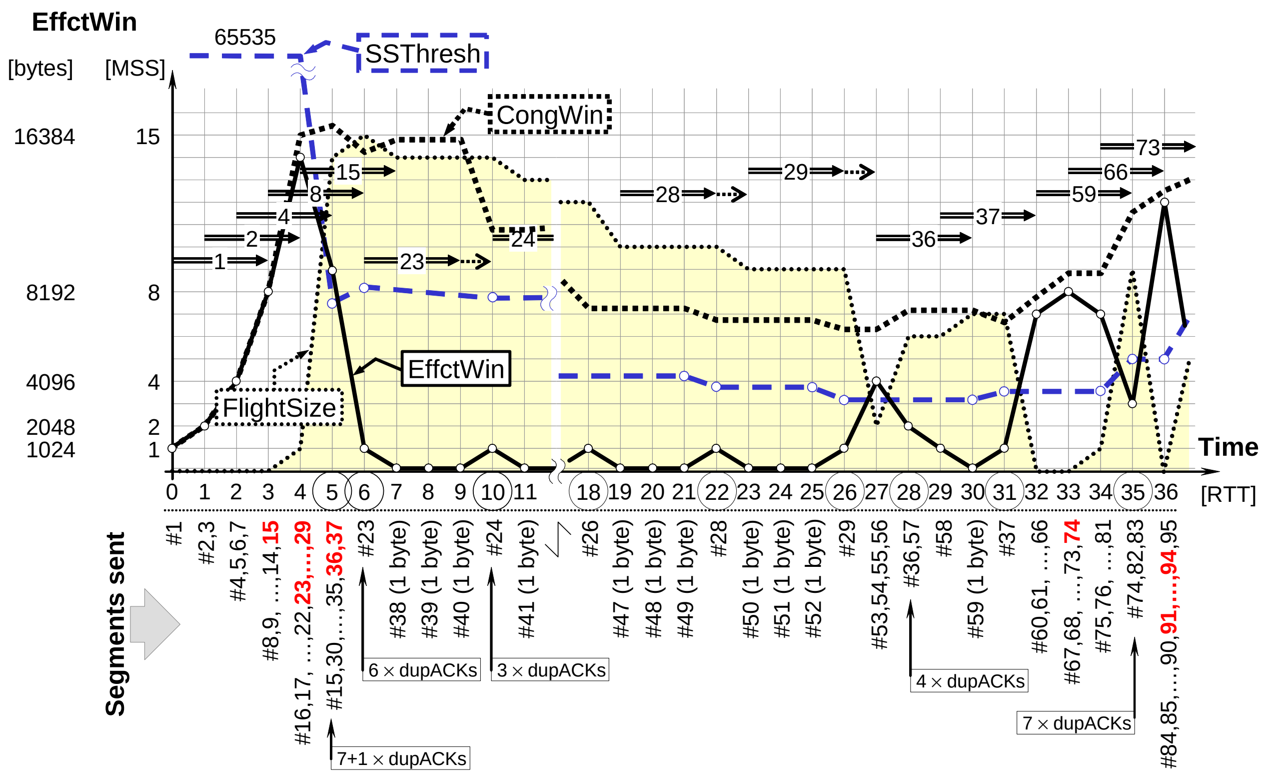

Figure 2-14: TCP Reno sender—the evolution of the effective and congestion window sizes for Example 2.1. The sizes are given on vertical axis (left) both in bytes and MSS units.

$$\texttt{CongWindow} = \max\ \{\lfloor ½\ \texttt{FlightSize} \rfloor, 2\times\texttt{MSS}\} + 3\times\texttt{MSS} \qquad (2.6)$$

where `FlightSize` is the amount of sent but unacknowledged data at the time of receiving the third dupACK. Compare this equation to (2.4), for computing `SSThresh`. This artificially "inflates" the congestion window by the number of segments (three) that have left the network and which the receiver has buffered. In addition, for each additional dupACK received *after the third dupACK*, increment `CongWindow` by `MSS`. This artificially inflates the congestion window in order to reflect the additional segment that has left the network (the TCP receiver has buffered it, waiting for the missing gap in the data to arrive).

As a result, in Figure 2-14 at 5×RTT when the sender receives 3+5 dupACKs, `CongWindow` becomes equal to $\frac{15}{2}+3+1+1+1+1+1 = 15.5 \times$ `MSS`. The last five 1's are due to 7+1−3 = 5 dupACKs received after the initial 3 ones. At 6×RTT the receiver requests the 23rd segment (thus cumulatively acknowledging up to the 22nd). `CongWindow` grows slightly to 17.75, but because there are 14 segments outstanding (#23 → #37), the effective window is shut up. The sender arrives at standstill and thereafter behaves similar to the TCP Tahoe sender (Figure 2-12).

Notice that, although at time = 10×RTT three dupACKs indicate that three segments that have left the network, these are only 1-byte segments, so it may be inappropriate to add 3×`MSS` as Eq. (2.6) postulates. RFC 2581 does not mention this possibility, so we continue applying Eq. (2.6) and because of this `CongWindow` converges to 6×`MSS` from above.

Figure 2-15 shows partial timeline at the time when the sender starts recovering. After receiving the 29th segment, the receiver delivers it to the application along with the buffered segments #30

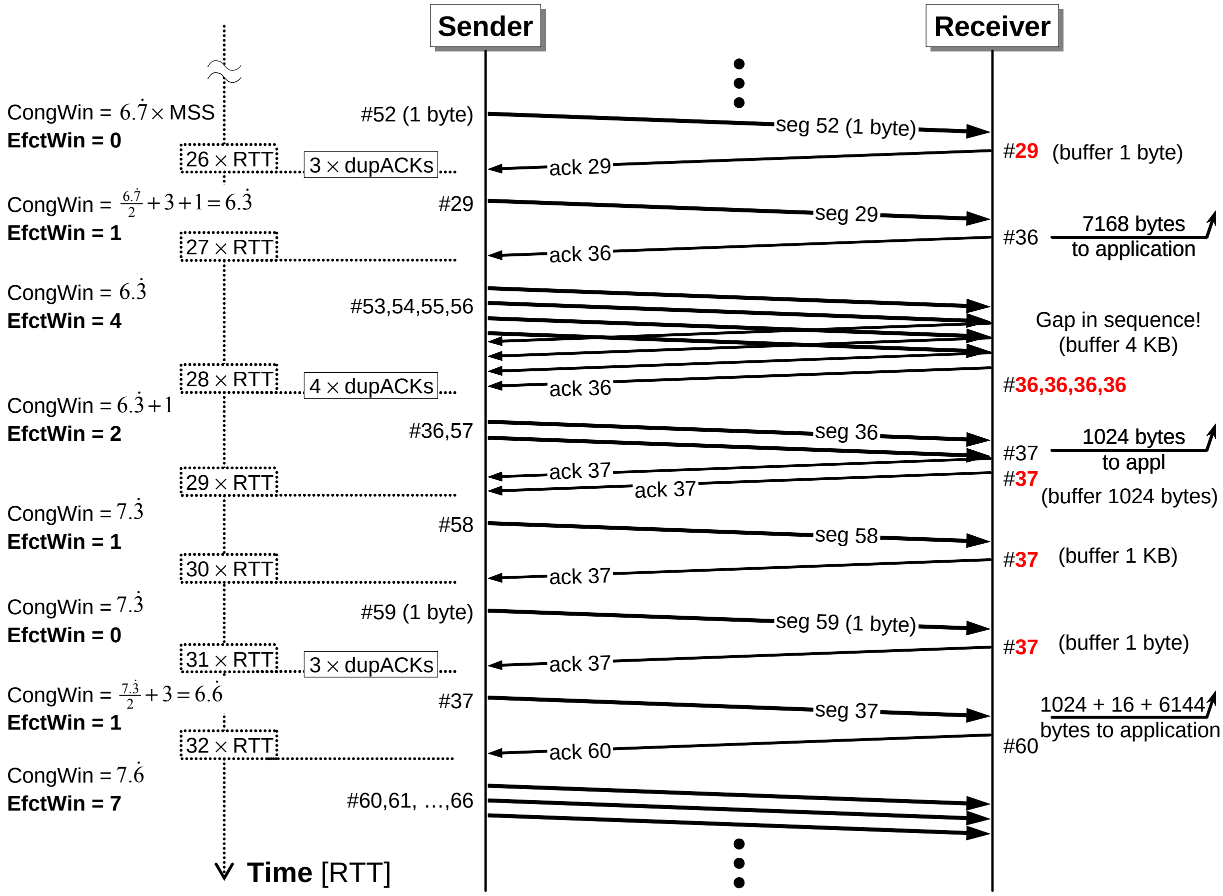

Figure 2-15: TCP Reno—partial timeline of segment and ACK exchanges for Example 2.1. (The slow start phase is the same as for Tahoe sender, Figure 2-10.)

→ #35 (a total of seven segments). At time = 27×RTT, a cumulative ACK arrives requesting the 36th segment (because segments #36 and #37 are lost at 5×RTT). Because `CongWindow` > 6×`MSS` and `FlightSize` = 2×`MSS`, the sender sends four new segments and each of the four makes the sender to send a dupACK. At 28×RTT, `CongWindow` becomes equal to $\frac{\dot{6}}{2}+3+1=\dot{7}$ × `MSS` and `FlightSize` = 6×`MSS` (we are assuming that the size of the unacknowledged 1-byte segments can be neglected).

Regarding the delay, TCP Reno in this example needs 37×RTT to successfully transfer 74 segments (not counting 16 one-byte segments to keep the connection alive, which makes a total of 90 segments—segment #91 and the consecutive ones are lost). This is somewhat better that TCP Tahoe and TCP Reno should better stabilize for a large number of segments.

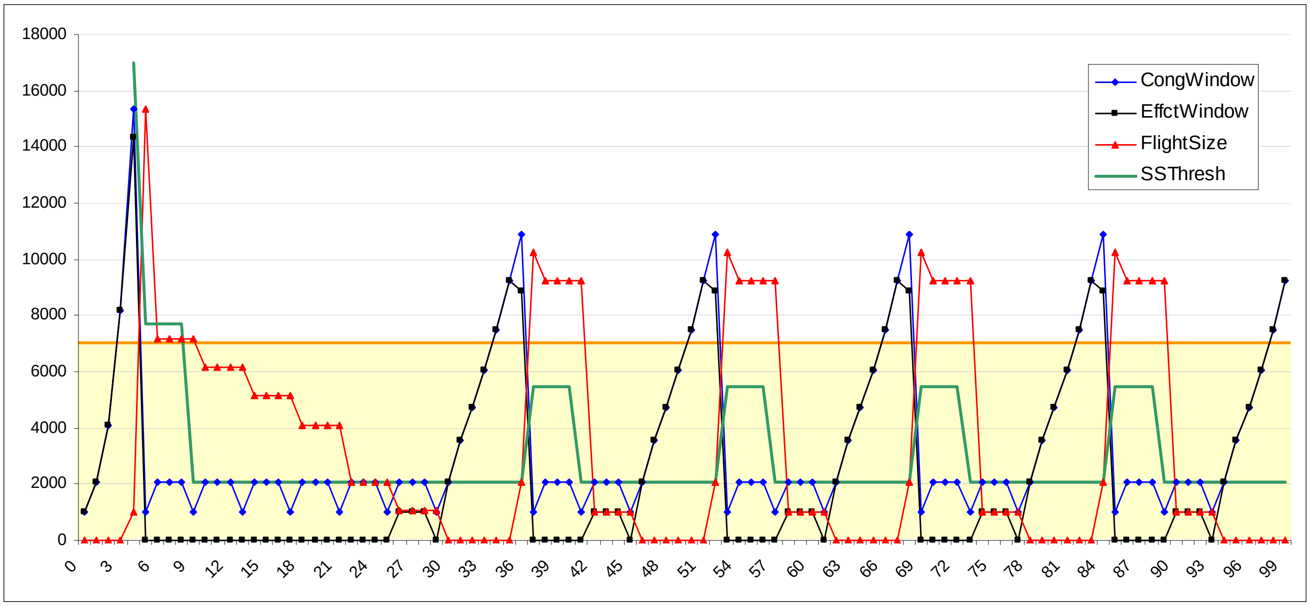

Figure 2-16: TCP Tahoe congestion parameters for Example 2.1 over the first 100 transmission rounds. The overall sender utilization comes to only 25 %. The lightly shaded background area shows the bottleneck router's capacity, which, of course, is constant.

2.2.3 TCP NewReno

Problems related to this section: Problem 2.15 → ??

The so-called **NewReno** version of TCP introduces a further improvement on fast recovery, which handles a case where two or more segments are lost within a single window. Same as the ordinary TCP Reno, the NewReno begins the **fast recovery** procedure when three duplicate ACKs are received, and ends it when either a retransmission timeout occurs or an ACK arrives that acknowledges all of the data up to and including the data that was outstanding when the fast recovery procedure began. After the presumably lost segment is retransmitted by fast retransmit, if the corresponding ACK arrives, there are two possibilities:

(4) The ACK specifies the sequence number at the end of the current window, in which case the retransmitted segment was the only segment lost from the current window. We call this acknowledgement a **full acknowledgment**.

(5) The ACK specifies the sequence number higher than the lost segment, but lower than the end of the window, in which case (at least) one more segment from the window has also been lost. We call this acknowledgement a **partial acknowledgment**.

As with the ordinary Reno, for each additional dupACK received while in fast recovery, NewReno increments `CongWindow` by `MSS` to reflect the additional segment that has left the network. The concept of partial acknowledgements is illustrated in Figure 2-17. In this scenario, the sender sends six segments, of which three are lost: segments #1, #3, and #5. The receiver buffers the three segments that arrive out of order and send three duplicate acknowledgements. Upon receiving the three dupACKs, the sender retransmits the oldest outstanding segment (#1)

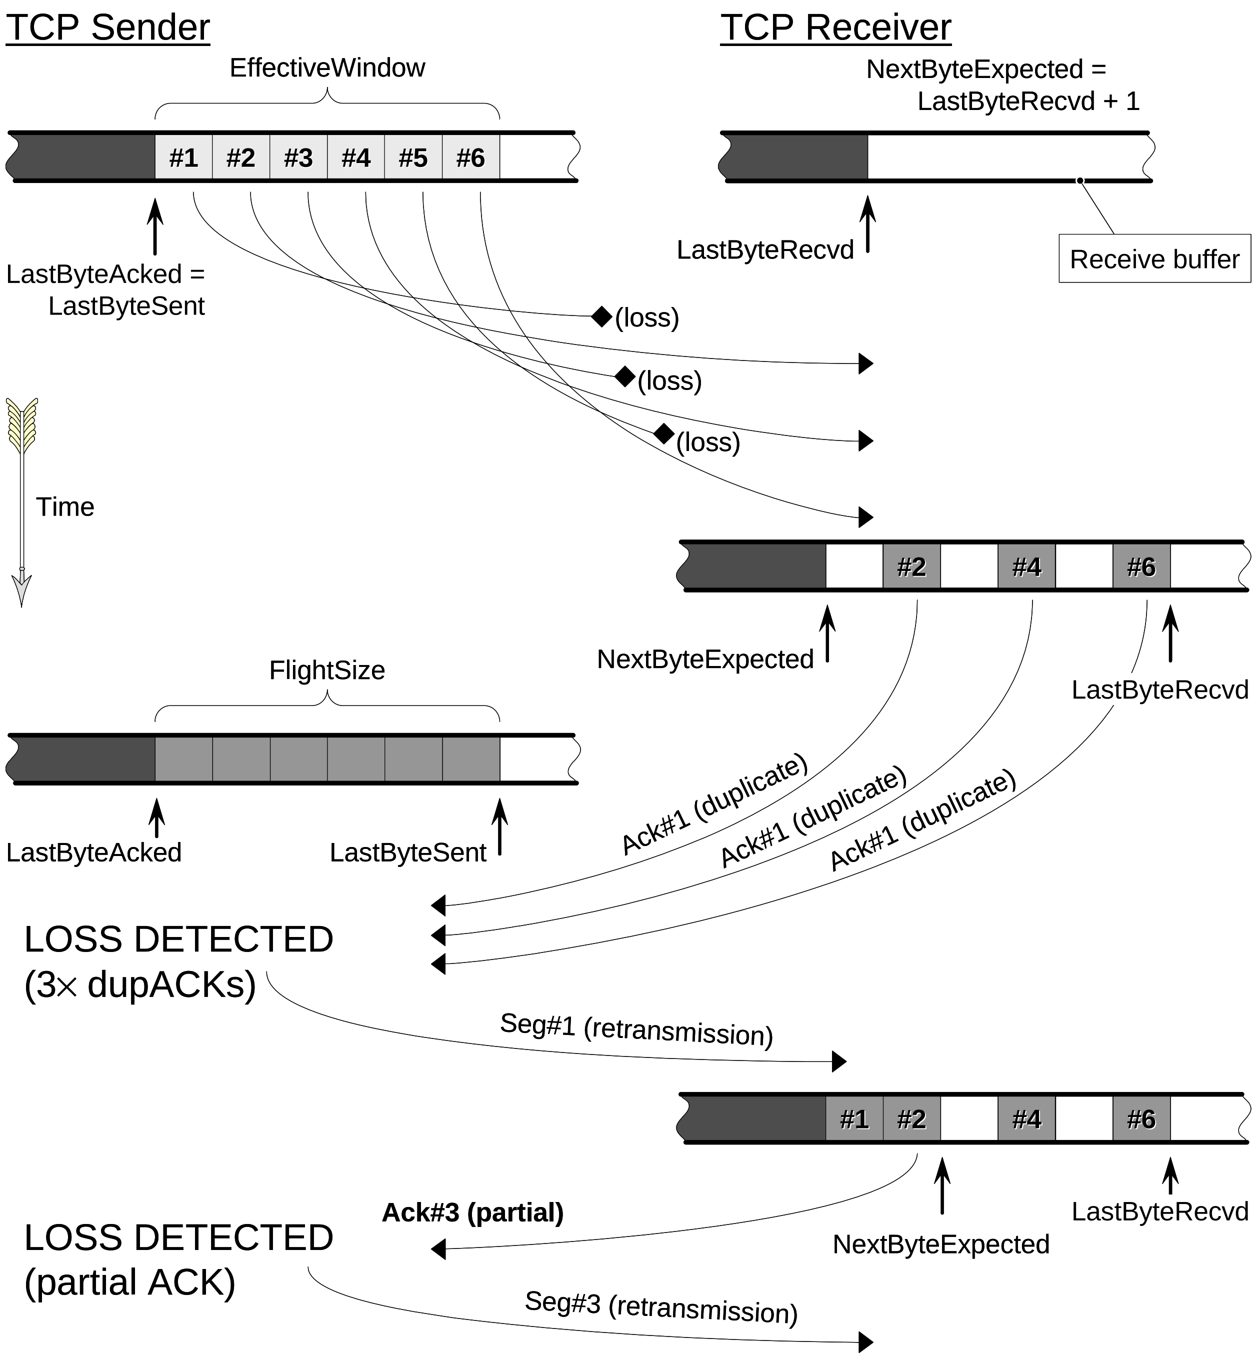

Figure 2-17: TCP NewReno partial acknowledgements.

and waits. The receiver fills only the first gap and now sends acknowledgement asking for segment #3. This is a *partial acknowledgement*, because it does not acknowledge *all* segments that were outstanding at the time the loss was detected.

The key idea of TCP NewReno is, if the TCP sender receives a *partial acknowledgment* during fast recovery, the sender should respond to the partial acknowledgment by inferring that the next in-sequence packet has been lost, and retransmitting that packet. In other words, NewReno proceeds to retransmit the second missing segment, without waiting for three dupACKs or RTO timer expiration. This means that TCP NewReno adds "partial acknowledgment" to the list of events in Table 2-2 by which the sender detects segment loss. The sender also deflates its congestion window by the amount of new data acknowledged by the cumulative acknowledgement, that is:

$$\text{NewlyAcked} = \texttt{LastByteAcked}(t) - \texttt{LastByteAcked}(t-1)$$

$$\texttt{CongWindow}(t)' = \texttt{CongWindow}(t-1) - \text{NewlyAcked} \tag{2.7a}$$

If (NewlyAcked ≥ `MSS`), then add back `MSS` bytes to the congestion window:

$$\texttt{CongWindow}(t) = \texttt{CongWindow}(t)' + \texttt{MSS} \tag{2.7b}$$

As with duplicate acknowledgement, this artificially inflates the congestion window in order to reflect the additional segment that has left the network. This "partial window deflation" attempts to ensure that, when fast recovery eventually ends, approximately `SSThresh` amount of data will be outstanding in the network. Finally, the sender sends a new segment if permitted by the new value of `EffectiveWin`.

When a *full acknowledgement* arrives, it acknowledges all the intermediate segments sent after the original transmission of the lost segment until the loss is discovered (the sender received the third duplicate ACK). This does not mean that there are no more outstanding data (i.e., `FlightSize` = 0), because the sender might have sent some new segments after it discovered the loss (if its `EffectiveWin` permitted transmission of news segments). At this point, the sender calculates its congestion window as:

$$\texttt{CongWindow} = \texttt{SSThresh} \tag{2.8}$$

Recall that `SSThresh` is computed using Eq. (2.4), where `FlightSize` is the amount of data outstanding when fast recovery was entered, *not* the current amount of data outstanding[17]. This reduction of the congestion window size is termed *deflating the window*. At this point, the TCP sender exits fast recovery and enters *congestion avoidance*.

An example of NewReno behavior is given below. Example 2.2 works over a similar network configuration as the one in Example 2.1. Again, we have a high-speed link from the TCP sender to the bottleneck router and a low-speed link from the router to the TCP receiver. However, there are some differences in the assumptions. In Example 2.1, we assumed a very large RTT, so that all packet transmission times are negligible compared to the RTT. We also assumed that cumulative ACKs acknowledged all the segments sent in individual RTT-rounds (or, bursts). Conversely, in Example 2.2, we will assume that the RTT is on the same order of magnitude as the transmission time on the *second* link. In addition, each segment is acknowledged individually. This results in a somewhat more complex, but also more accurate, analysis of the TCP behavior.

Example 2.2 Analysis of the Slow-Start Phase in TCP NewReno

Consider an application that is engaged in a lengthy file transfer using the TCP NewReno protocol over the network shown in the figure.

---

[17] RFC-3782 suggests an alternative option to set `CongWindow` = min{`SSThresh`, `FlightSize` + `MSS`}, where `FlightSize` is the *current* amount of data outstanding. Check RFC 3782 for details.

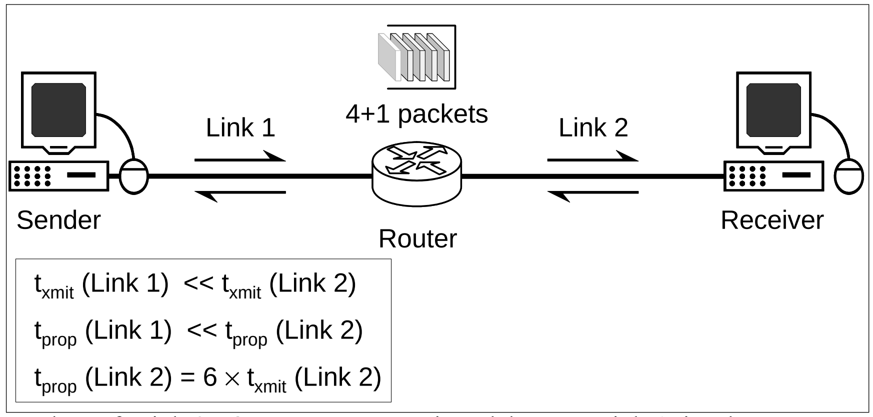

The following assumptions are made:

A1. Full duplex links connect the router to each endpoint host so that simultaneous transmissions are possible in both directions on each link. The transmission rate of Link-1 is much greater than that of Link-2. One-way propagation delay on Link-1 is also negligible compared to the propagation delay of Link-2. Assume that all packet transmissions are error free.

A2. The propagation delay on Link-2 (from the router to the receiver equals six times the transmission delay for data packets on the same link. Also assume that the ACK packet size is negligible, i.e., their transmission delay is approximately zero.

A3. The router buffer can hold up to four packets plus one packet currently in transmission. The packets that arrive to a full buffer are dropped. However, this does *not* apply to ACK packets, i.e., ACKs do not experience congestion or loss.

A4. The receiver does not use delayed ACKs, i.e., it sends an ACK immediately after receiving a data segment.

A5. The receiver has set aside a large receive buffer for the received segments, so this will never be a limiting factor for the sender's window size.

Considering only the slow-start phase (until the first segment loss is detected), we would like to know the following:

(a) The evolution of the parameters such as congestion window size, router buffer occupancy, and how well the communication pipe is filled with packets.

(b) The ordinal packet numbers of all the packets that will be dropped at the router (due to the lack of router memory space).

(c) The maximum congestion widow size that will be achieved after the first packet loss is detected (but before the time t = 60).

(d) How many packets will be sent after the first packet loss is detected until the time t = 60? Explain the reason for transmission of each of these packets.

The solution is shown in Figure 2-18 and discussed in the following text.

Figure 2-18 shows the evolution of four parameters over the first 20 time units of the slow-start phase. The four parameters are: (*i*) congestion window size; (*ii*) slow start threshold; (*iii*) current number of packets in the router, both in transmission or waiting for transmission; and (*iv*) current number of packets in flight on Link-2, that is the packets that neither are in the router nor acknowledged. Notice that the last parameter is *not* the same as the `FlightSize` defined at the beginning of Section 2.2. `FlightSize` is maintained by the sender to know how many segments are outstanding. Unlike this, the current number of packets in flight on Link-2 (bottom chart in Figure 2-18) represents only the packets that were transmitted by the router, but for which the ACK has not yet arrived at the TCP sender.

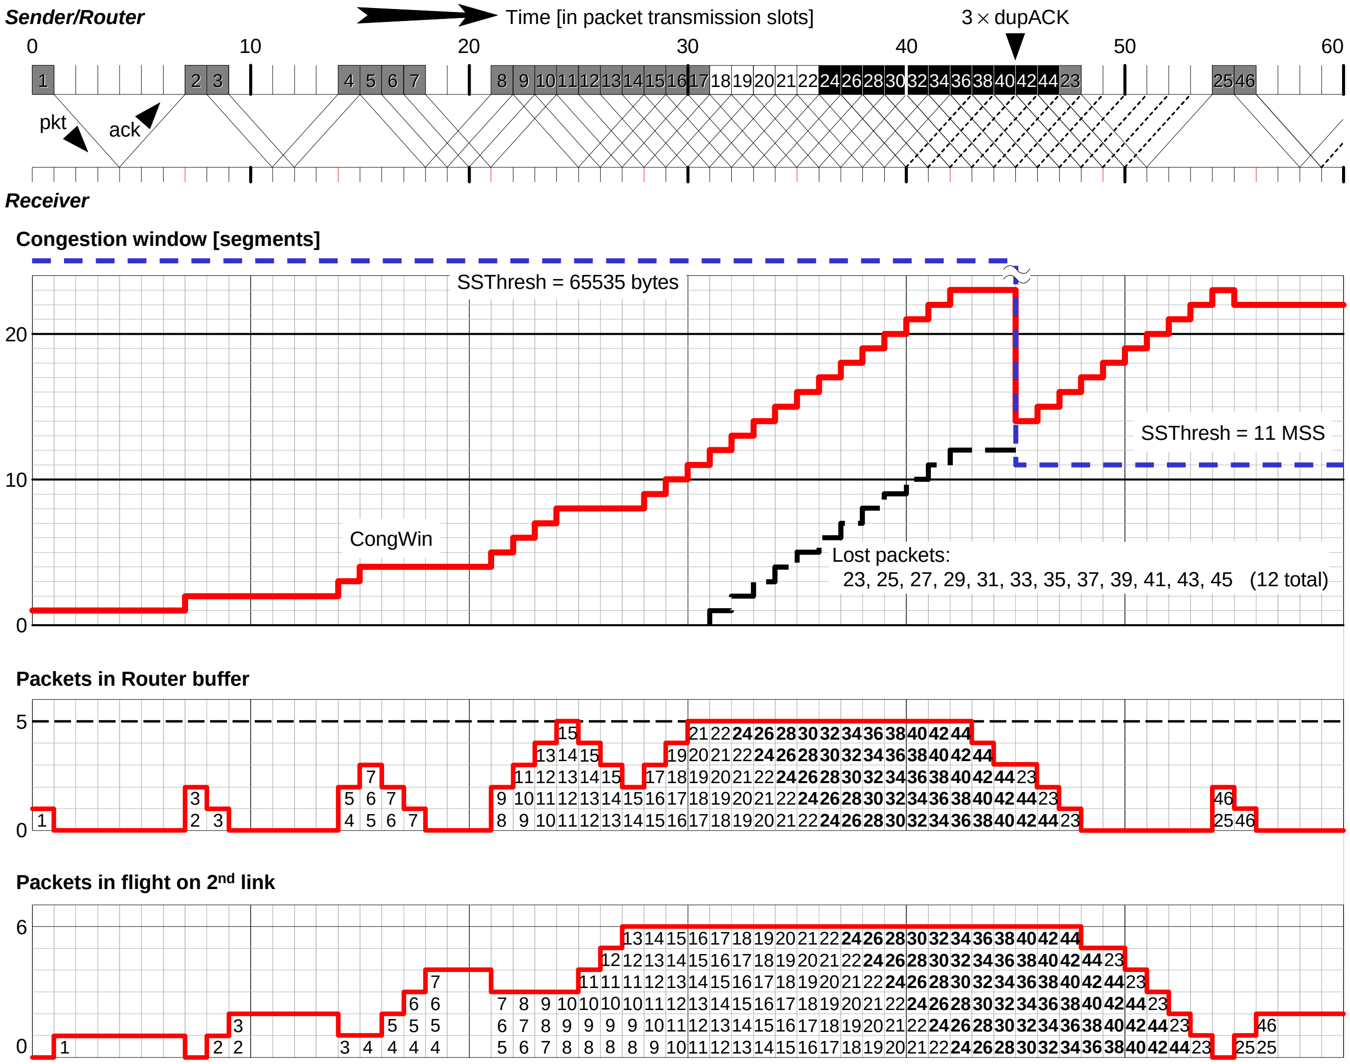

Figure 2-18: Evolution of the key parameters for the TCP NewReno sender in Example 2.2. (*Continued in* Figure 2-19.)

The gray boxes on the top line symbolize packet transmissions on the Link-2. This is because the transmission delay on Link-1 is negligible, so any packet sent by the TCP sender immediately ends up on the router.

Recall that during the slow start, the sender increments its congestion window size by one MSS for each successfully received acknowledgement. We can see on the top of Figure 2-18 how the acknowledgment for packet #1 arrives at t = 7 and the congestion window size rises to 2. The sender sends two more segments (#2 and #3) and they immediately end up on the router. Notice how packets that are sent back-to-back (in bursts) are separated by the packet transmission time on Link-2. When the ACK arrives for #2, the sender is in slow start, so `FlightSize` = 2 − 1 = 1 and `CongWin` = 2 + 1 = 3. According to equation (2.3),

$$\texttt{EffectiveWin} = \texttt{CongWin} - \texttt{FlightSize} = 3 - 1 = 2$$

In other words, during slow start, every time the sender receives a non-duplicate acknowledgement for one segment, the sender can send two new segments. After sending segments #4 and #5, we have: `CongWin` = 3, `FlightSize` = 3, and `EffectiveWin` = 0. When the ACK for segment #3 arrives (delayed by t_x(Link-2) after the first ACK, because the segment #3 traveled back-to-back behind #2) we have: `FlightSize` = 3 – 1 = 2 and `CongWin` = 3 + 1 = 4. Therefore,

`EffectiveWin` = `CongWin` – `FlightSize` = 4 – 2 = 2

and the sender will send two new segments (#6 and #7). Now, we have `CongWin` = 4, `FlightSize` = 4, and `EffectiveWin` = 0. The sender is waiting for an acknowledgement for segment #4.

Notice that when an in the chart for the current number of packets in the router shows below the curve the ordinal numbers of the packets. The bottommost packet is the one that is in transmission during the current time slot. The packets above it are currently waiting for transmission. For example, at time t = 7, packet #2 is currently in transmission and packet #3 is waiting in the router memory. At t = 8, packet #2 is traversing Link-2 (shown under the bottom curve) and packet #3 is in transmission on the router. The attentive reader might notice that the packet numbers in the bottommost row of the router buffer curve are identical to the ones at the top of Figure 2-18.

(a)

Because we are considering a single connection in slow start, packet arrivals on the router occur in bursts of exactly two packets. This is because for every received acknowledgment, `FlightSize` is reduced by 1 and `CongWin` is incremented by 1, which means that effectively the sender can send two new packets. The sender sends two new packets and they immediately end up on the router. Therefore, when a buffer overflow occurs, exactly one packet is dropped. This is because at the beginning of the preceding time slot, the buffer would have been full and the router would transmit one packet, therefore freeing up space for one new packet. When two packets arrive, the first is stored and the second is dropped.

The first loss happens at time t = 31. In the previous time slot (t =30) the router had five packets (#17, #18, #19, #20, and #21), of which packet #17 was transmitted by the router. At t = 31, the acknowledgement for packet #11 will arrive and it will increment congestion window by one, to the new value of 12×MSS. The sender sends packets #22 and #23 and they immediately arrive to the router. The router just transmitted packet #17 and has space for only one new packet. Packet #22 joins the tail of the waiting line and packet #23 is dropped.

The top row of Figure 2-18 shows white boxes for the transmission periods of the five packets that the router will transmit after the loss of packet #23. Black boxes symbolize packets that are sent out of order, for which the preceding packet was dropped at the router (due to the lack of router memory space). There will be a total of 12 packets fro which the preceding packet was lost, starting with packet #24 and ending with packet #44.

The TCP sender receives three duplicate acknowledgements (asking for packet #23) at time t = 45. The reader should notice that packet #23 was lost at the router at time t = 31, but the sender learned about the loss only at time t = 45 (by receiving three dupACKs)! When the sender discovers the loss, it sets the congestion window size to one half of the number of segments in flight, which is 23, plus 3 for three duplicate acknowledgements—remember equation (2.6) from Section 2.2.2. The slow-start threshold is set to one-half of the number of segments in flight—

remember equation (2.4)—so `SSThresh` becomes equal to 11. In addition, because this is a TCP Reno sender, the congestion window is incremented by one for each new duplicate acknowledgement that is received after the first three.

Upon detecting loss at time $t = 45$, the TCP sender immediately retransmits segment #23, but the packet joins the queue at the router behind packets #42 and #44, which arrived before #23. The router is not aware that these are TCP packets, lest that some of them are retransmitted, so it does not give preferential treatment to retransmitted packets. As seen in the top row of Figure 2-18, the router will transmit packet #23 over Link-2 during the time slot $t = 47$. Therefore, generally it takes longer than one RTT for the sender to receive the acknowledgment for a retransmitted segment.

The acknowledgement for the retransmitted segment #23 arrives at time $t = 54$ and it asks for segment #25 (because #24 was received correctly earlier). This is only a *partial acknowledgment* because a full acknowledgement would acknowledge segment #45. Therefore, the TCP NewReno sender immediately sends packet #25 without waiting for three duplicate acknowledgements. In addition, the sender adjusts its congestion window size according to Eq. (2.7):

$$\text{NewlyAcked} = \text{segment\#24} - \text{segment\#22} = 2\ \text{MSS}$$

Because (NewlyAcked ≥ `MSS`), the sender uses Eq. (2.7b):

$$\texttt{CongWindow}(54) = \texttt{CongWindow}(53) - \text{NewlyAcked} + \texttt{MSS} = 23 - 2 + 1 = 22 \times \texttt{MSS}$$

Because the current `FlightSize` = 21 × `MSS` (segments #25 through #45 are unacknowledged), one new segment can be sent. As a result, the sender transmits segment #46.

(b)

There will be a total of 12 packets lost during the considered period. The lost packets are (also indicated in Figure 2-18): 23, 25, 27, 29, 31, 33, 35, 37, 39, 41, 43, and 45.

TCP NewReno Fast Recovery Phase

Continuing with Example 2.2, the TCP NewReno sender will enter the fast recovery phase when it discovers the loss of packet #23 at time $t = 45$ (by receiving three dupACKs). The sender will exit the fast recovery phase when it receives a full acknowledgement. The sender originally transmits packet #23 at time $t = 31$ and it immediately arrives at the router where it is lost (Figure 2-18). From time $t = 31$ until the loss is discovered at time $t = 45$, the sender sends a total of 22 "intermediate segments" (segments #24, #25, …, #45). Therefore, the TCP sender will consider it a "full acknowledgement" when it receives an acknowledgement for packet #45. At this point, the sender will *exit fast recovery and enter congestion avoidance*. Notice that segment #46, and any segments transmitted thereafter might still be outstanding.

Figure 2-19 shows the continuation of Figure 2-18 for the same Example 2.2. As seen, the sender has not yet received a "full acknowledgement" until time $t = 120$ (packet #45 has not been retransmitted); therefore, the sender is still in the fast recovery state.

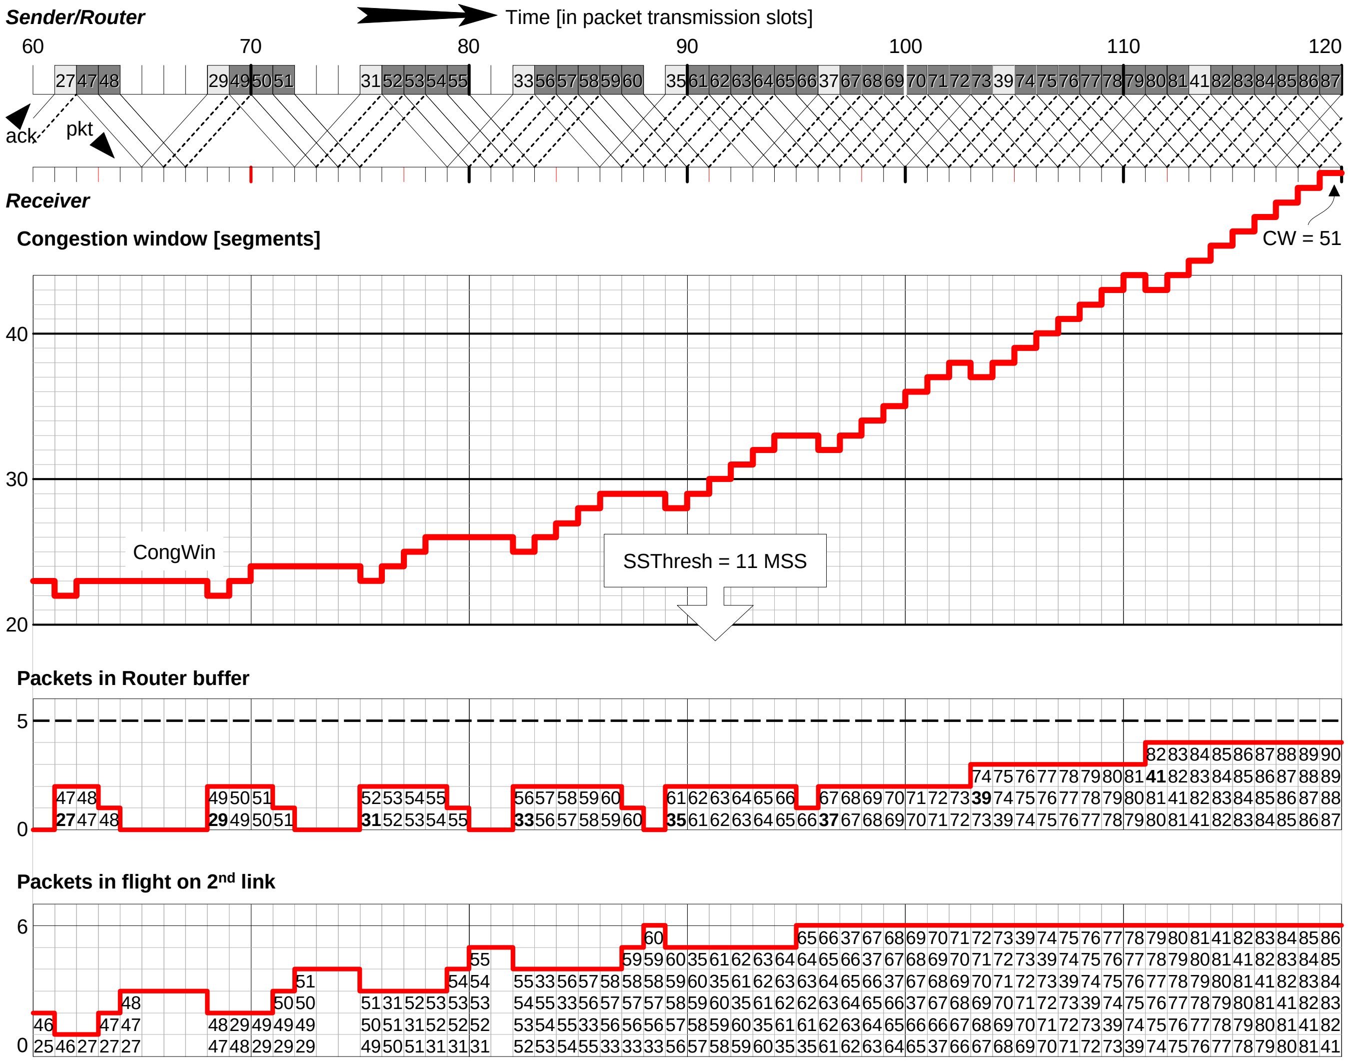

Figure 2-19: Evolution of the key parameters for the TCP NewReno sender in Example 2.2. (*Continued from* Figure 2-18.)

2.3 Fairness

2.4 Recent TCP Versions

Early TCP versions, Tahoe and Reno, perform relatively simple system observation and control. The TCP Tahoe performance is illustrated in Figure 2-16 over the first 100 transmission rounds. Although the obvious inefficiency (sender utilization is only 25 %) can be somewhat attributed to the contrived scenario of Example 2.1, this is not far from reality. By comparison, a simple Stop-and-Wait protocol would achieve the sender utilization of ?? %. Recent TCP versions introduce sophisticated observation and control mechanisms to improve performance.

TCP Vegas [Brakmo & Peterson 1995] watches for the signs of incipient congestion—before losses occur—and takes actions to avert it.

TCP Westwood [Mascolo *et al*. 2001] uses bandwidth estimates to compute the congestion window and slow start threshold after a congestion episode.

FAST TCP [Jin *et al*. 2003] detects congestion by measuring packet delays.

2.5 TCP over Wireless Links

The TCP congestion control algorithms presented in Section 2.2 assume most packet losses are caused by routers dropping packets due to traffic congestion. However, packets may be also dropped if they are corrupted in their path to destination. In wired networks the fraction of packet loss due to transmission errors is generally low (less than 1 percent). Communication over wireless links is often characterized by sporadic high bit-error rates, and intermittent connectivity due to handoffs. TCP performance in such networks suffers from significant throughput degradation and very high interactive delays

Several factors affect TCP performance in mobile ad-hoc networks (MANETs):

- Wireless transmission errors
- Power saving operation
- Multi-hop routes on shared wireless medium (for instance, adjacent hops typically cannot transmit simultaneously)
- Route failures due to mobility

Figure 2-20

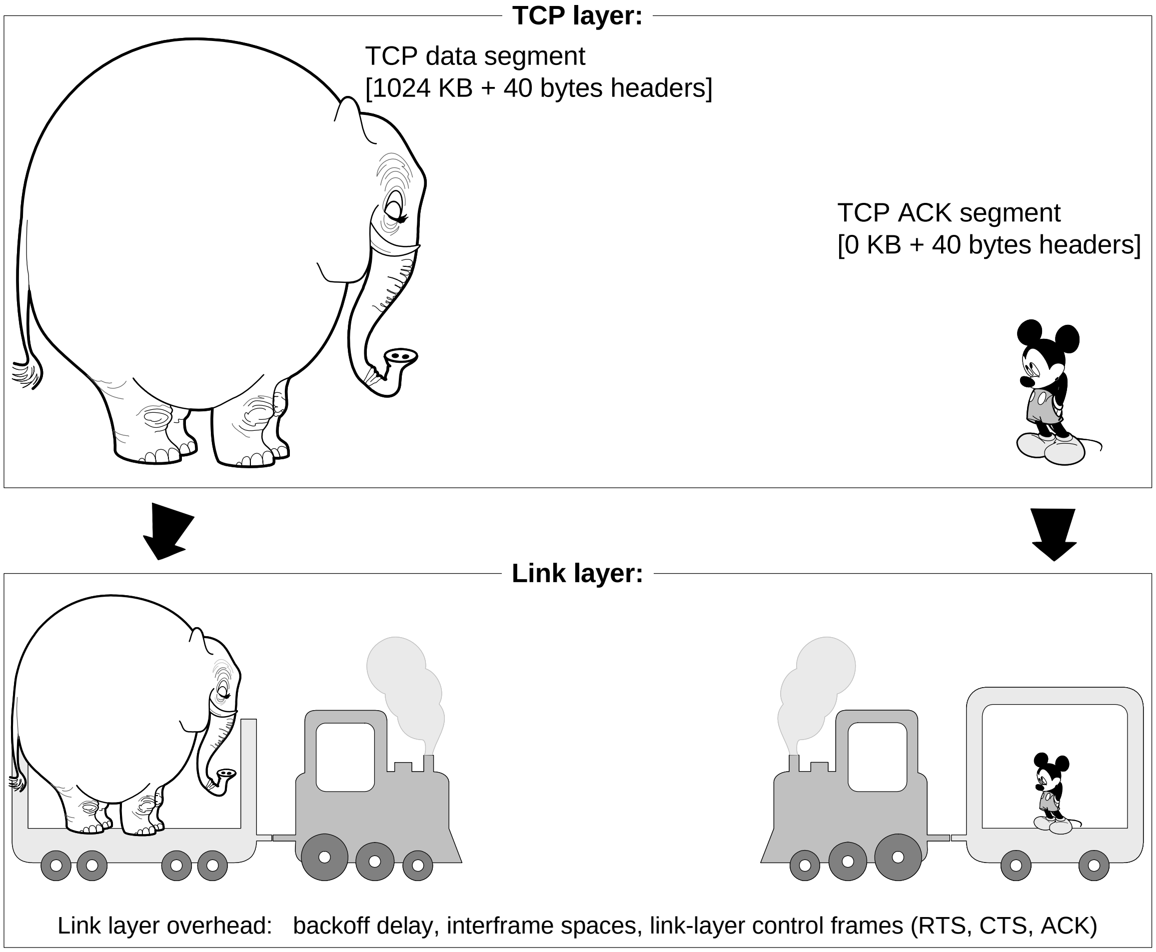

Figure 2-20: Due to significant wireless link-layer overhead, TCP data segments and TCP acknowledgements (which are of greatly differing sizes) appear about the same size at the link layer.

2.6 Summary and Bibliographical Notes

The TCP service model provides a communication abstraction that is reliable, ordered, point-to-point, duplex, byte-stream, and flow and congestion controlled. TCP's notion of "duplex" is that the same logical connection handles reliable data delivery in both directions. Unlike ARQ protocols described in Section 1.3, which treat data packets as atomic units, TCP treats bytes as the fundamental unit of reliability.

TCP sender uses the received cumulative acknowledgments to determine which packets have reached the receiver, and provides reliability by retransmitting lost packets. The sender detects the loss of a packet either by the arrival of several duplicate acknowledgments or the expiration of the timeout timer due to the absence of an acknowledgment for the packet. To accurately set the timeout interval, the sender maintains a running average of the estimated roundtrip delay and the mean linear deviation from it. The timeout interval is calculated as the sum of the smoothed round-trip delay plus four times its mean deviation. TCP reacts to packet losses by decreasing its transmission (congestion) window size before retransmitting packets, initiating congestion control or avoidance mechanisms (e.g., slow start), and backing off its retransmission timer (Karn's algorithm). These actions result in a reduction in the load on the intermediate links, thereby controlling the congestion in the network.

[Stevens, 1994] provides the most comprehensive coverage of TCP in a single book. It appears that this whole book is available online at http://www.uniar.ukrnet.net/books/tcp-ip_illustrated/.

[Comer, 2006] is also very good, although does not go in as much detail.

TCP described in 1974 by Vinton Cerf and Robert Kahn in IEEE Transactions on Communication. Three-way handshake described by Raymond Tomlinson in SIGCOMM 1975.

TCP & IP initial standard in 1982 (RFC-793 & RFC-791). BSD Unix 4.2 released in 1983 supported TCP/IP. Van Jacobson's algorithms for congestion avoidance and congestion control published in 1988; most implemented in 4.3 BSD Tahoe.

In the original TCP specification (RFC-793), the retransmission timeout (RTO) was set as a multiple of a running average of the RTT. For example, it might have been set as $\texttt{TimeoutInterval} = \eta \times \texttt{EstimatedRTT}$, with η set to a constant such as 2. However, this simple choice failed to take into account that at high loads round trip variability becomes high, leading to unnecessary retransmissions. The solution offered by Jacobson, see Eq. (2.2), factors in both average and standard deviation.

The TCP Reno Fast Recovery algorithm was described in RFC 2581 and first implemented in the 1990 BSD Reno release.

In 1996, Janey Hoe [Hoe, 1996] proposed an enhancement to TCP Reno, which subsequently became known as *NewReno*. The main idea here is for the TCP sender to remain in fast recovery until all the losses in a window are recovered.

There have been other enhancements proposed to TCP over the past few years, such as TCP Vegas congestion control method [Brakmo & Peterson, 1995], various optimizations for wireless networks, optimizations for small windows (e.g., RFC-3042), etc.

RFC-3168, 3155, 3042, 2884, 2883, 2861, 2757, 2582 (NewReno)

The NewReno modification of TCP's Fast Recovery algorithm is described in RFC-3782 [Floyd et al. 2004].

TCP over wireless:

[Balakrishnan, *et al*., 1997], [Holland & Vaidya, 1999], [Fu, *et al*., 2005].

TCP has supported ongoing research since it was written. As a result, the End-to-End research group has published a Roadmap for TCP Specification Documents [RFC-4614] which will guide expectations in that area.

SSFnet.org, "TCP Regression Tests," Online at: http://www.ssfnet.org/Exchange/tcp/tcpTestPage.html

The SSFnet.org tests show the behavior of SSF TCP Tahoe and Reno variants for different networks, TCP parameter settings, and loss conditions.

NASA Jet Propulsion Laboratory (JPL) and Vinton Cerf recently jointly developed Disruption-Tolerant Networking (DTN) protocol to transmit images to and from a spacecraft more than 20 million miles from Earth. DTN is intended for reliable data transmissions over a deep space communications network, for which the TCP/IP protocol suite is unsuitable. An interplanetary Internet needs to be strong enough to withstand delays, disruptions, and lost connections that space can cause. For example, errors can happen when a spacecraft slips behind a planet, or when solar storms or long communication delays occur. Even traveling at the speed of light, communications sent between Mars and Earth take between three-and-a-half minutes to 20 minutes. Unlike TCP, DTN does not assume there will be a constant end-to-end connection. DTN is designed so that if a destination path cannot be found, the data packets are not discarded but are kept in a network node until it can safely communicate with another node. The interplanetary Internet could allow for new types of complex space missions that involve multiple landed, mobile, and orbiting spacecraft, as well as ensure reliable communications for astronauts on the surface of the moon.

http://www.jpl.nasa.gov/news/news.cfm?release=2008-216

Problems

Problem 2.1

Consider the TCP procedure for estimating RTT with $\alpha = 0.125$ and $\beta = 0.25$. Assume that the `TimeoutInterval` is initially set as 3 seconds. Suppose that **all** measured RTT values equal 5 seconds, no segment loss, and the segment transmission time is negligible. The sender starts sending at time zero.

(a) What values will `TimeoutInterval` be set to for the segments sent during the first 11 seconds?
(b) Assuming a TCP Tahoe sender, how many segments will the sender transmit (including retransmissions) during the first 11 seconds?
(c) Repeat steps (a) and (b) but this time around assume that the sender picked the initial `TimeoutInterval` as 5 seconds?

Show the work.

Problem 2.2

Consider two hosts connected by a local area network with a negligible round-trip time. Assume that one is sending to the other a large amount of data using TCP with RcvBuffer = 20 Kbytes and MSS = 1 Kbytes. Also assume error-free transmission, high-speed processors in the hosts, and reasonable values for any other parameters that you might need.

(a) Draw the congestion window diagram during the slow-start (until the sender enters congestion avoidance) for the network speed of 100 Mbps.

(b) How different the diagram becomes if the network speed is reduced to 10 Mbps? 1 Mbps?

(c) What will be the average throughput (amount of data transmitted per unit of time) once the sender enters congestion avoidance?

Explain your answers.

Problem 2.3

Suppose that the hosts from Problem 2.2 are connected over a satellite link with RTT = 20 ms (low earth orbit satellites are typically 850 km above the Earth surface). Draw the congestion window diagram during the slow-start for the network speed of 100 Mbps. Explain any similarities or differences compared to the one from Problem 2.2(a).

Problem 2.4

Consider the network shown in the figure. TCP senders at hosts *A* and *B* have 3.6 KB of data each to send to their corresponding TCP receivers, both running at host *C*. Assume MTU = 512 bytes for all the links and `TimeoutInterval` = 2×RTT = 2×1 sec. The router buffer size is 3 packets in addition to the packet currently being transmitted; should the router need to drop a packet, it drops the last arrived from the host which currently sent more packets. Sender *A* runs TCP Tahoe and sender *B* runs TCP Reno and assume that sender *B* starts transmission 2×RTTs after sender *A*.

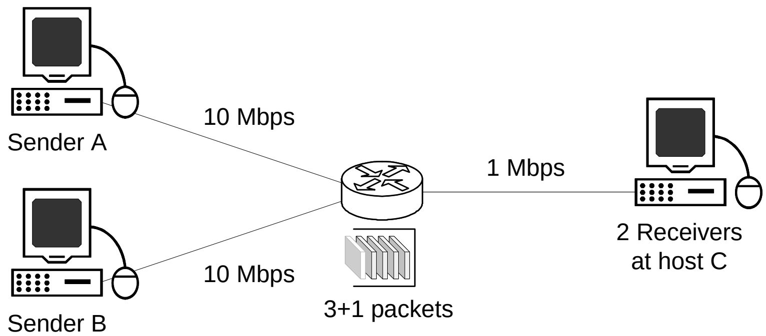

(a) Trace the evolution of the congestion window sizes on both senders until all segments are successfully transmitted.

(b) What would change if `TimeoutInterval` is modified to 3×RTT = 3×1 sec?

Assume a large `RcvWindow` and error-free transmission on all the links. Finally, to simplify the graphs, assume that all ACK arrivals occur exactly at unit increments of RTT and that the associated `CongWindow` update occurs exactly at that time, too.

Problem 2.5

Consider a TCP Tahoe sender working on the network with RTT = 1 sec, `MSS` = 1 KB, and the bottleneck link bandwidth equal to 128 Kbps. Ignore the initial slow-start phase and assume that the sender exhibits periodic behavior where a segment loss is always detected in the congestion avoidance phase via duplicate ACKs when the congestion window size reaches `CongWindow` = 16×`MSS`.

(a) What is the min/max range in which the window size oscillates?

(b) What will be the average rate at which this sender sends data?

(c) Determine the utilization of the bottleneck link if it only carries this single sender.

[Hint: When computing the average rate, draw the evolution of the congestion window. Assume `RcvWindow` large enough not to matter.]

Problem 2.6

Specify precisely a system that exhibits the same behavior as in Problem 2.5:

- What is the buffer size at the bottleneck router?
- What is the minimum value of `TimeoutInterval`?

Demonstrate the correctness of your answer by graphing the last two transmission rounds before the segment loss is detected and five transmission rounds following the loss detection.

Problem 2.7

Consider two hosts communicating using the TCP-Tahoe protocol. Assume RTT = 1, `MSS` = 512 bytes, `TimeoutInterval` = 3×RTT, `SSThresh` = 3×`MSS` to start with, and `RcvBuffer` = 2 KB. Also, assume that the bottleneck router has available buffer size of 1 packet in addition to the packet currently being transmitted.

(a) Starting with `CongWindow` = 1×`MSS`, determine the congestion window size when the first packet loss will happen at the router (not yet detected at the sender).

(b) What will be the amount of unacknowledged data at the sender at the time the sender detects the loss? What is the total number of segments acknowledged by that time?

Assume that no cumulative ACKs are sent, i.e., each segment is acknowledged individually.

Problem 2.8

Consider two hosts communicating by TCP-Reno protocol. Assume RTT = 1, `MSS` = 256 bytes, `TimeoutInterval` = 3×RTT, `RcvBuffer` = 2 KB, and the sender has a very large file to send. Start considering the system at the moment when it is in slow start state, `CongWin` = 8×`MSS`, `SSThresh` = 10×`MSS` and the sender just sent eight segments, each 1×`MSS` bytes long. Assume that there were no lost segments before this transmission round and currently there are no buffered segments at the receiver.

Assuming that, of the eight segments just sent, the fourth segment is lost, trace the evolution of the congestion window sizes for the subsequent five transmission rounds. Assume that no more segments are lost for all the considered rounds. For every step, indicate the transmitted segments and write down the numeric value of `CongWin` (in bytes). To simplify the charts, assume that ACK arrivals occur exactly at unit increments of RTT and that the associated `CongWin` update occurs exactly at that time, too.

Problem 2.9

Consider the network configuration shown in the figure below. The mobile node connects to the server using the TCP protocol to download a large file. Assume MSS = 1024 bytes, error-free transmission, and sufficiently large storage spaces at the access point and the receiver. Assume that the Assuming that the TCP receiver sends only cumulative acknowledgements. Calculate how long time it takes to deliver the first 15 Kbytes of data from that moment the TCP connection is established. In addition, draw the timing diagram of data and acknowledgement transmissions. (You can exploit the fact that TCP sends cumulative acknowledgements.)

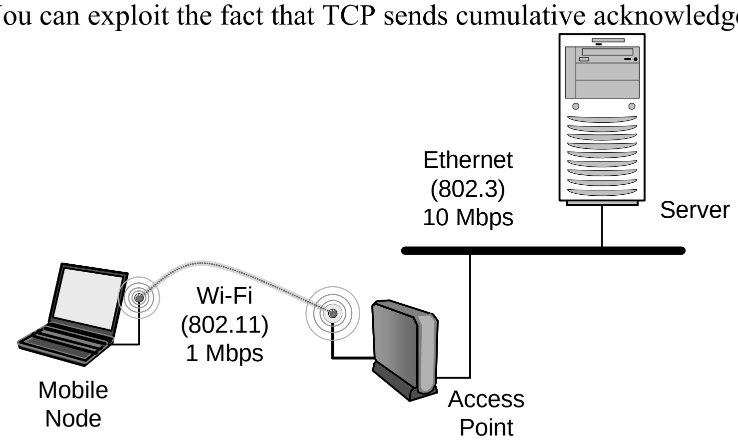

(In case you need these, assume the distance between the mobile node and the access point equal to 100 m, and the same from the access point to the server. Also, the speed of light in the air is 3×10^8 m/s, and in a copper wire is 2×10^8 m/s.)

Problem 2.10

Consider an application that is engaged in a lengthy file transfer using the TCP Tahoe protocol over the following network.

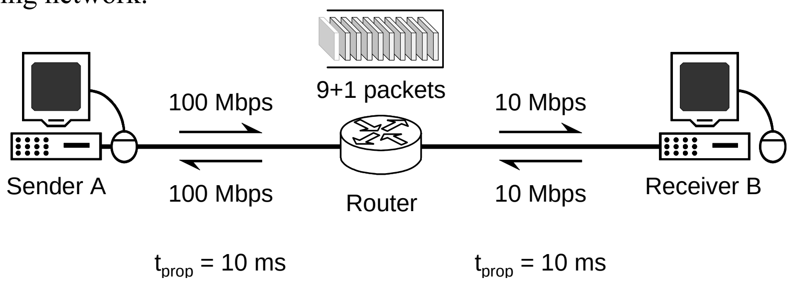

The following assumptions are made:

A1. Full duplex links connect the router to each endpoint host so that simultaneous transmissions are possible in both directions on each link. The link transmission rates are as indicated. One-way propagation delay on each link equals 10 ms. Assume that all packet transmissions are error free.

A2. Each data segment sent by the sender is 1250 bytes long. You can ignore all header overheads, so the transmission delay for data packets over a 100 Mbps link is exactly 0.1 ms and over 10 Mbps is exactly 1 ms. Also assume that the ACK packet size is negligible, i.e., their transmission delay is approximately zero.

A3. The router buffer can hold up to nine packets plus one packet currently in transmission. The packets that arrive to a full buffer are dropped. However, this does *not* apply to ACK packets, i.e., ACKs do not experience congestion or loss.

A4. The receiver does not use delayed ACKs, i.e., it sends an ACK immediately after receiving a data segment.

A5. The receiver has set aside a buffer of RcvBuffer = 64 Kbytes for the received segments.

Answer the following questions:

(a) What is the minimum possible time interval between receiving two consecutive ACKs at the sender?

(b) Write down the transmission start times for the first 7 segments.

(c) Write down the congestion widow sizes for the first 6 transmission rounds, i.e., the first 6 RTTs. (Hint: Try to figure out the pattern of packet arrivals and departures on the router, to understand how the queue of packets grows and when the buffer is fully occupied, so the next packet is dropped.)

(d) In which round will the first packet be dropped at the router? What is the ordinal number of the first dropped packet, starting with #1 for the first packet? Explain your answer.

(e) What is the congestion window size at the 11th transmission round?

(f) What is the long-term utilization of the TCP sender (ignore the initial period until it stabilizes)?

(g) What is the long-term utilization of the link between the router and the receiver (again, ignore the initial period until it stabilizes)?

(h) What will change if delayed ACKs are used to acknowledge cumulatively multiple packets?

(i) Estimate the sender utilization under the delayed ACKs scenario.

Problem 2.11

Consider a TCP Tahoe sender working with MSS = 1 KB, and the bottleneck link bandwidth equal to 1 Mbps. Ignore the initial slow-start phase and assume that the network exhibits periodic behavior where every tenth packet is lost. Consider three different scenarios where all parameters remain the same except for the round-trip time, which changes as: RTT_1 = 0.01 sec, RTT_2 = 0.1 sec, and RTT_3 = 1 sec.

What will be the average rate at which this sender sends data for the different scenarios? Provide an explanation in case you observe any differences between the three scenarios.

Problem 2.12

Calculate the total time required for transferring a 1-MB file from a server to a client in the following cases, assuming an RTT of 100 ms, a segment size of 1 KB, and an initial 2×RTT of "handshaking" (initiated by the client) before data is sent. Assume error-free transmission.

(a) The bottleneck bandwidth is 1.5 Mbps, and data packets can be sent continuously (i.e., without waiting for ACKs)

(b) The bottleneck bandwidth is 1.5 Mbps, but Stop-and-wait ARQ is employed

(c) The bandwidth is infinite, meaning that we take transmission time to be zero, and Go-back-20 is employed

(d) The bandwidth is infinite, and TCP Tahoe is employed

Problem 2.13

Consider a TCP Reno sender, which is in the middle of sending a large amount of data and assume that you are observing it at time t_i. Let t_i, t_{i+1}, t_{i+2}, …, t_{i+7} denote times when the TCP sender will send the subsequent 8 data segments, as governed by its congestion control algorithm. The following assumptions are made:

A1. The TCP sender's segment size equals `MSS` = 200 bytes. At time t_i, the sender is in the *slow start* phase and the congestion window size is already updated as `CongWin`(t_i) = 400 bytes. There are currently no unacknowledged segments. The slow start threshold `SSThresh`(t_i) = 64 Kbytes and the receiver's buffer size `RcvWindow`(t_i) = 1000 bytes.

A2. The sender's sequence number for the next segment that will be transmitted at time t_i equals 30. Assume that the sender transmits back-to-back all the segments that are permitted by its current `EffectiveWindow` size (i.e., the segments are sent in "bursts"). Assume that the segment transmission time is much smaller than the propagation time, i.e., $t_x \ll t_p$ and $t_p \approx$ ½ RTT.

A3. The receiver does not use delayed ACKs, i.e., it sends an ACK immediately after receiving a data segment. All in-order segments are immediately delivered to the application and they never linger in the receive buffer.

A4. The estimated round-trip time at time t_{i-1} equals `EstimatedRTT`(t_{i-1}) = 100 milliseconds, the standard deviation equals `DevRTT`(t_{i-1}) = 10 milliseconds, and `SampleRTT`(t_i) = 106 ms.
Any subsequent transmissions will experience the following round-trip times (from the moment a data segment is transmitted from the sender until the corresponding ACK is received at the sender): `RTT`(t_i) = 105 ms, `RTT`(t_{i+1}) = 93 ms, `RTT`(t_{i+2}) = 179 ms,

$\mathtt{RTT}(t_{i+3})$ = 182 ms, $\mathtt{RTT}(t_{i+4})$ = 165 ms, $\mathtt{RTT}(t_{i+5})$ = 193 ms, $\mathtt{RTT}(t_{i+6})$ = 154 ms, and $\mathtt{RTT}(t_{i+7})$ = 171 ms.
Note: the above values RTT(t) are different from SampleRTT(t), which is the RTT value measured at time t.

Starting at time t_i, consider the subsequent 8 segment transmissions and do the following:

(a) Show the congestion window sizes CongWin(t) and the sequence numbers of the segments transmitted from the sender at times $t = t_i$, t_{i+1}, t_{i+2}, …, t_{i+7}.

(b) Show the sequence numbers of the corresponding acknowledgements and indicate the times when the ACKs will arrive. Also show the values of RcvWindow(t) as carried in each acknowledgement packet.

(c) Show the values of EstimatedRTT(t) and DevRTT(t) as measured by the TCP retransmission-timer management algorithm.

(d) Indicate the times when the TCP sender will set its retransmission timer, if any, as dictated by the TCP algorithm and write down the values of TimeoutInterval(t).

Problem 2.14

Problem 2.15

TCP NewReno RTO Timeout Timer Calculation

Consider the evolution of TCP NewReno parameters shown in Figure 2-19 for Example 2.2. Starting with time t = 89 when segment #35 is retransmitted, show the values of TimeoutInterval(t), calculated using Eq. (2.2). Stop when the ACK for the retransmitted #41 arrives, which will happen at t =120 and show the value of TimeoutInterval(120). Assume that at time t = 88, EstimatedRTT(88) = 6, DevRTT(88) = 0.05, and the values of the control parameters α = 0.125 and β = 0.25.

Follow the procedure for computing TimeoutInterval(t) explained in Section 2.1.2 (and summarized in the pseudocode at the end of this section) as closely as possible. Explain how you obtained every new value of TimeoutInterval(t).

Chapter 3
Multimedia and Real-time Applications

Contents

3.1 Application Requirements

People needs determine the system requirements.

In some situations it is necessary to consider human users as part of an end-to-end system, treating them as active participants, rather than passive receivers of information. For instance, people have thresholds of boredom, and finite reaction times. A specification of user's perceptions is thus required, as it is the user that ultimately defines whether the result has the right quality level.

A *traffic model* summarizes the expected "typical" behavior of a source or an aggregate of sources. Of course, this is not necessarily the ultimate source of the network traffic. The model may consider an abstraction by "cutting" a network link or a set of link at any point in the network and considering the aggregate "upstream" system as the source(s).

Traffic models fall into two broad categories. Some models are obtained by detailed traffic measurements of thousands or millions of traffic flows crossing the physical link(s) over days or years. Others are chosen because they are amenable to mathematical analysis. Unfortunately, only a few models are both empirically obtained and mathematically tractable.

Two key traffic characteristics are:

- Message arrival rate

- Message servicing time

Message (packet) arrival rate specifies the average number of packets generated by a given source per unit of time. Message servicing time specifies the average duration of servicing for messages of a given source at a given server (intermediary). Within the network, packet-servicing time comprises not much more than inspection for correct forwarding plus the transmission time, which is directly proportional to the packet length.

In the following analysis, we will usually assume that the traffic source is infinite, because an infinite source is easier to describe mathematically. For a finite source, the arrival rate is affected by the number of messages already sent; indeed, if all messages are already sent, the arrival rate drops to zero. If the source sends finite but large number of messages, we assume an infinite source to simplify the analysis.

Traffic models commonly assume that packets arrive as a Poisson process, that is, the interarrival time between calls is drawn from an exponential distribution.

Packet servicing times have traditionally been modeled as drawn from an exponential distribution. That is, the probability that the servicing lasts longer than a given length x decreases exponentially with x. However, recent studies have shown servicing times to be *heavy-tailed*. Intuitively, this means that many packets are very long. More precisely, if T_p represents the packet servicing time, and $c(t)$ is defined to be a slowly varying function of t when t is large, the probability that the packet is serviced longer than t is given by:

$$P(T > t) = c(t)\cdot t^{-\alpha} \text{ as } t\to\infty,\ 1 < \alpha < 2$$

As Figure 3-xxx shows, a heavy-tailed distribution has a significantly higher probability mass at large values of t than an exponential function.

3.1.1 Application Types

Multimedia application bandwidth requirements range from G.729 8Kbps speech codec and H.263 64Kbps video codec to 19.2 Mbps for MPEG2, P, 4:2:0 (US standard) based videoconferencing and 63Mbps SXGA 3D computer games [DuVal & Siep 2000]. In general, the higher the speech sampling rate, the better the potential call quality (but at the expense of more bandwidth being consumed). For example, G.711 encoding standard for audio provides excellent quality. Data is delivered at 64 Kbps, and the codec imposes no compression delay. Technically, G.711 delivers 8,000 bytes per second without compression so that full Nyquist-dictated samples are provided.

Applications may also have periodic traffic for real-time applications, aperiodic traffic for web browsing clients, aperiodic traffic with maximum response times for interactive devices like the mouse and keyboard, and non-real time traffic for file transfers. Thus, we see that the range of bandwidth and timeliness requirements for multimedia applications is large and diverse.

Table 3-1: Characteristics of traffic for some common sources/forms of information.

Source	Traffic type	Arrival rate/Service time	Size or Rate
Voice	CBR	Deterministic/ Deterministic	64 Kbps

Video	CBR	Deterministic/ Deterministic	64 Kbps, 1.5 Mbps
	VBR	Deterministic/Random	Mean 6 Mbps, peak 24 Mbps
Text	ASCII	Random/Random	2 KB/page
	Fax	Random/ Deterministic	50 KB/page
Picture	600 dots/in, 256 colors, 8.5 × 11 in	Random/ Deterministic	33.5 MB
	70 dots/in, b/w, 8.5 × 11 in	Random/ Deterministic	0.5 MB

Table 3-1 presents some characteristics about the traffic generated by common forms of information. Notice that the bit streams generated by a video signal can vary greatly depending on the compression scheme used. When a page of text is encoded as a string of ASCII characters, it produces a 2-Kbyte string; when that page is digitized into pixels and compressed as in facsimile, it produces a 50-KB string. A high-quality digitization of color pictures (similar quality to a good color laser printer) generates a 33.5-MB string; a low-quality digitization of a black-and-white picture generates only a 0.5-MB string.

We classify all traffic into three types. A user application can generate a constant bit rate (CBR) stream, a variable bit rate (VBR) stream, or a sequence of messages with different temporal characteristics. We briefly describe each type of traffic, and then consider some examples.

Constant Bit Rate (CBR)

To transmit a voice signal, the telephone network equipment first converts it into a stream of bits with constant rate of 64 Kbps. Some video-compression standards convert a video signal into a bit stream with a constant bit rate (CBR). For instance, MPEG-1 is a standard for compressing video into a constant bit rate stream. The rate of the compressed bit stream depends on the parameters selected for the compression algorithm, such as the size of the video window, the number of frames per second, and the number of quantization levels. MPEG-1 produces a poor quality video at 1.15 Mbps and a good quality at 3 Mbps.

Voice signals have a rate that ranges from about 4 Kbps when heavily compressed and low quality to 64 Kbps. Audio signals range in rate from 8 Kbps to about 1.3 Mbps for CD quality.

Variable Bit Rate (VBR)

Some signal-compression techniques convert a signal into a bit stream that has variable bit rate (VBR). For instance, MPEG-2 is a family of standards for such variable bit rate compression of video signals. The bit rate is larger when the scenes of the compressed movies are fast moving than when they are slow moving. Direct Broadcast Satellite (DBS) uses MPEG-2 with an average rate of 4 Mbps.

To specify the characteristics of a VBR stream, the network engineer specifies the average bit rate and a statistical description of the fluctuations of that bit rate. More about such descriptions will be said later.

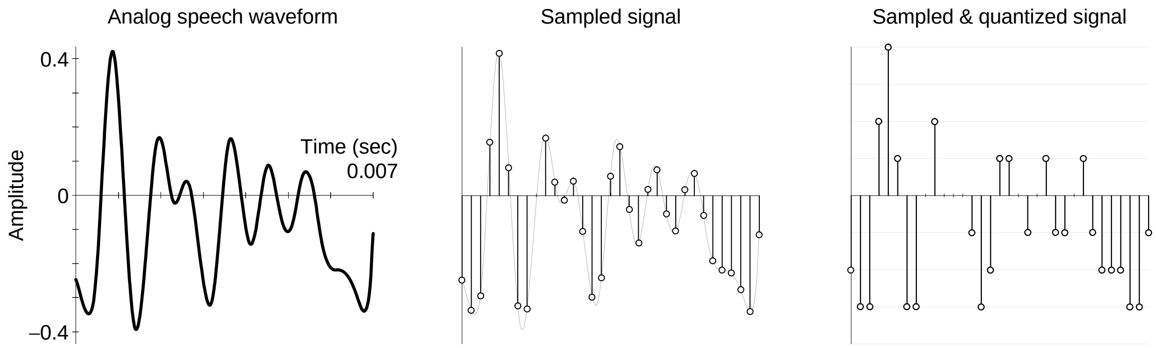

Figure 3-1: Analog speech signal sampling, and quantization to 4 bits.

Messages

Many user applications are implemented as processes that exchange messages over a network. An example is Web browsing, where the user sends requests to a web server for Web pages with embedded multimedia information and the server replies with the requested items. The message traffic can have a wide range of characteristics. Some applications, such as email, generate isolated messages. Other applications, such as distributed computation, generate long streams of messages. The rate of messages can vary greatly across applications and devices.

To describe the traffic characteristics of a message-generating application, the network engineer may specify the average traffic rate and a statistical description of the fluctuations of that rate, in a way similar to the case of a VBR specification.

See definition of fidelity in:

B. Noble, "System support for mobile, adaptive applications," IEEE Personal Communications, 7(1), pp.44-49, February 2000.

E. de Lara, R. Kumar, D. S. Wallach, and W. Zwaenepoel, "Collaboration and Multimedia Authoring on Mobile Devices," *Proc. First Int'l Conf. Mobile Systems, Applications, and Services (MobiSys 2003)*, San Francisco, CA, pp. 287-301, May 2003.

In any scenario where information is communicated, two key aspects of information are fidelity and timeliness. Higher fidelity implies greater quantity of information, thus requiring more resources. The system resources may be constrained, so it may not be possible to transmit, store, and visualize at a particular fidelity. If memory and display are seen only as steps on information's way to a human consumer, then they are part of the communication channel. The user could experience pieces of information at high fidelity, sequentially, one at a time, but this requires time and, moreover, it requires the user to assemble in his or her mind the pieces of the puzzle to experience the whole. Some information must be experienced within particular temporal and or spatial (structural?) constraints to be meaningful. For example, it is probably impossible to experience music one note at a time with considerable gaps in between. Or, a picture cannot be experienced one pixel at a time. Therefore, the user has to trade fidelity for temporal or spatial capacity of the communication channel.

Information loss may sometimes be tolerable; e.g., if messages contain voice or video data, most of the time the receiver can tolerate some level of loss.

Shannon had to introduce fidelity in order to make problem tractable [Shannon & Weaver 1949].

Information can be characterized by fidelity ~ info content (entropy). The effect of a channel can be characterized as deteriorating information's fidelity and increasing the latency:

$$\text{fidelity}_{IN} + \text{latency}_{IN} \rightarrow ()\_\_\_\_\_) \rightarrow \text{fidelity}_{OUT} + \text{latency}_{OUT}$$

Wireless channels in particular suffer from limitations reviewed in Volume 2. Increasing the channel capacity to reduce latency is usually not feasible—either it is not physically possible or it is too costly.

Information qualities can be considered in many dimensions. We group them in two opposing ones:

- Those that tend to increase the information content
- Delay and its statistical characteristics

The computing system has its limitations as well. If we assume *finite buffer length*, then in addition to *delay* problem, there is a *random loss* problem. This further affects the fidelity. Fidelity has different aspects, such as:

- Spatial (sampling frequency in space and quantization – see Brown&Ballard CV book)
- Temporal (sampling frequency in time)
- Structural (topologic, geometric, …)

Delay or latency may also be characterized with more parameters than just instantaneous value, such as the amount of variability of delay, also called delay jitter. In real life both fidelity and latency matter and there are thresholds for each, below which information becomes useless. The system is forced to manipulate the fidelity in order to meet the latency constraints. A key question is, how faithful should signal be in order to be quite satisfactory without being too costly? In order arrive at a right tradeoff between the two, the system must know:

1. Current channel quality parameters, e.g., capacity, which affect fidelity and latency
2. User's tolerances for fidelity and latency

The former determines *what* can be done, i.e., what fidelity/latency can be achieved with the channel at hand, and the latter determines *how* to do it, i.e., what matters more or less to the user at hand. Of course, both channel quality and user preferences change with time.

Example with telephone: sound quality is reduced to meet the delay constraints, as well as reduce the costs.

Targeted reduction of information fidelity in a controlled manner helps meet the latency constraints and averts random loss of information. Common techniques for reducing information fidelity include:

- Lossless and lossy data compression

- Packet dropping (e.g., RED congestion-avoidance mechanism in TCP/IP)
- …?

The above presentation is a simplification in order to introduce the problem. Note that there are many other relevant parameters, such as security, etc., that characterize the communicated information and will be considered in detail later.

Organizational concerns:

- Local traffic that originates at or terminates on nodes within an organization (also called autonomous system, AS)
- Transit traffic that passes through an AS

3.1.2 Standards of Information Quality

In text, the entropy per character depends on how many values the character can assume. Because a continuous signal can assume an infinite number of different value at a sample point, we are led to assume that a continuous signal must have an entropy of an infinite number of bits per sample. This would be true if we required absolutely accurate reproduction of the continuous signal. However, signals are transmitted to be heard, seen, or sensed. Only a certain degree of fidelity of reproduction is required. Thus, in dealing with the samples which specify continuous signals, Shannon introduces *fidelity criterion*. To reproduce the signal in a way meeting the fidelity criterion requires only a finite number of binary digits per sample per second, and hence we say that, within the accuracy imposed by a particular fidelity criterion, the entropy of a continuous source has a particular value in bits per sample or bits per second.

Standards of information quality help perform ordering of information bits by importance (to the user).

Man best handles information if encoded to his abilities. (Pierce, p.234)

In some cases, we can apply common sense in deciding user's servicing quality needs. For example, in applications such as voice and video, users are somewhat tolerable of information loss, but very sensitive to delays. Conversely, in file transfer or electronic mail applications, the users are expected to be intolerable to loss and tolerable to delays. Finally, there are applications where both delay and loss can be aggravating to the user, such as in the case of interactive graphics or interactive computing applications.

For video, expectations are low

For voice, ear is very sensitive to jitter and latencies, and loss/flicker

Voice communication requires a steady, predictable packet delivery rate in order to maintain quality. Jitter, which is variation in packet delivery timing, is the most common culprit that reduces call quality in Internet telephony systems. Jitter causes the audio stream to become broken, uneven or irregular. As a result, the listener's experience becomes unpleasant or intolerable.

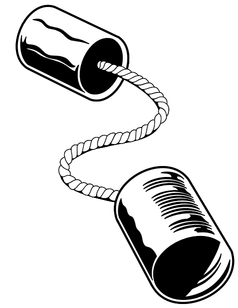

The end results of packet loss are similar to those of jitter but are typically more severe when the rate of packet loss is high. Excessive latency can result in unnatural conversation flow where there is a delay between words that one speaks versus words that one hears. Latency can cause callers to talk over one another and can also result in echoes on the line. Hence, jitter, packet loss and latency can have dramatic consequences in maintaining normal and expected call quality.

Human users are not the only recipients of information. For example, network management system exchanges signaling packets that may never reach human user. These packets normally receive preferential treatment at the intermediaries (routers), and this is particularly required during times of congestion or failure.

It is particularly important during periods of congestion that traffic flows with different requirements be differentiated for servicing treatments. For example, a router might transmit higher-priority packets ahead of lower-priority packets in the same queue. Or a router may maintain different queues for different packet priorities and provide preferential treatment to the higher priority queues.

User Studies

User studies uncover the degree of service degradation that the user is capable of tolerating without significant impact on task-performance efficiency. A user may be *willing* to tolerate inadequate QoS, but that does not assure that he or she will be *able* to perform the task adequately.

Psychophysical and cognitive studies reveal population levels, not individual differences. Context also plays a significant role in user's performance.

The human senses seem to perceive the world in a roughly logarithmic way. The eye, for example, cannot distinguish more than six degrees of brightness; but the actual range of physical brightness covered by those six degrees is a factor of $2.5 \times 2.5 \times 2.5 \times 2.5 \times 2.5 \times 2.5$, or about 100. A scale of a hundred steps is too fine for human perception. The ear, too, perceives approximately logarithmically. The physical intensity of sound, in terms of energy carried through the air, varies by a factor of a trillion (10^{12}) from the barely audible to the threshold of pain; but because neither the ear nor the brain can cope with so immense a gamut, they convert the unimaginable multiplicative factors into comprehensible additive scale. The ear, in other words, relays the physical intensity of the sound as logarithmic ratios of loudness. Thus a normal conversation may seem three times as loud as a whisper, whereas its measured intensity is actually 1,000 or 10^3 times greater.

Fechner's law in psychophysics stipulates that the magnitude of sensation—brightness, warmth, weight, electrical shock, any sensation at all—is proportional to the logarithm of the intensity of the stimulus, measured as a multiple of the smallest perceptible stimulus. Notice that this way the stimulus is characterized by a pure number, instead of a number endowed with units, like seven pounds, or five volts, or 20 degrees Celsius. By removing the dependence on specific units, we have a general law that applies to stimuli of different kinds. Beginning in the 1950s, serious departures from Fechner's law began to be reported, and today it is regarded more as a historical curiosity than as a rigorous rule. But even so, it remains important approximation …

Define j.n.d. (just noticeable difference)

For the voice or video application to be of an acceptable quality, the network must transmit the bit stream with a short delay and corrupt at most a small fraction of the bits (i.e., the BER must be small). The maximum acceptable BER is about 10^{-4} for audio and video transmission, in the absence of compression. When an audio and video signal is compressed, however, an error in the compressed signal will cause a sequence of errors in the uncompressed signal. Therefore the tolerable BER is much less than 10^{-4} for transmission of compressed signals.

The end-to-end delay should be less than 200 ms for real-time video and voice conversations, because people find larger delay uncomfortable. That delay can be a few seconds for non-real-time interactive applications such as interactive video and information on demand. The delay is not critical for non-interactive applications such as distribution of video or audio programs.

Typical acceptable values of delays are a few seconds for interactive services, and many seconds for non-interactive services such as email. The acceptable fraction of messages that can be corrupted ranges from 10^{-8} for data transmissions to much larger values for noncritical applications such as junk mail distribution.

Among applications that exchange sequences of messages, we can distinguish those applications that expect the messages to reach the destination in the correct order and those that do not care about the order.

3.1.3 User Models

User Preferences

User Utility Functions

Example: Augmented Reality (AR)

{PROBLEM STATEMENT}

Inaccuracy and delays on the alignment of computer graphics and the real world are one of the greatest constrains in registration for augmented reality. Even with current tracking techniques it is still necessary to use software to minimize misalignments of virtual and real objects. Our augmented reality application represents special characteristics that can be used to implement better registration methods using an adaptive user interface and possibly predictive tracking.

{CONSTRAINS}

AR registration systems are constrained by perception issues in the human vision system.

An important parameter of continuous signals is the acceptable frame rate. For virtual reality applications, it has been found that the acceptable frame rate is 20 frames per second (fps), with periodical variations of up to 40% [Watson 97], and maximum delays of 10 milliseconds [Azuma, 1995]. The perception of misalignment by the human eye is also restrictive. Azuma found experimentally that it is about 2-3 mm of error at the length of the arm (with an arm length of about 70 cm) is acceptable [Azuma 95]. However, the human eye can detect even smaller differences as of one minute of arc [Doenges 85]. Current commercially available head-mounted displays used for AR cannot provide more than 800 by 600 pixels, this resolution makes impossible to provide an accuracy one minute of arc.

{SOURCES OF ERROR}

Errors can be classified as static and dynamic. Static errors are intrinsic on the registration system and are present even if there is no movement of the head or tracked features.

Most important static errors are optical distortions and mechanical misalignments on the HMD, errors in the tracking devices (magnetic, differential, optical trackers), incorrect viewing parameters as field of view, tracker-to-eye position and orientation. If vision is used to track, the optical distortion of the camera also has to be added to the error model.

Dynamic errors are caused by delays on the registration system. If a network is used, dynamic changes of throughput and latencies become an additional source of error.

{OUR AUGMENTED REALITY SYSTEM}

Although research projects have addressed some solutions for registrations involving predictive tracking [Azuma 95] [Chai 99] we can extended the research because our system has special characteristics (many of these approaches). It is necessary to have accurate registration most of its usage however it is created for task where there is limited movement of the user, as in a repairing task. Delays should be added to the model if processing is performed on a different machine. Also there is the necessity of having a user interface that can adapt to registration changes or according to the task being developed, for example removing or adding information only when necessary to avoid occluding the view of the AR user.

{PROPOSED SOLUTION}

The proposed solution is based on two approaches: predictive registration and adaptive user interfaces. Predictive registration allows saving processing time, or in case of a networked system it can provide better registration in presence of latency and jitter. With predictive registration delays as long as 80ms can be tolerated [Azuma 94]. A statistical model of Kalman filters and extended Kalman filters can be used to optimize the response of the system when multiple tracking inputs as video and inertial trackers are used [Chai 99].

Adaptive user interfaces can be used to improve the view of the augmented world. This approach essentially takes information form the tracking system to determine how the graphics can be gracefully degraded to match the real world. Estimation of the errors was used before to get and approximated shape of the 3D objects being displayed [MacIntyre 00]. Also some user interface techniques based on heuristics where used to switch different representations of the augmented world [Höllerer 01]. The first technique has a strict model to get an approximated AR view but it degrades the quality of the graphics, specially affecting 3D models. The second technique degrades more gracefully but the heuristics used are not effective for all the AR systems. A combination would be desirable.

{TRACKING PIPELINE}

This is a primary description of our current registration pipeline

Image Processing

[Frame capture] => [Image threshold] => [Subsampling] => [Features Finding] =>

[Image undistortion] => [3D Tracking information] => [Notify Display]

Video Display

[Get processed frame] => [Frame rendering in a buffer] => [3D graphics added to Buffer]

=> [Double buffering] => [Display]

These processes are executed by two separated threads for better performance and resource usage.

{REFERENCES}

[Watson 97]

Watson, B., Spaulding, V., Walker, N., Ribarsky W., "Evaluation of the effects of frame time variation on VR task performance," IEEE VRAIS'96, 1996, pp. 38-52.

http://www.cs.northwestern.edu/~watsonb/school/docs/vr97.pdf

[Azuma 95]

R. Azuma, "Predictive Tracking for Augmented Reality," UNC Chapel Hill Dept. of Computer Science Technical Report TR95-007 (February 1995), 262 pages.

http://www.cs.unc.edu/~azuma/dissertation.pdf

[Doenges 85]

P. K. Doenges, “Overview of Computer Image Generation in Visual Simulation,” SIGGRAPH '85 Course Notes #14 on High Performance Image Generation Systems (San Francisco, CA, 22 July 1985).

(Not available on the web, cited in Azuma’s paper)

[Chai 99]

L. Chai, K. Nguyen, W. Hoff, and T. Vincent, “An adaptive estimator for registration in augmented reality,” Proc. of 2nd IEEE/ACM Int'l Workshop on Augmented Reality, San Franscisco, Oct. 20-21, 1999.

http://egweb.mines.edu/whoff/projects/augmented/iwar1999.pdf

[Azuma 94]

R. Azuma and G. Bishop, “Improving Static and Dynamic Registration in an Optical See-Through HMD,” Proceedings of SIGGRAPH '94 (Orlando, FL, 24-29 July 1994), Computer Graphics, Annual Conference Series, 1994, 197-204.

http://www.cs.unc.edu/~azuma/sig94paper.pdf

[MacIntyre 00]

B. MacIntyre, E. Coelho, S. Julier, “Estimating and Adapting to Registration Errors in Augmented Reality Systems,” In IEEE Virtual Reality Conference 2002 (VR 2002), pp. 73-80, Orlando, Florida, March 24-28, 2002.

http://www.cc.gatech.edu/people/home/machado/papers/vr2002.pdf

[Hollerer 01]

T. Höllerer, D. Hallaway, N. Tinna, S. Feiner, “Steps toward accommodating variable position tracking accuracy in a mobile augmented reality system,” In Proc. 2nd Int. Workshop on Artificial Intelligence in Mobile Systems (AIMS '01), pages 31-37, 2001.

http://monet.cs.columbia.edu/publications/hollerer-2001-aims.pdf

3.1.4 Performance Metrics

Delay (the average time needed for a packet to travel from source to destination), statistics of delay (variation, jitter), packet loss (fraction of packets lost, or delivered so late that they are considered lost) during transmission, packet error rate (fraction of packets delivered in error);

Bounded delay packet delivery ratio (BDPDR): Ratio of packets forwarded between a mobile node and an access point that are successfully delivered within some pre-specified delay constraint. The delay measurement starts at the time the packet is initially queued for transmission (at the access point for downstream traffic or at the originating node for upstream traffic) and ends when it is delivered successfully at either the mobile node destination (downstream traffic) or AP (upstream traffic).

Quality of Service

QoS, Keshav p.154

Cite Ray Chauduhuri's W-ATM paper {cited in Goodman}

Quality of Service (QoS)

Performance measures

Throughput

Latency

Real-time guarantees

Other factors

Reliability

Availability

Security

Synchronization of data streams

Etc.

A networking professional may be able to specify what quality-of-service metrics are needed, and can specify latency, packet loss and other technical requirements. However, the consumer or independent small-office-home-office (SOHO) user would more easily understand service classifications such as "High-Definition Movie Tier" or an "Online Gamer Tier." Few consumers will be able to specify service-level agreements, but they may want to know if they are getting better services when they pay for them, so a consumer-friendly reporting tool would be needed. In addition, although enterprises are increasingly likely to buy or use a premise-based session border controller to better manage IP traffic, service providers will need to come up with an easier and less expensive alternative to classify consumer IP packets based on parameters such as user profiles and service classes.

3.2 Source Characteristics and Traffic Models

Different media sources have different traffic characteristics.

3.2.1 Traffic Descriptors

Some commonly used traffic descriptors include peak rate and average rate of a traffic source.

Average Rate

The average rate parameter specifies the average number of packets that a particular flow is allowed to send per unit of time. A key issue here is to decide the interval of time over which the average rate will be regulated. If the interval is longer, the flow can generate much greater number of packets over a short period than if the interval is short. In other words, a shorter averaging interval imposes greater constraints. For example, average of 100 packets per second is different from an average of 6,000 packets per minute, because in the later can the flow is allowed to generate all 6,000 over one 1-second interval and remain silent for the remaining 59 seconds.

Researchers have proposed two types of average rate definitions. Both […]

Burst size: this parameter constrains the total number of packets (the "burst" of packets) that can be sent by a particular flow into the network over a short interval of time.

Peak Rate

The peak rate is the highest rate at which a source can generate data during a communication session. Of course, the highest data rate from a source is constrained by the data rate of its outgoing link. However, by this definition even a source that generates very few packets on a 100-Mbps Ethernet would be said to have a peak rate of 100 Mbps. Obviously, this definition does not reflect the true traffic load generated by a source. Instead, we define the **peak rate** as the maximum number of packets that a source can generate over a very short period of time. In the above example, one may specify that a flow be constrained to an average rate of 6,000 packets/minute and a peak rate of 100 packets/second.

Primitive traffic characterization is given by the source entropy.

See also MobiCom'04, p. 174: flow characterization

For example, image transport is often modeled as a two state on-off process. While on, a source transmits at a uniform rate. For more complex media sources such as variable bit rate (VBR) video coding algorithms, more states are often used to model the video source. The state transitions are often assumed Markovian, but it is well known that non-Markovian state transitions could also be well represented by one with more Markovian states. Therefore, we shall adopt a general Markovian structure, for which a deterministic traffic rate is assigned for each state. This is the well-known Markovian fluid flow model [Anick *et al*. 1982], where larger communication entities, such as an image or a video frame, is in a sense "fluidized" into a fairly smooth flow of very small information entities called cells. Under this fluid assumption, let $X_i(t)$ be the rate of cell emission for a connection i at the time t, for which this rate is determined by the state of the source at time t.

The most common modeling context is queuing, where traffic is offered to a queue or a network of queues and various performance measures are calculated.

Simple traffic consists of single arrivals of discrete entities (packets, frames, etc.). It can be mathematically described as a point process, consisting of a sequence of arrival instants T_1, T_2, …, T_n, … measured from the origin 0; by convention, $T_0 = 0$. There are two additional equivalent descriptions of point processes: *counting processes* and *interarrival time processes*. A counting process $\{N(t)\}_{t=0}^{\infty}$ is a continuous-time, non-negative integer-valued stochastic process, where $N(t) = \max\{n : T_n \le t\}$ is the number of (traffic) arrivals in the interval (0, t]. An interarrival time process is a non-negative random sequence $\{A_n\}_{n=1}^{\infty}$, where $A_n = T_n - T_{n-1}$ is the length of the time interval separating the n-th arrival from the previous one. The equivalence of these descriptions follows from the equality of events:

$$\{N(t) = n\} = \{T_n \le t < T_{n+1}\} = \left\{\sum_{k=1}^{n} A_k \le t < \sum_{k=1}^{n+1} A_k\right\}$$

because $T_n = \sum_{k=1}^{n} A_k$. Unless otherwise stated, we assume throughout that $\{A_n\}$ is a stationary sequence and that the common variance of the A_n is finite.

3.2.2 Self-Similar Traffic

3.3 Approaches to Quality-of-Service

This section reviews some end-to-ed mechanisms for providing quality-of-service (QoS), and hints at mechanisms used in routers. Chapter 5 details the router-based QoS mechanisms.

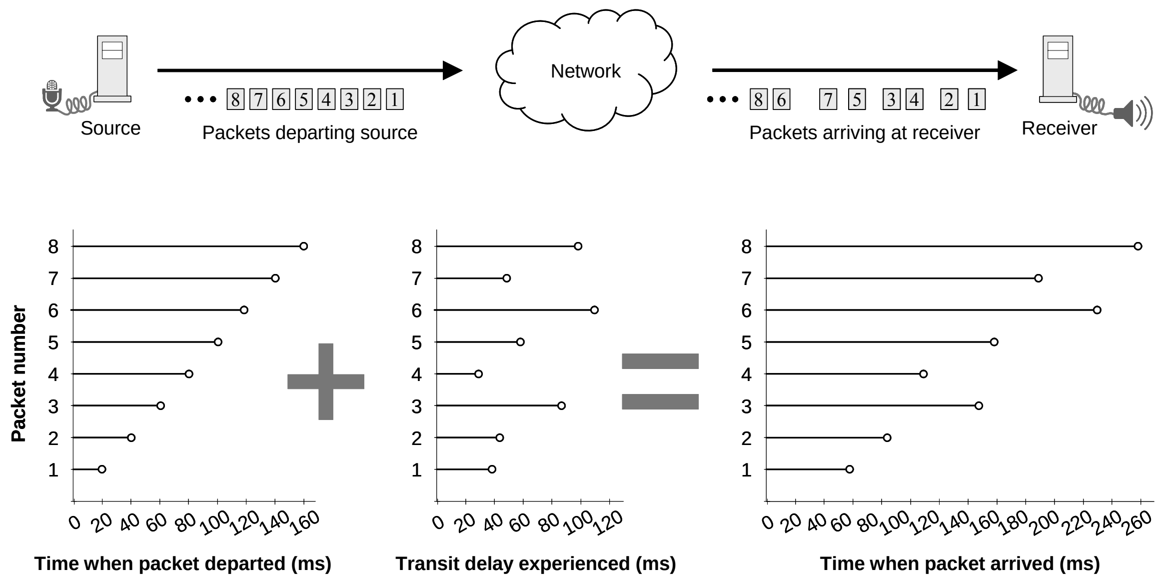

Figure 3-2: Packets depart with a uniform spacing, but they experience variable amount of delay (jitter) and arrive at the receiver irregularly (packets #3 and #6 arrive out of order).

3.3.1 End-to-End Delayed Playout

Problems related to this section: Problem 3.2 → Problem 3.5

Removing Jitter by Delayed Playout

Consider the example shown in Figure 3-2 where a source sends audio signal to a receiver for playout. Let us assume that the source segments the speech stream every 20 milliseconds and creates data packets. The source outputs the packets with a uniform spacing between them, but they arrive at the receiver at irregular times due to random network delays.

Speech packetization at the intervals of 20 ms seems to be a good compromise. If the interval were longer, flicker due to lost or late packets would be more noticeable; conversely, if the interval were shorter, the packet-header overhead would be too high, with the header size possibly exceeding the payload.

Playing out the speech packets as they arrive (with random delays) would create significant distortions and impair the conversation. One way to deal with this is to buffer the packets at the receiving host to smooth out the jitter. Packets are buffered for variable amounts of time in an attempt to play out each speech segment with a constant amount of delay relative to the time when it was pacektized and transmitted from the source. Let us introduce the following notation (see Figure 3-2):

t_i = the time when the i^{th} packet departed its source

d_i = the amount of delay experienced by the i^{th} packet while in transit

r_i = the time when the i^{th} packet is received by receiver (notice that $r_i = t_i + d_i$)

p_i = the time when the i^{th} packet is played at receiver

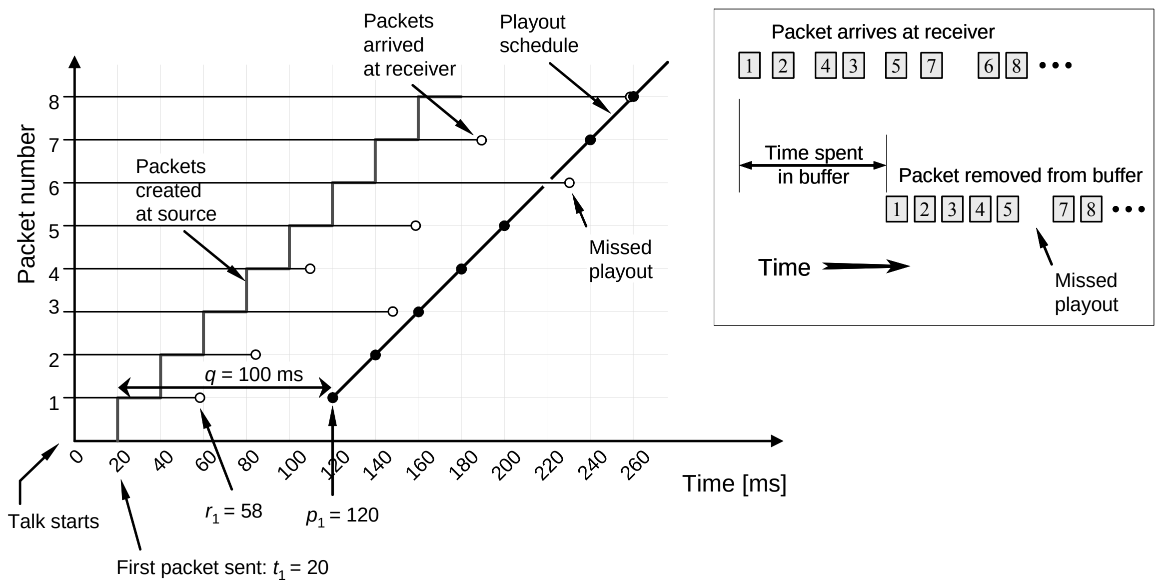

Figure 3-3: Removing jitter at receiver by delaying the playout.

Let q denote the constant delay introduced to smoothen out the playout times. Then $p_i = t_i + q$. The time difference between the i^{th} packet's playout time and the time it is received equals $\Delta_i = (t_i + q) - r_i$. If $\Delta_i \geq 0$, the i^{th} packet should be buffered for this duration before it is played out. If $\Delta_i < 0$, the i^{th} packet should be discarded because it arrived too late for playout. Figure 3-3 illustrates jitter removal for the example given in Figure 3-2. In this case, the constant playout delay of q = 100 ms is selected. With this choice, the sixth packet does not arrive by its scheduled playout time, and the receiver considers it lost.

We could try selecting a large q so that all packets will arrive by their scheduled playout time. However, for applications such as Internet telephony, delays greater than 400 ms are not acceptable because of human psychophysical constraints. Ideally, we would like keep the playout delay less than 150 ms. Larger delays become annoying and it is difficult to maintain a meaningful conversation. We know from the discussion in Section 2.1.2 that average end-to-end network delays can change significantly during day or even during short periods. Therefore, the best strategy is to adjust the playout delay adaptively, so that we select the minimum possible delay for which the fraction of missed playouts is kept below a given threshold.

We can again use the approach described in Section 2.1.2 and estimate the average end-to-end network delay using *Exponential Weighted Moving Average* (EWMA). Similar to Eq. (2.2), we estimate the average network delay $\hat{\delta}_i$ upon reception of the i^{th} packet as

$$\hat{\delta}_i = (1-\alpha)\cdot\hat{\delta}_{i-1} + \alpha\cdot(r_i - t_i)$$

where α is a fixed constant, say, $\alpha = 0.001$. We also estimate the average standard deviation $\hat{\upsilon}_i$ of the delay as

$$\hat{\upsilon}_i = (1-\alpha)\cdot\hat{\upsilon}_{i-1} + \alpha\cdot| r_i - t_i - d_i |$$

Notice that the playout delay q is measured relative to packet's departure time (Figure 3-3). Therefore, we cannot adjust q for each packet individually, because this would still result in distorted speech. An option is to set the playout delay constant for an interval of time, but the

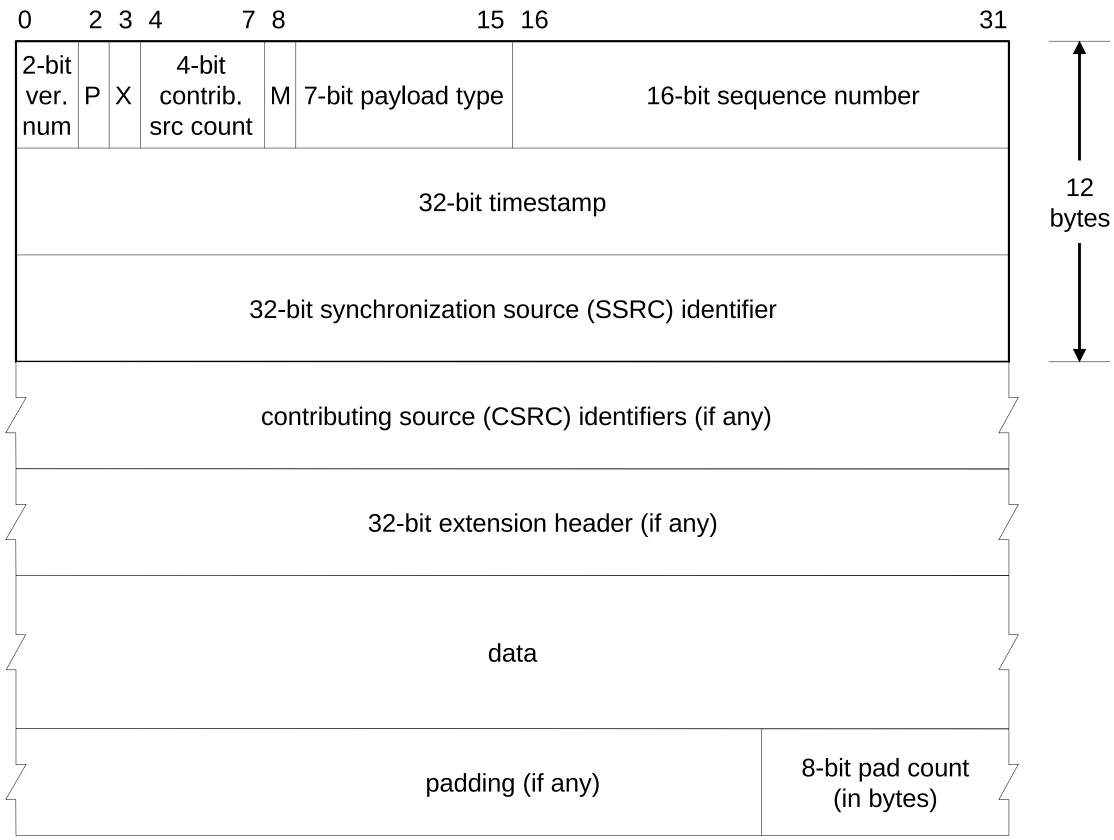

Figure 3-4: Real-time Transport Protocol (RTP) packet format.

question is when this interval should start and how long it should last. It turns out that humans do not notice if the periods of silence between the utterances are stretched or compressed. This fact is used to adjust the playout delay adaptively: the playout delay q is adjusted only at the start of an utterance (or, "talk spurt") and it is maintained constant until the next period of silence. The receiver maintains average delay $\hat{\delta}_i$ and average standard deviation $\hat{\upsilon}_i$ for each received packet. During a period of silence, the receiver calculates the playout delay for the subsequent talk spurt as follows. If packet k is the first packet of the next talk spurt, k^{th}'s playout delay is computed as

$$q_k = \hat{\delta}_k + K \cdot \hat{\upsilon}_k \tag{3.1}$$

where K is a positive constant, for example we can set $K = 4$ following the same reasoning as in Section 2.1.2 for Eq. (2.2). Then, the playout time of the k^{th} packet and all the remaining packets of the next spurt is computed as $p_i = t_i + q_k$.

RTP

The *Real-time Transport Protocol* (RTP) provides the transport of real-time data packets. To accommodate new real-time applications, the protocol specifies only the basics and it is somewhat incomplete. Unlike conventional protocols, RTP can be tailored to specific application needs through modifications and additions to headers. This allows the protocol to adapt easily to new audio and video standards.

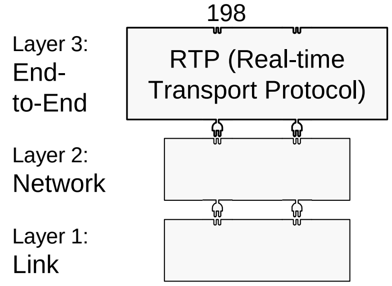

RTP implements the end-to-end layer (or, transport-layer in the OSI model) features needed to provide synchronization of multimedia data streams. Figure 3-4 shows the header format used by RTP.

The first two bits indicate the RTP version.

The "padding" (P) bit is set when the packet contains a set of padding octets that are not part of the payload. For example, RTP data might be padded to fill up a block of a certain size as required by some encryption algorithms.

The extension bit (X) is used to indicate the presence of an extension header, which can be defined for some application's special needs. Such headers are rarely used, because a payload-specific header can be defined as part of the payload format definition used by the application.

The 4-bit contributing-sources count represents the number of contributing source (CSRC) identifiers, if any are included in the header.

The M bit allows significant events to be marked in the packet stream (that is, frame boundaries).

The 7-bit payload type specifies the format of the payload in the RTP packet. An RTP sender emits a single RTP payload type at any given time. An RTP packet can contain portions of either audio or video data streams. To differentiate between these streams, the sending application includes a payload type identifier within the RTP header. The identifier indicates the specific encoding scheme used to create the payload.

The sequence number is used by the receiver when removing jitter at the receiver, as described earlier. It is used to restore the original packet order and detect packet loss. The sequence number increments by one for each RTP data packet sent. The initial value of the sequence number is randomly determined. This makes hacking attacks on encryption more difficult. A random number is used even if the source device does not encrypt the RTP packet. The packets can flow through a translator host or router that does provide encryption services.

The timestamp is used along with the sequence number to detect gaps in a packet sequence. Timestamps are also used in RTP to synchronize packets from different sources. The timestamp represents the sampling (creation) time of the first octet in each RTP data packet. It is derived from a clock that increments monotonically and linearly. The resolution of the timer depends on the desired synchronization accuracy required by the application. It is possible that several consecutive RTP packets have the same timestamp. For example, this can occur when a single video frame is transmitted in multiple RTP packets. Because the payloads of these packets were logically generated at the same instant, their time stamps remain constant. The initial value of the time stamp is random.

The synchronization source (SSRC) identifier is a randomly chosen identifier for an RTP host. All packets from the same source contain the same SSRC identifier. Each device in the same RTP session must have a unique SSRC identifier. This enables the receiver to group packets for playback.

The contributing source (CSRC) identifiers field contains a list of the sources for the payload in the current packet. This field is used when a mixer combines different streams of packets. The information contained in this field allows the receiver to identify the original senders.

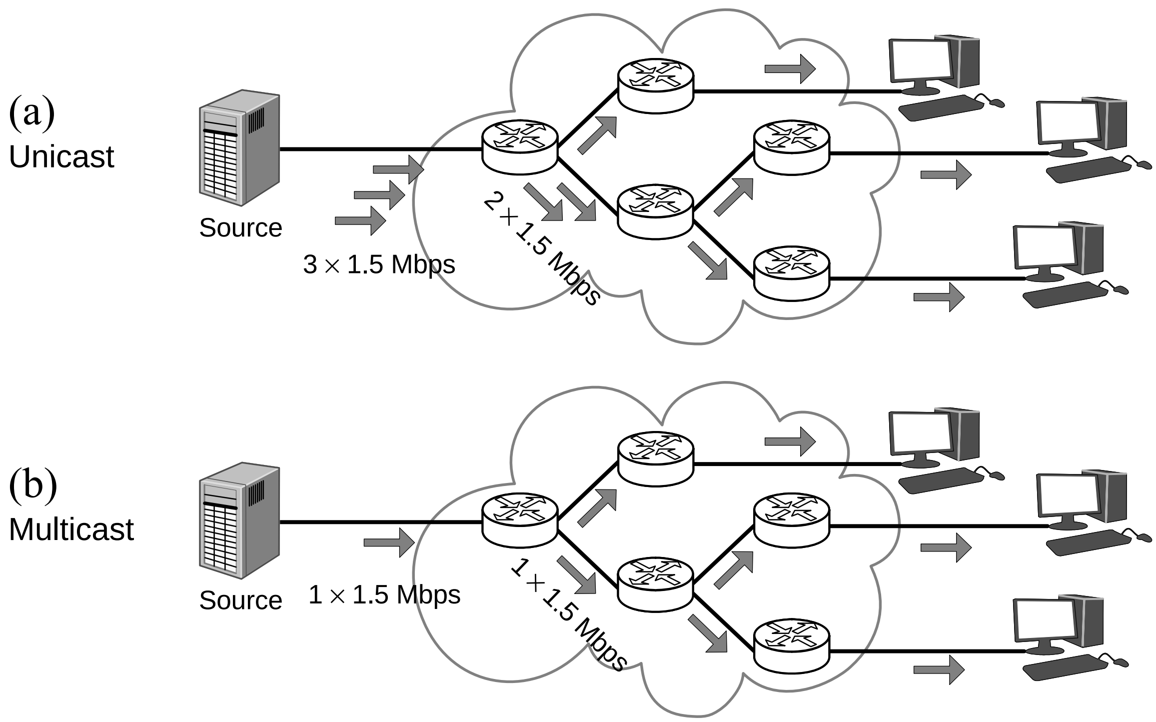

Figure 3-5: Unicast vs. multicast routing.

RTCP

The *Real-Time Control Protocol* (RTCP) monitors the quality of service provided to existing RTP sessions. The primary function of RTCP is to provide feedback about the quality of the RTP data distribution. This is comparable to the flow and congestion control functions provided by other transport protocols, such as TCP. Feedback provided by each receiver is used to diagnose stream-distribution faults. By sending feedback to all participants in a session, the device observing problems can determine if the problem is local or remote. This also enables a managing entity (that is, a network service provider) that is not a participant in the session to receive the feedback information. The network provider can then act as a third-party monitor to diagnose network problems.

3.3.2 Multicast Routing

Problems related to this section: Problem 3.7 → ??

When multiple receivers are required to get the same data at approximately the same time, multicast routing is a more efficient way of delivering data than unicast. A unicast packet has a single source IP address and a single destination IP address. Data are delivered to a single host. A multicast packet has a single source IP, but it has a multicast destination IP address that will be delivered to a group of receivers. (Recall the multicast address class D in Figure 1-45.) The advantage is that multiple hosts can receive the same multicast stream (instead of several individual streams), thereby saving network bandwidth. In general, the bandwidth saving with multicast routing becomes more substantial as the number of destinations increases.

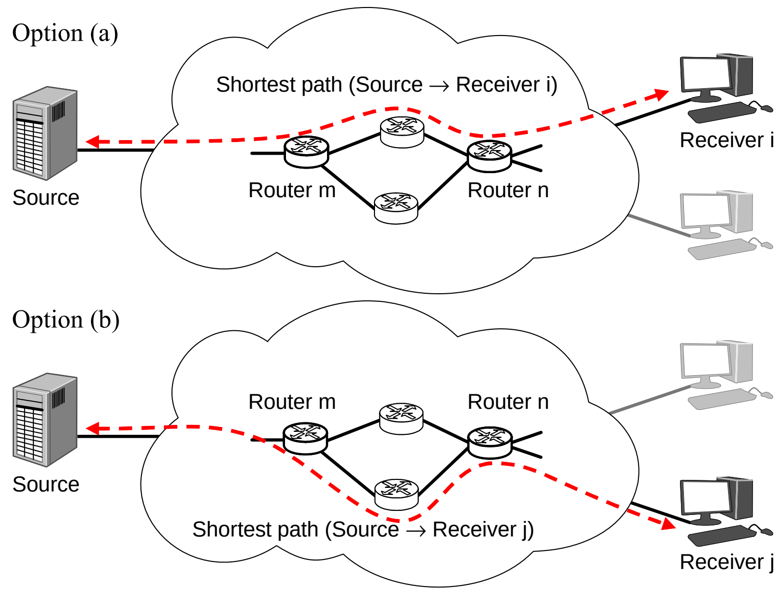

Figure 3-6: The superimposed shortest paths must form a tree rooted in the source.

Figure 3-5 shows an example where three users are simultaneously watching the same video that is streamed from the same video source. In Figure 3-5(a), all users receive their individual streams via a unicast delivery. As illustrated, the source must send a total of 3 unicast streams, each targeted to a different user. Obviously, this is a waste of bandwidth. If the compressed video takes approximately 1.5 Mbps of bandwidth per stream, then the first link must support data rate of at least $3 \times 1.5 = 4.5$ Mbps. Similarly, the lower link from the first router relays two streams which consume $2 \times 1.5 = 3$ Mbps of bandwidth. If, on the other hand, the network supports multicast routing as in Figure 3-5(b), the source sends only a single stream and all links would be using 1.5 Mbps for this video. Of course, this kind of resource saving is possible only if multiple users are downloading simultaneously the same video from the same source.

There are two key issues for multicast routing protocols:

1. **Multicast Group Management:** identifying which hosts are members of a multicast group and supporting dynamic changes in the group membership; multiple sources and multiple receivers may need to be supported
2. **Multicast Route Establishment:** setting up the (shortest path) route from each source to each receiver

A *multicast group* relates a set of sources and receivers with each other, but conceptually exists independently of them. Such group is identified by a unique IP multicast address of Class D (Figure 1-45). It is created either when a source starts sending to the group address (even if no receivers are present) or when a receiver expresses its interest in receiving packets from the group (even if no sources are currently active).

To establish the multicasting routes, we start by superimposing all the shortest paths connecting the source with all the receivers. The result will be a tree, i.e., it cannot be a graph with cycles. To see why, consider a contrary possibility, illustrated in Figure 3-6, where the shortest paths from the source to two receivers *i* and *j* share two intermediary nodes (routers *m* and *n*), but not the nodes between these two nodes. We know from Section 1.4 that the shortest path between two intermediate points does not depend on where this path extends beyond the intermediate points. In other words, it does not depend on the endpoints. Hence, if we superimpose all the shortest paths from the source host to any destination, we will obtain a tree structure, for which the source host is the root node. (Notice that alternative paths between *m* and *n* could be equally long, but there should be a uniform policy to resolve the tied cases.) The next issue is, how the multicast-capable routers should construct this tree.

Reverse Path Forwarding (RPF) Algorithm

We assume that all routers in the network are running a unicast routing algorithm (described in Section 1.4) and maintain unicast routing tables independently of the multicast algorithm. Thus, the routers know either the shortest unicast paths to all nodes in the network, or at least the next hop on the shortest path to any other node in the network.

In **reverse path forwarding (RPF)** algorithm, when a router receives a packet, it forwards the packet to all outgoing links (except the one on which it was received) only if the packet arrived on the link that is on this router's shortest unicast path back to the source. Otherwise, the router simply discards the incoming packet without forwarding it on any of its outgoing links. (A tie between two routers is broken by selecting the router with the smallest network address.)

A problem with RPF is that it essentially *floods* every router in the network, regardless of whether it has hosts attached to it that are interested in receiving packets from the multicast group. To avoid these unnecessary transmissions, we perform **pruning**: the router that no longer has attached hosts interested in receiving multicast packets from a particular source informs the next-hop router on the shortest path to the source that it is not interested.

Here is an example:

Example 3.1 Multicast Using Reverse Path Forwarding (RPF) Algorithm

Consider the network shown in Figure 3-7, in which radio broadcast source *A* distributes a radio program to a multicast group *Γ*, whose members are all the shown hosts. Assume that all link costs are equal to 1, including the Ethernet links, i.e., any two nodes that are separated by one hop.

(a) Draw the shortest path multicast tree for the group *Γ*.

(b) Assuming that the reverse path forwarding (RPF) algorithm is used, how many packets are forwarded in the entire network per every packet sent by the source *A*? To avoid ambiguities, describe how you counted the packets.

(c) Assuming that the RPF algorithm uses pruning to exclude the networks that do not have hosts that are members of *Γ*, how many packets are forwarded in the entire network per every packet sent by the source *A*?

The solutions for (a) and (b) are shown in Figure 3-8. The shortest path multicast tree is drawn by thick lines in Figure 3-8(a). Router R3 is two hops from R2 (the root of the tree) both via R1 and via R4. Router R1 is selected because it has a smaller address than R4. (Therefore, link R4–R3 is not part of the tree!) Notice that router R6 is not connected to R3 (via the multihomed host *E*), because multihomed hosts do *not* participate in routing or forwarding of transit traffic.

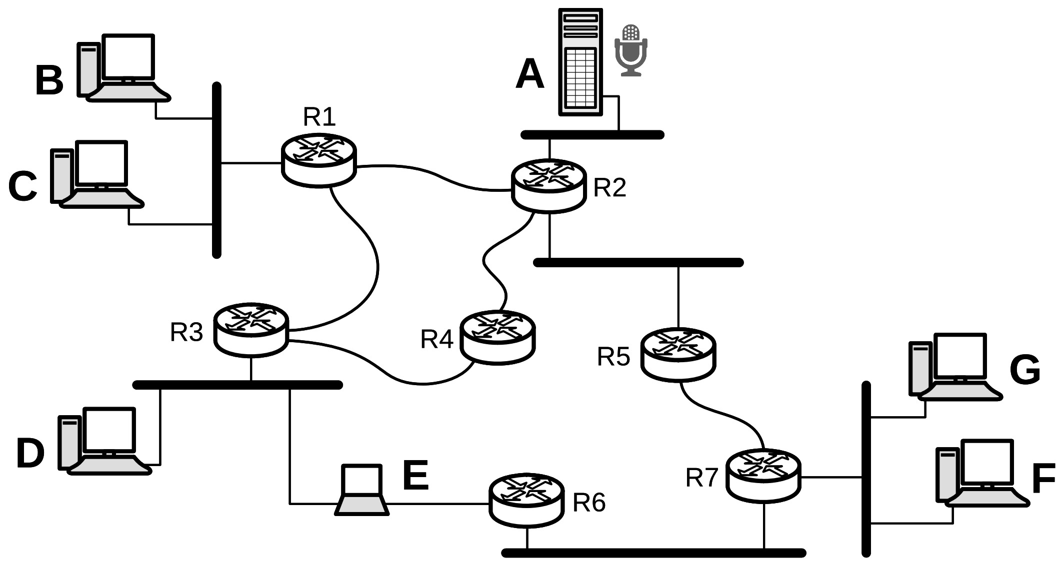

Figure 3-7: Example network used in the multicast Example 3.1.

The root router R2 sends packets to routers R1, R4, and R5, which in turn forward them to all outgoing links (except the one on which it was received) only if the packet arrived on the link that is on its own shortest unicast path back to the source. Otherwise, the router simply discards the incoming packet without forwarding it on any of its outgoing links. In Figure 3-8(b), R3 will receive the packets from both R1 and R4, but it will forward only the one from R1 and discard the one from R4. The way we count packets is how many packets leave the router, and the router has to forward a different packet on each different outgoing link. Therefore, the number of forwarded packets in the entire network is 8 per each sent packet. If we include the packets forwarded to end hosts, the total is 8 + 6 = 14 (shown in Figure 3-8(b)).

As for part (c), routers R4 and R6 do not have any host for which either one is on the shortest path to the source *A*. The shortest path for host *E* is via R3–R1–R2. Therefore, R4 and R6 should send a *prune* message to R2 and R7, respectively, to be removed from the multicast tree. This reduces the number of forwarded packets by 4, so the total number is 4 per each sent packet, or 4 + 6 = 10, if the end hosts are counted.

What if a router is pruned earlier in the session, but later it discovers that some of its hosts wish to receive packets from that multicast group? One option is that the router explicitly sends a **graft message** to the next-hop router on the shortest path to the source. Another option is for the source(s) and other downstream routers to flood packets periodically from the source in search for receivers that may wish to join the group later in the session. This extended version of RPF is called **flood-and-prune** approach to multicast-tree management.

A key property of RPF is that routing loops are automatically suppressed and each packet is forwarded by a router exactly once. The basic assumption underlying RPF is that the shortest path is symmetric in both directions. That is, the shortest path from the source to a given router contains the same links as the shortest path from this router to the source. This assumption requires that each link is symmetric (roughly, that each direction of the link has the same cost). If links are not symmetric, then the router must compute the shortest path from the source to itself, given the information from its unicast routing tables. Notice that this is possible only if a link-state protocol (Section 1.4.2) is used as the unicast routing algorithm.

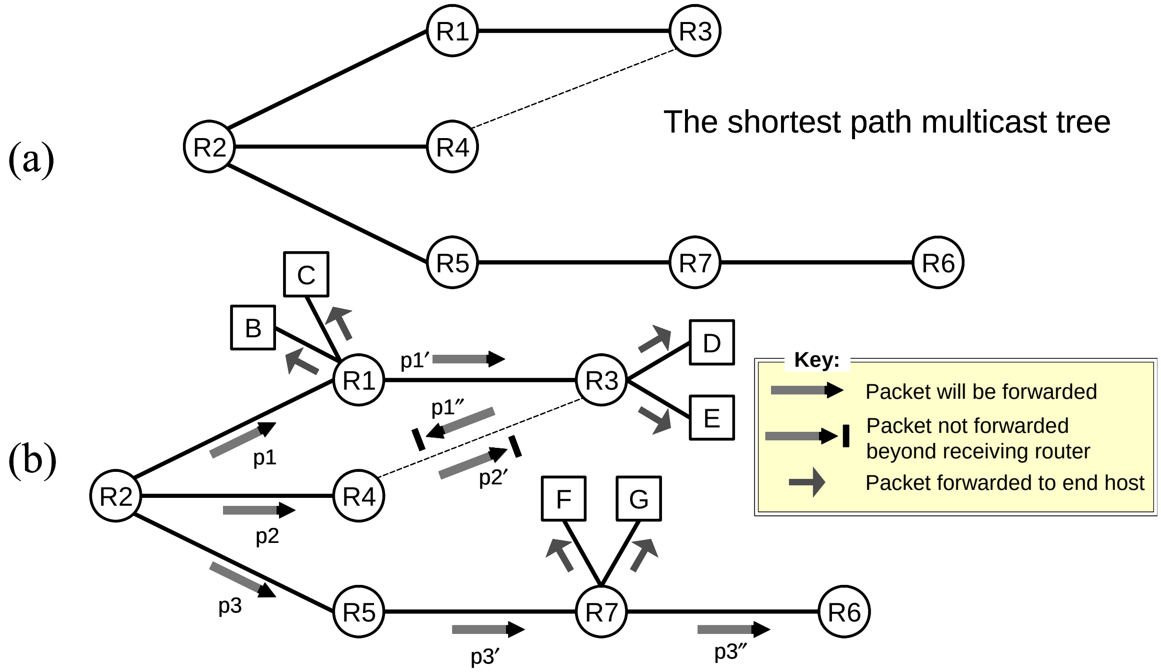

Figure 3-8: The shortest path multicast tree for the example network in Figure 3-7.

Spanning Tree Algorithms

The reverse path forwarding algorithm, even with pruning, does not completely avoid transmission of redundant multicast packets. Consider the network in Figure 3-9(a), which is similar to Figure 3-7 but slightly more complex. The shortest-path multicast tree is shown in Figure 3-9(b). Router R4 can be pruned because it does not have attached hosts that are interested in multicast packets from the source *A*. (Router R5 is relaying packets for R6 and R7, so it stays.) As seen, routers R3, R5, R6, R7, and R8 will receive either one or two redundant packets. Ideally, every node should receive only a single copy of the multicast packet. This would be the case if the nodes were connected only by the thick lines in Figure 3-9(b). The reason is that the thick lines form a *tree* structure, so there are no multiple paths for the packet to reach the same node. A tree that is obtained by removing alternative paths, while keeping connected the nodes that were originally connected, is called a **spanning tree**. If a multicast packet were forwarded from the root of the tree to all other nodes, every node would receive exactly one copy of the packet. If links have associated costs and the total cost of the tree is the sum of its link costs, then the spanning three with a minimum total cost is called a **minimum spanning tree**.

Therefore, an alternative to RPF is to construct a spanning tree and have each source send the packets out on its incident link that belongs to the spanning tree. Any node that receives a multicast packet then forwards it to all of its neighbors in the spanning tree (except the one from where the packet came). Multicasting on a spanning tree requires a total of only $N - 1$ packet transmissions per packet multicast, where N is the number of nodes. Notice that a single spanning tree is sufficient for any number of sources. This is true because any node of a tree can serve as its root. To convince yourself about this, take an arbitrary tree and select any of its nodes. Now imagine that you pull this node up and the other nodes remain hanging from the selected node. What you get is a tree rooted in the selected node.

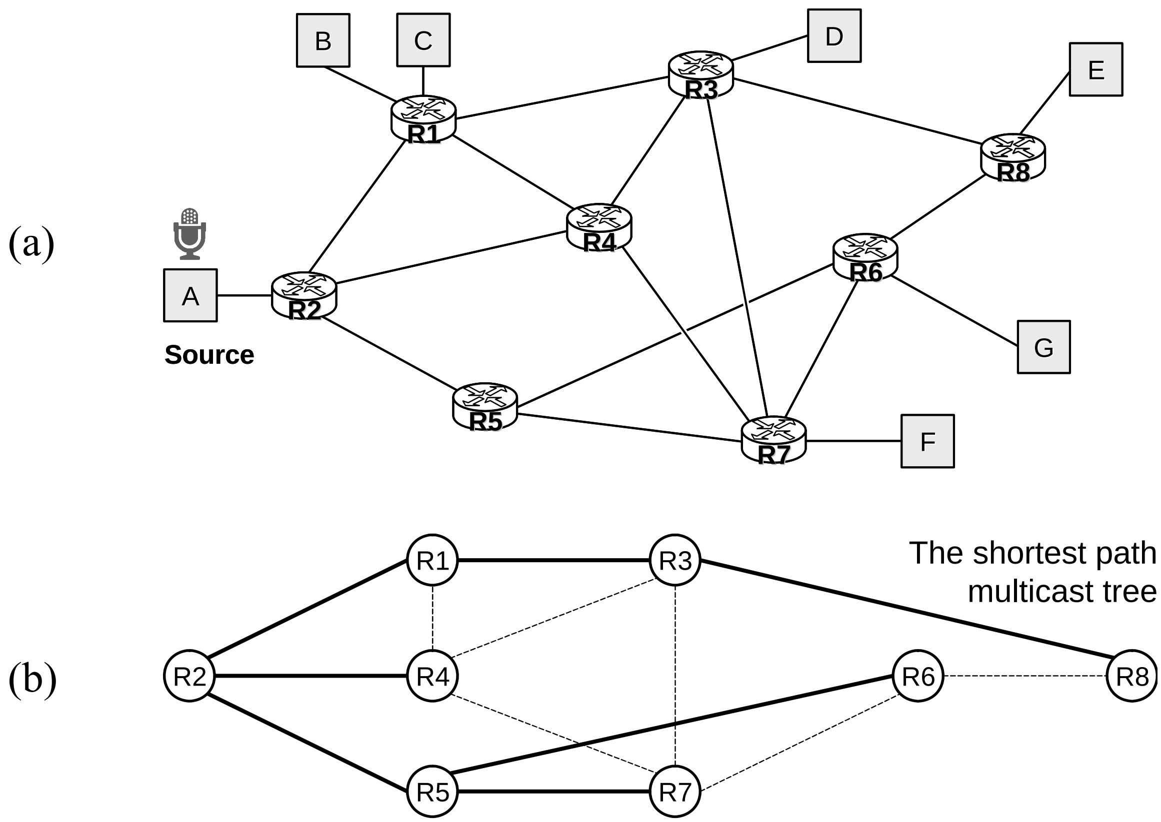

Figure 3-9: Example network for multicast routing.

The main complexity of the spanning-tree multicasting lies in the creation and maintenance of the spanning tree, as sources and receivers dynamically join or leave the multicast group or the network topology changes. (Notice that RPF does not have this problem, because it relies on flooding.) One algorithm that builds and maintains the spanning-tree efficiently is known as **core-based trees (CBTs)**. With CBTs, the spanning tree is formed starting from a "core router" (also known as a "rendezvous point" or a "center node"), which can be statically configured or automatically selected. Other routers are added by growing "branches" of the tree, consisting of a chain of routers, away from the core router out towards the routers directly adjoining the multicast group members. The core router is also known as a "center" and CBT is sometimes called *center-based approach*.

The tree building process starts when a host joins the multicast group. The host sends a join-request packet addressed to the core router. The information about the core router is statically configured. This join-request packet travels hop-by-hop towards the target core, forwarded using unicast routing. The process stops when the packet either arrives at an intermediate router that already belongs to the spanning tree or arrives at the destination (the core router). In either case, the path that the join-request packet has followed defines the branch of the spanning tree between the leaf node that originated the join-request and the core router. The node at which the message terminated confirms the packet by sending a join-acknowledgement message. The join-acknowledgement message travels the same route in the opposite direction the join-request message traveled earlier.

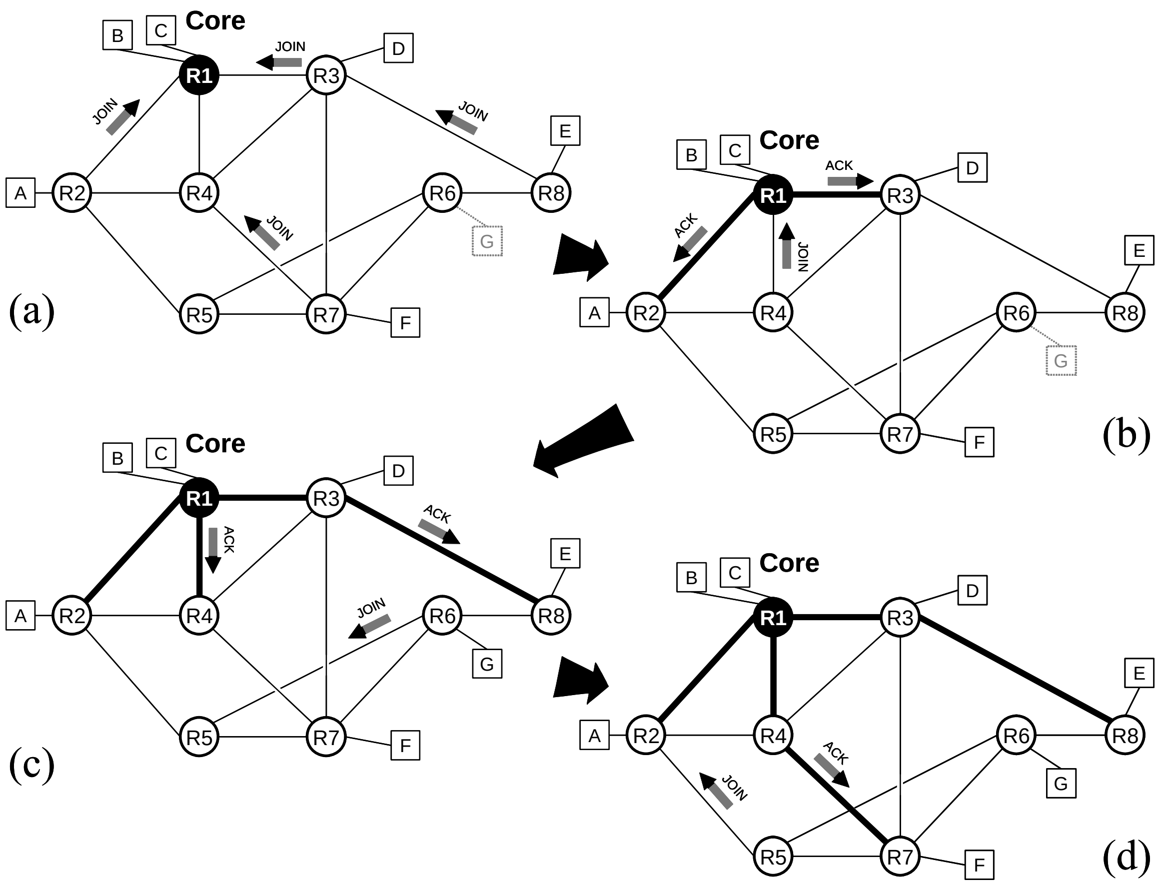

Figure 3-10: Forming the spanning tree by CBTs approach for the example in Figure 3-9. Thick lines represent the spanning tree, as its branches are grown.

Figure 3-10 illustrates the process for the network in Figure 3-9(a), assuming that R1 is configured as the core router. Suppose that the source *A* and receivers *B*, *C*, *D*, *E*, and *F* join simultaneously. (Receiver *G* will join later.) Figure 3-10(a) shows how each router that has attached a group member unicasts the join-request message to the next hop on the unicast path to the group's core. In Figure 3-10(b), the core R1 sends join-acknowledgement messages to R2 and R3, and R4 relays R7's join request. Notice that R3 does not forward the join request by R8. This is because R3 already sent its own join request. Subsequent join requests received for the same group are cached until this router has received a join acknowledgement for the previously sent join, at which time any cached joins can also be acknowledged. This happens in Figure 3-10(c), where after receiving a join acknowledgement, R3 in turn acknowledges R8's join. We assume that at this time receiver *G* decides to join the multicast group and R6 sends a join request on its behalf. There are three shortest paths from R6 to R1: paths R6-R5-R2-R1, R6-R7-R4-R1, and R6-R8-R3-R1; we assume that the tie was broken by the unicast routing algorithm and R6-R5-R2-R1 was selected. (The reader may notice that there were several other shortest-path ties, which again we assume were broken by the unicast algorithm.) Figure 3-10(d) shows that the branch from the core to R7 is established, and at the same time, the join request from R6 reaches R2. R2 will not propagate R6's join request because it is already on the spanning tree for the same group. Therefore, R2 will respond with a join acknowledgement, which will travel opposite the join request until it reaches R6 (not shown in Figure 3-10).

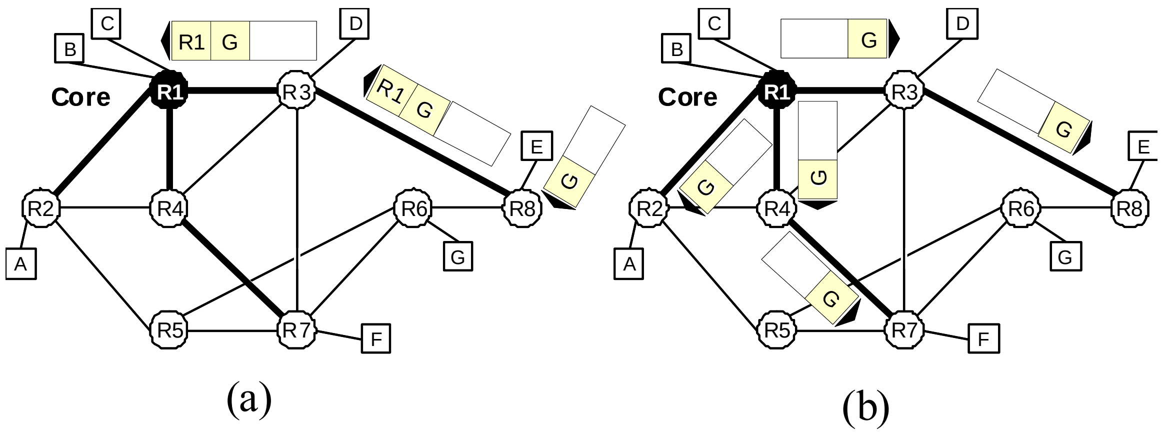

Figure 3-11: Packet forwarding from source *E* to the multicast group *G* in Figure 3-10.

The resulting spanning tree is known as a **group-shared multicast tree** because any multicast source in the multicast group *G* can use this tree. All routers that are part of the spanning tree create a forwarding table entry for the shared tree, called $\langle *, G\rangle$ entry, where the wildcard $*$ stands for "any source" (within the group *G*). The outgoing network port for this entry is the network interface on which the Join message arrived during the spanning tree construction. All data packets that arrive for group *G* are forwarded out to this port. Each source first sends its traffic to the core router, which then multicasts it down the spanning tree. Consider the example in Figure 3-10 and assume that host *E* multicasts a message to the group. Host *E* constructs an IP packet and uses the group *G* IP address (Figure 1-45). *E* sends the packet to a router on its local network known as the *designated router*, in our case R8. R8 encapsulates the packet into a unicast IP packet and sends it to the core R1 (Figure 3-11). When the packet reaches R1, the core removes the unicast IP header and forwards it down the tree. As a pefromance optimization, packets destined for the group do not need to reach the core before they are multicast. As soon as a packet reaches the tree, it can be forwarded upstream toward the root, as well as downstream to all other branches.

If any router or a link goes down, the downstream router that used this router as the next hop towards the core will have to rejoin the spanning tree individually on behalf of each group present on their outgoing interfaces. Further, during reconfiguring a new router as the core a situation can occur where a leaf router finds that the next hop towards the new core is the router that is downstream to it relative to the prior core. Such a situation is depicted in Figure 3-12. Here, after reconfiguration, router R7 finds that in order to join the new core it has to send a join request towards R5, which is downstream to it (i.e., R7 is still the next-hop router for R5 toward the old core). To deal with this situation, R7 sends a "flush-tree" message downstream to teardown the old tree, i.e., to break the spanning-tree branch from R7 to R5. The downstream routers then perform explicit Rejoin if they have group members attached to them.

CBTs has several advantages over RPF's flood-and-prune approach when the multicast group is sparse (i.e., relatively few routers in the network have group members attached). First, routers that are not members of the multicast group will never know of its existence, so we avoid the overhead of flooding. Second, join and leave messages are explicit, so the hosts can join or leave without waiting for the next flooded packet. Third, each router needs to store only one record per

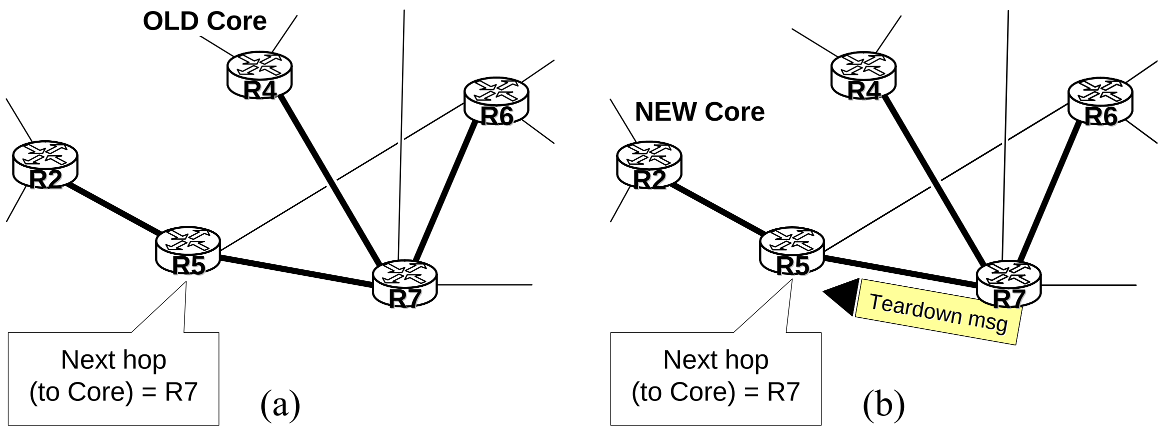

Figure 3-12: Reconfiguration of a CBTs core router from (a) to (b) requires the spanning-tree teardown and rebuilding a new spanning-tree.

group (the interfaces on which to forward packets for that group). It does not need to store per-source prune information or compute a shortest path tree explicitly.

However, CBTs has several issues of its own. All traffic for the group must pass through the core router, which can become a bottleneck. A shared spanning tree is not the most efficient solution for different sources, as discussed below. Additionally, there is a reliability issue: if the core router goes down, every multicast group that goes through it also goes down. Other issues include: unidirectional vs. bidirectional shared trees; core placement and selection; multiple cores; and, dynamic cores.

A shared spanning tree based on a core router is not optimal for all sources. For example in Figure 3-11, a packet from R8 will reach R7 in four hops (via the core R1), instead of two hops (via R6). In general, the path from a source to receiver via the core might be significantly longer than the shortest possible path. The degree of inefficiency depends on where the core and sources are located relative to each other. If the core is in the "middle," the inefficiency is reasonably small. A possible optimization is to build a **source-specific tree**. Instead of sending a wildcard Join message to join the group *G*, a receiver router sends a source-specific Join towards the source. As this message follows the shortest path towards the source *S*, the routers along the way create an $\langle S, G\rangle$ entry for this tree in their forwarding table. The resulting tree has the root at the source *S* rather than the core router, which may not be part of the new source-specific tree at all. However, the group-shared tree rooted in the core should remain untouched so that other nodes in the group *G* may become sources at a later point.

Core-based tree approach to building and maintaining spanning trees is implemented in the Internet multicast protocol called Protocol-Independent Multicast (PIM), in the variation called Sparse Mode (PIM-SM). See Section 8.2.4 for more information.

3.3.3 Peer-to-Peer Routing

Skype, etc.

3.3.4 Resource Reservation and Integrated Services

Integrated Services (IntServ) is an architecture that specifies the elements to guarantee quality-of-service (QoS) on data networks. IntServ requires that every router in the system implements IntServ, and every application that requires some kind of guarantees has to make an individual reservation. Flow specifications (**Flowspecs**) describe what the reservation is for, while RSVP is the underlying mechanism to signal it across the network.

There are two parts to a Flowspec:

(i) What does the traffic look like, specified in the **Traffic SPECification** or T_{spec} part.

(ii) What guarantees does it need, specified in the service **Request SPECification** or R_{spec} part.

T_{spec}s include token bucket algorithm parameters (Section 5.2). The idea is that there is a token bucket which slowly fills up with tokens, arriving at a constant rate. Every packet that is sent requires a token, and if there are no tokens, then it cannot be sent. Thus, the rate at which tokens arrive dictates the average rate of traffic flow, while the depth of the bucket dictates how “bursty” the traffic is allowed to be.

T_{spec}s typically just specify the token rate and the bucket depth. For example, a video with a refresh rate of 75 frames per second, with each frame taking 10 packets, might specify a token rate of 750Hz, and a bucket depth of only 10. The bucket depth would be sufficient to accommodate the “burst” associated with sending an entire frame all at once. On the other hand, a conversation would need a lower token rate, but a much higher bucket depth. This is because there are often pauses in conversations, so they can make do with fewer tokens by not sending the gaps between words and sentences. However, this means the bucket depth needs to be increased to compensate for the traffic being burstier.

R_{spec}s specify what requirements there are for the flow: it can be normal internet “best effort,” in which case no reservation is needed. This setting is likely to be used for webpages, FTP, and similar applications. The “controlled load” setting mirrors the performance of a lightly loaded network: there may be occasional glitches when two people access the same resource by chance, but generally both delay and drop rate are fairly constant at the desired rate. This setting is likely to be used by soft QoS applications. The “guaranteed” setting gives an absolutely bounded service, where the delay is promised to never go above a desired amount, and packets never dropped, provided the traffic stays within the specification.

Resource Reservation Protocol (RSVP)

The RSVP protocol (Resource ReSerVation Protocol) is a transport layer protocol designed to reserve resources across a network for an integrated services Internet. RSVP defines how applications place reservations for network resources and how they can relinquish the reserved resources once they are not need any more. It is used by a host to request specific qualities of service from the network for particular application data streams or flows. RSVP is also used by routers to deliver quality-of-service (QoS) requests to all nodes along the path(s) of the flows and to establish and maintain state to provide the requested service. RSVP requests will generally result in resources being reserved in each node along the data path.

RSVP is not used to transport application data but rather to control the network, similar to routing protocols. A host uses RSVP to request a specific QoS from the network, on behalf of an application data stream. RSVP carries the request through the network, visiting each node the network uses to carry the stream. At each node, RSVP attempts to make a resource reservation for the stream.

To make a resource reservation at a node, the RSVP daemon communicates with two local decision modules, admission control and policy control. Admission control determines whether the node has sufficient available resources to supply the requested QoS. Policy control determines whether the user has administrative permission to make the reservation. If either check fails, the RSVP program returns an error notification to the application process that originated the request. If both checks succeed, the RSVP daemon sets parameters in a packet classifier and packet scheduler to obtain the desired QoS. The packet classifier determines the QoS class for each packet and the scheduler orders packet transmission to achieve the promised QoS for each stream.

The routers between the sender and listener have to decide if they can support the reservation being requested, and, if they cannot, they send a reject message to let the listener know about it. Otherwise, once they accept the reservation they have to carry the traffic.

The routers store the nature of the flow, and then police it. This is all done in soft state, so if nothing is heard for a certain length of time, then the reader will time out and the reservation will be cancelled. This solves the problem if either the sender or the receiver crash or are shut down incorrectly without first canceling the reservation. The individual routers have an option to police the traffic to ascertain that it conforms to the flowspecs.

Summary of the key aspects of the RSVP protocol:

1. Shortest-path multicast group/tree

 * Require a shortest-path multicast group/tree to have already been created.

 * Tree created by Dijkstra algorithm (Section 1.4.2) for link state routing protocols, or via reverse path broadcast procedure, for distance vector routing protocols.

2. PATH message

 * Source sends a PATH message to group members with T_{spec} info

 * T_{spec} = Description of traffic flow requirements

3. Router inspection of PATH message

 * Each router receiving the PATH message inspects it and determines the reverse path to the source.

 * Each router also may include a QoS advertisement, which is sent downstream so that the receiving hosts of the PATH message might be able to more intelligently construct, or dynamically adjust, their reservation request.

4. RESV message

 * Receiver sends RESV message "back up the tree" to the source.

 * RESV message contains the (T_{spec}, R_{spec}) info (the FlowSpec pair) and Filter spec.

 * R_{spec} = Description of service requested from the network (i.e., the receiver's requirements)

* Thus, the FlowSpec (T_{spec}, R_{spec}) specifies a desired QoS.

* The Filter spec, together with a session specification, defines the set of data packets—the "flow"—to receive the QoS defined by the FlowSpec.

5. Router inspection of RESV message

* Each router inspects the (T_{spec}, R_{spec}) requirements and determines if the desired QoS can be satisfied.

* If yes, the router forwards the RESV message to the next node up the multicast tree towards the source.

* If no, the router sends a rejection message back to the receiving host.

6. RSVP session

* If the RESV message makes its way up the multicast tree back to the source, the reservation flow request has been approved by all routers in the flow path, and transmission of the application data can begin.

* PATH/RESV messages are sent by source/receiver every 30 seconds to maintain the reservation.

* When a timeout occurs while routers await receipt of a RESV message, then the routers will free the network resources that had been reserved for the RSVP session.

RSVP runs over IP, both IPv4 and IPv6. Among RSVP's other features, it provides opaque transport of traffic control and policy control messages, and provides transparent operation through non-supporting regions.

Limitations of Integrated Services

IntServ specifies a fine-grained QoS system, which is often contrasted with DiffServ's coarse-grained control system (Section 3.3.5).

The problem with IntServ is that many states must be stored in each router. As a result, IntServ works on a small-scale, but as you scale up to a system the size of the Internet, it is difficult to keep track of all of the reservations. As a result, IntServ is not very popular.

One way to solve this problem is by using a multi-level approach, where per-microflow resource reservation (i.e., resource reservation for individual users) is done in the edge network, while in the core network resources are reserved for aggregate flows only. The routers that lie between these different levels must adjust the amount of aggregate bandwidth reserved from the core network so that the reservation requests for individual flows from the edge network can be better satisfied. See RFC 3175.

3.3.5 Traffic Classes and Differentiated Services

DiffServ (Differentiated Services) is an IETF model for QoS provisioning. There are different DiffServ proposals, and some simply divide traffic types into two classes. The rationale behind this approach is that, given the complexities of the best effort traffic, it makes sense to add new complexity in small increments.

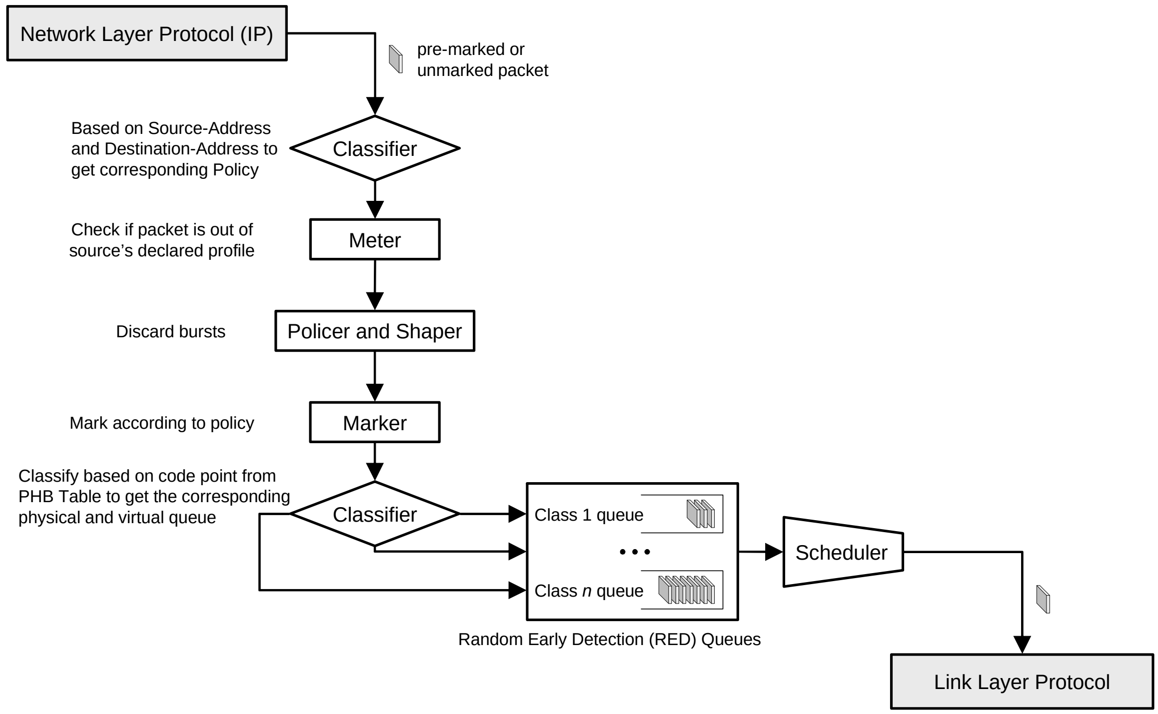

Figure 3-13: DiffServ architecture.

Suppose that we have enhanced the best-effort service model by adding just one new class, which we call "premium."

Assuming that packets have been marked in some way, we need to specify the router behavior on encountering a packet with different markings. This can be done in different ways and IETF is standardizing a set of router behaviors to be applied to marked packets. These are called "per-hop behaviors" (PHBs), a term indicating that they define the behavior of individual routers rather than end-to-end services.

DiffServ mechanisms (Figure 3-13):

* Lies between the network layer and the link layer

* Traffic *marked*, *metered*, *policed*, and *shaped* at source

* Packets *queued* for preferential forwarding, based on:

  - Delay bounds marking

  - Throughput guarantees marking

* Queue for each class of traffic, varying parameters

* Weighted *scheduling* preferentially forwards packets to link layer

DiffServ Traffic Classes

One PHB is “expedited forwarding” (EF), which states that the packets marked as EF should be forwarded by the router with minimal delay and loss. Of course, this is only possible if the arrival rate of EF packets at the router is always less than the rate at which the router can forward EF packets.

Another PHB is known as “assumed forwarding” (AF).

3.4 Adaptation Parameters

3.5 QoS in Wireless Networks

3.6 Summary and Bibliographical Notes

Latency, jitter and packet loss are the most common ills that plague real-time and multimedia systems. The remedy is in the form of various quality-of-service (QoS) provisions. Chapter 4 analyzes store-and-forward and queuing congestion in switches and routers. Congestion can lead to packets spacing unpredictably and thus resulting in jitter. The more hops a packet has to travel, the worse the jitter. Latency due to distance (propagation delay) is due to the underlying physics and nothing can be done to reduce propagation delay. However, devices that interconnect networks (routers) impose latency that is often highly variable. Jitter is primarily caused by these device-related latency variations. As a device becomes busier, packets must be queued. If those packets happen to be real-time audio data, jitter is introduced into the audio stream and audio quality declines.

Chapter 5 describes techniques for QoS provisioning.

The material presented in this chapter requires basic understanding of probability and random processes. [Yates & Goodman 2004] provides an excellent introduction and [Papoulis & Pillai 2001] is a more advanced and comprehensive text.

For video, expectations are low

For voice, ear is very sensitive to jitter and latencies, and loss/flicker

QoS: [Wang, 2001]

→ Multicast Routing

[Bertsekas & Gallagher, 1992] describe several algorithms for spanning-tree construction. Ballardie, *et al.*, [1993] introduced *core based trees* (CBT) algorithm for forming the delivery tree—the collection of nodes and links that a multicast packet traverses. Also see RFC-2189.

[Gärtner, 2003] reviews several distributed algorithms for computing the spanning tree of a network. He is particularly focusing on *self-stabilizing* algorithms that are guaranteed to recover from an arbitrary perturbation of their local state in a finite number of execution steps. This means that the variables of such algorithms do not need to be initialized properly.

→ IntServ

RSVP by itself is rarely deployed in data networks as of this writing (Fall 2009), but the traffic engineering extension of RSVP, called RSVP-TE [RFC 3209], is becoming accepted recently in many QoS-oriented networks.

As an important research topic: show that multihop can or cannot support multiple streams of voice.

RFC-2330 [Paxson, *et al.*, 1998] defines a general framework for performance metrics being developed by the IETF's IP Performance Metrics effort, by the IP Performance Metrics (IPPM) Working Group.

RFC-3393 [Demichelis & Chimento, 2002] defines one-way delay jitter across Internet paths.

RFC 2205: Resource ReSerVation Protocol (RSVP) -- Version 1 Functional

There was no notion of QoS in Ethernet until 1998 when IEEE published 802.1p as part of the 802.1D-1998 standard. 802.1p uses a three-bit field in the Ethernet frame header to denote an eight-level priority. One possible service-to-value mapping is suggested by RFC-2815, which describes Integrated Service (IntServ) mappings on IEEE 802 networks.

[Thomson, *et al.*, 1997]

Problems

Problem 3.1

Problem 3.2

Consider an internet telephony session, where both hosts use pulse code modulation to encode speech and sequence numbers to label their packets. Assume that the user at host *A* starts speaking at time zero, the host sends a packet every 20 ms, and the packets arrive at host *B* in the order shown in the table below. If *B* uses fixed playout delay of $q = 210$ ms, write down the playout times of the packets.

Packet sequence number	Arrival time r_i [ms]	Playout time p_i [ms]
#1	195	
#2	245	
#3	270	
#4	295	
#6	300	
#5	310	
#7	340	
#8	380	
#9	385	
#10	405	

Problem 3.3

Consider an internet telephony session over a network where the observed propagation delays vary between 50–200 ms. Assume that the session starts at time zero and both hosts use pulse code modulation to encode speech, where voice packets of 160 bytes are sent every 20 ms. Also, both hosts use a fixed playout delay of $q = 150$ ms.

(a) Write down the playout times of the packets received at one of the hosts as shown in the table below.

(b) What size of memory buffer is required at the destination to hold the packets for which the playout is delayed?

Packet sequence number	Arrival time r_i [ms]	Playout time p_i [ms]
#1	95	
#2	145	
#3	170	
#4	135	
#6	160	
#5	275	
#7	280	

#8	220	
#9	285	
#10	305	

Problem 3.4

Consider the same internet telephony session as in Problem 3.2, but this time the hosts use adaptive playout delay. Assume that the first packet of a new talk spurt is labeled *k*, the current estimate of the average delay is $\hat{\delta}_k$ = 90 ms, the average deviation of the delay is $\hat{\upsilon}_k$ = 15 ms, and the constants α = 0.01 and K = 4.

The table below shows how the packets are received at the host *B*. Write down their playout times, keeping in mind that the receiver must detect the start of a new talk spurt.

Packet seq. #	Timestamp t_i [ms]	Arrival time r_i [ms]	Playout time p_i [ms]	Average delay $\hat{\delta}_i$ [ms]	Average deviation $\hat{\upsilon}_i$
k	400	480			
k+1	420	510			
k+2	440	570			
k+3	460	600			
k+4	480	605			
k+7	540	645			
k+6	520	650			
k+8	560	680			
k+9	580	690			
k+10	620	695			
k+11	640	705			

Problem 3.5

Consider an internet telephony session using adaptive playout delay, where voice packets of 160 bytes are sent every 20 ms. Consider one of the receivers during the conversation. Assume that at the end of a previous talk spurt, the current estimate for the average delay is $\hat{\delta}_k$ = 150 ms, the average deviation of the delay is $\hat{\upsilon}_k$ = 50 ms. Because of high delay and jitter, using the constant K = 4 would produce noticeable playout delays affecting the perceived quality of conversation. If the receiver decides to maintain the playout delay at 300 ms, what will be the percentage of packets with missed playouts (approximate)? Explain your answer.

(*Hint*: Use the chart shown in the figure below to derive the answer.)

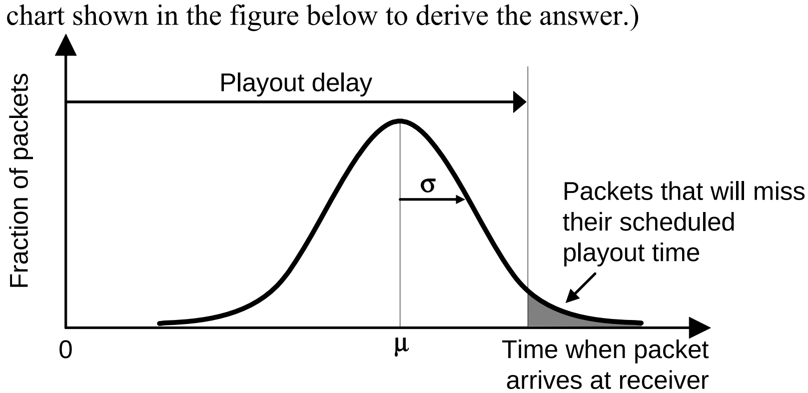

Problem 3.6

Consider the Internet telephony (VoIP) conferencing session shown in the figure. Assume that each audio stream is using PCM encoding at 64 Kbps. The packet size is 200 bytes.

Determine the periods for transmitting RTCP packets for all senders and receivers in the session.

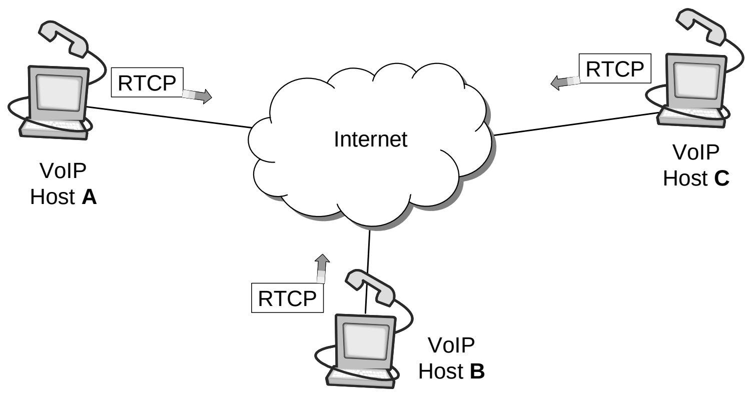

Problem 3.7

Consider the following network where router *A* needs to multicast a packet to all other routers in the network. Assume that the cost of each link is 1 and that the routers are using the reverse path forwarding (RPF) algorithm.

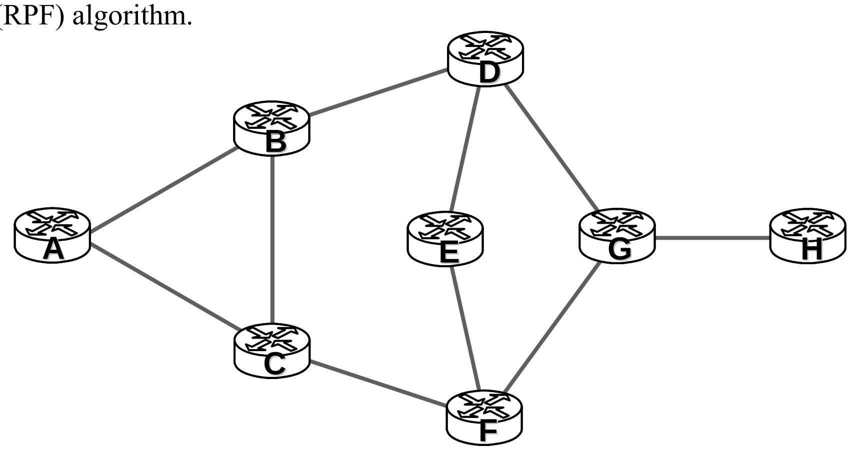

Do the following:

(e) Draw the shortest path multicast tree for the network.

(f) How many packets are forwarded in the entire network per every packet sent by the source *A*?

(g) Assuming that RPF uses pruning and routers *E* and *F* do not have attached hosts that are members of the multicast group, how many packets are forwarded in the entire network per every packet sent by *A*?

For each item, show the work, not only the final result.

Problem 3.8

Chapter 4
Switching and Queuing Delay Models

Contents

A key problem that switches and routers must deal with are the finite physical resources. Network nodes and links are never provisioned to support the maximum traffic rates because that is not economical. We all know that highway networks are not designed to support the maximum possible vehicular traffic. Because of economic reasons, highway networks are built to support average traffic rates. Traffic congestion normally occurs during the rush hours and later subsides. If congestion lasts too long, it can cause a major disruption of travel in the area. Although this is very unpleasant for the travelers and can cause economic loss, it is simply too expensive to provision road networks to avoid such situations altogether. Similar philosophy guides the design of communication networks.

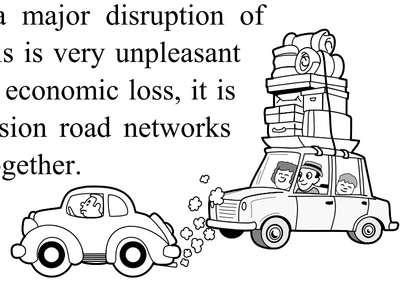

Data packets need two types of services on their way from the source to the final destination:

- *Computation* (or *processing*), which involves adding guidance information (or headers) to packets and looking up this information to deliver the packet to its correct destination
- *Communication* (or *transmission*) of packets over communication links

Figure 4-1 compares the total time to delay a packet if the source and destination are connected with a direct link vs. total time when an intermediate router relays packets. As seen, the router introduces both processing and transmission delays.

Both services are offered by physical servers (processing units and communication links) which have limited servicing capacity. When a packet arrives to a server that is already servicing packets that arrived previously, then we have a problem of *contention* for the service. The new packet is placed into a *waiting line* (or *queue*) to wait for its turn. The delay experienced while

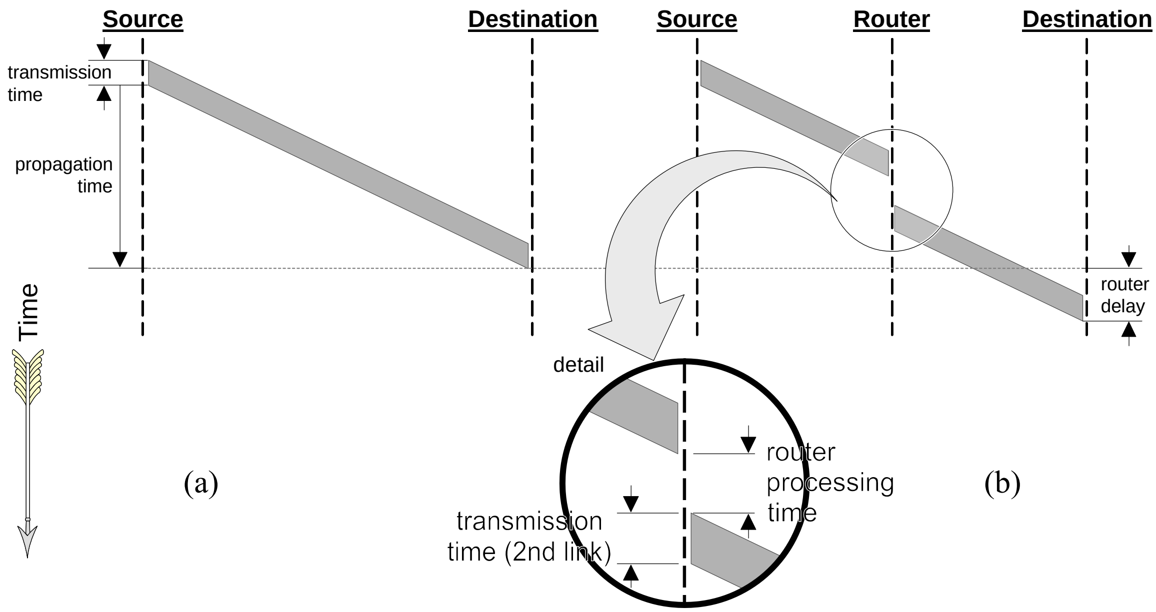

Figure 4-1: Delays in datagram packet switching (or, forwarding, or, routing). (a) Single hop source-to-destination connection without intermediary nodes. (b) Intermediary nodes introduce additionally processing and transmission delays.

waiting in line before being serviced is part of the total delay experienced by packets when traveling from source to destination. Generally, packets in a networked system experience these types of delays:

processing + transmission + propagation + queuing

The first three types of delays are described in Section 1.3. This chapter deals with the last kind of delay, **queuing delays**. Queuing models are used as a prediction tool to estimate this waiting time:

- Computation (queuing delays while waiting for processing)
- Communication (queuing delays while waiting for transmission)

This chapter studies what contributes to routing delays and presents simple analytical models of router queuing delays. Chapter 5 describes techniques to reduce routing delays or redistribute them in a desired manner over different types of data traffic.

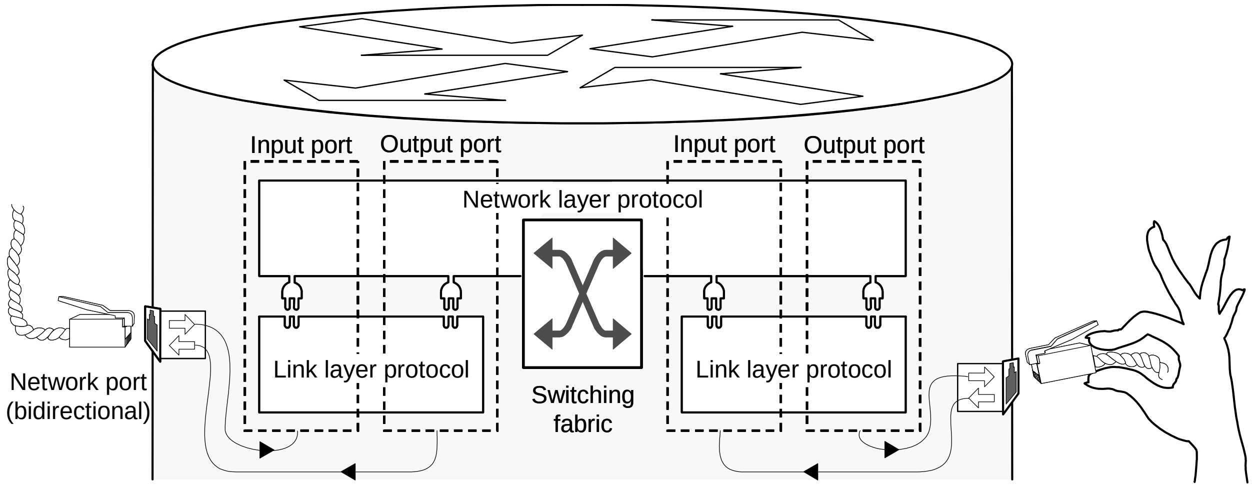

Figure 4-2: Hardware and software components of a router.

4.1 Packet Switching in Routers

Section 1.4 introduced routers and mainly focused on the control functions of routers. This section describes the datapath functions of routers. A key problem to deal with is the resource limitation. If the router is too slow to move the incoming packets, the packets will experience large delays and may need to be discarded if the router runs out of the memory space (known as buffering space). When packets are discarded too frequently, the router is said to be *congested*. The ability of a router to handle successfully the resource contention is a key aspect of its performance.

Figure 4-2 illustrates key hardware and software components of a router. Section 4.1.1 describes how these components function to forward data packets. Although network ports are bidirectional, it is useful to logically separate input and output ports. Router implements only the bottom two layers of the software protocol stack: link and network layer. Each network port has an associated link-layer protocol, but the network-layer protocol is common for all ports. This property will be explained later with Figure 4-4.

4.1.1 How Routers Forward Packets

Routers have two main functions:

1. **Forwarding or switching packets** (*datapath functions*) that pass through the router. One could think of these functions as using maps to direct the packets to their destinations. These operations are performed very frequently and are most often implemented in special purpose hardware.
2. **Maintaining routing tables** (*control functions*) by exchanging network connectivity information with neighboring routers, as well as system configuration and management. One could think of these functions as surveying and cartography to build the maps that

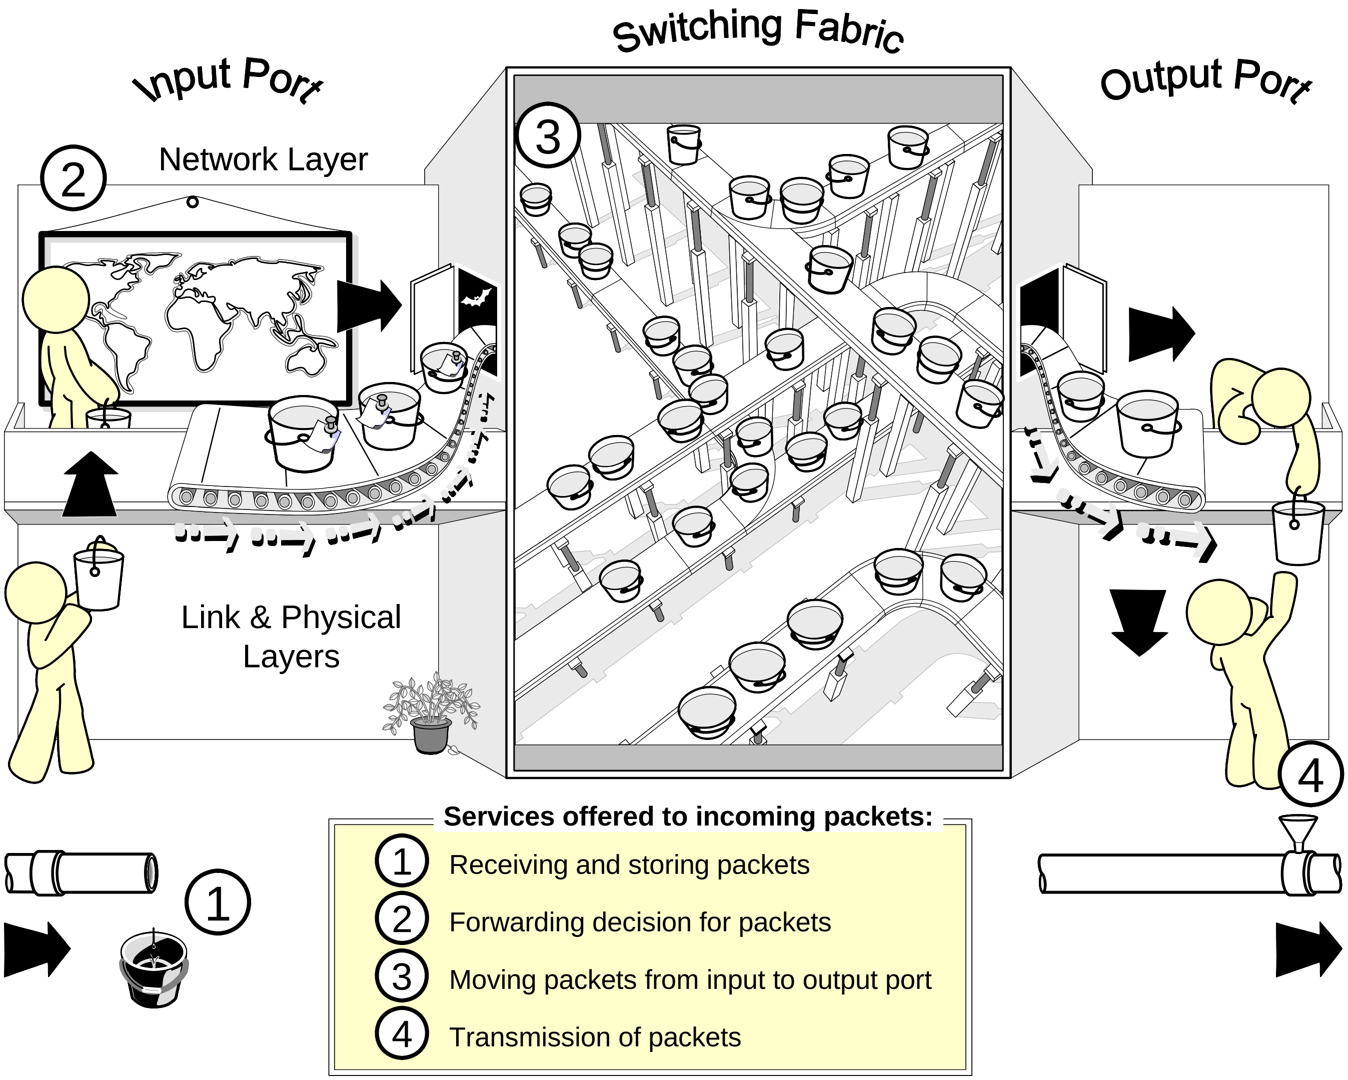

Figure 4-3: How router forwards packets: illustrated are four key services offered to incoming data packets.

are used for packet forwarding. These operations are performed relatively infrequently and are invariably implemented in software.

When trying to improve the per-packet performance of a router, we focus on the datapath functions because they must be fast. Routing table maintenance is described in Section 1.4. This chapter focuses on the datapath functions of packet forwarding (or, switching).

The datapath architecture consists of three main units: (1) *input ports* where incoming packets are received; (2) *switching fabric* that transfers packets from input to output ports; and, (3) *output ports* that transmit packets to an outgoing communication link. Routers offer four key functions to incoming data packets (illustrated in Figure 4-3):

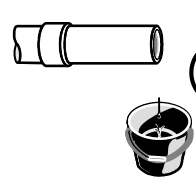

1. A packet is **received and stored** in the local memory on the input port at which the packet arrived.
2. The packet guidance information (destination address, stored in the packet

header) is looked up and the **forwarding decision** is made as to which output port the packet should be transferred to.

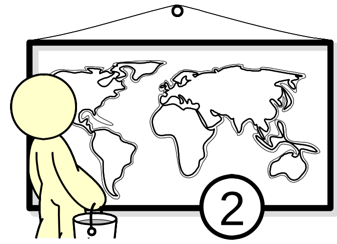

3. The packet is **transferred** across the switching fabric (also known as the *backplane*) to the appropriate output port, as decided in the preceding step.

4. The packet is **transmitted** on the output port over the outgoing communication link.

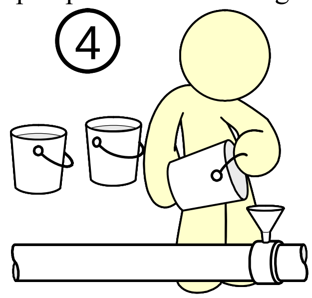

4.1.2 Router Architecture

Figure 1-12 shows protocol layering in end systems and intermediate nodes (switches or routers). Figure 4-4 shows the same layering from a network perspective. Each router runs an independent layer-1 (Link layer) protocol for each communication line. A single network layer (layer-2) is common for all link layers of the router.

The key architectural question in router design is about the implementation of the (shared) network layer of the router's protocol stack. The network layer binds together the link layers and performs packets switching. Link layers are terminating different communication links at the router and they essentially provide data input/output operations. They function independently of one another and are implemented using separate hardware units.

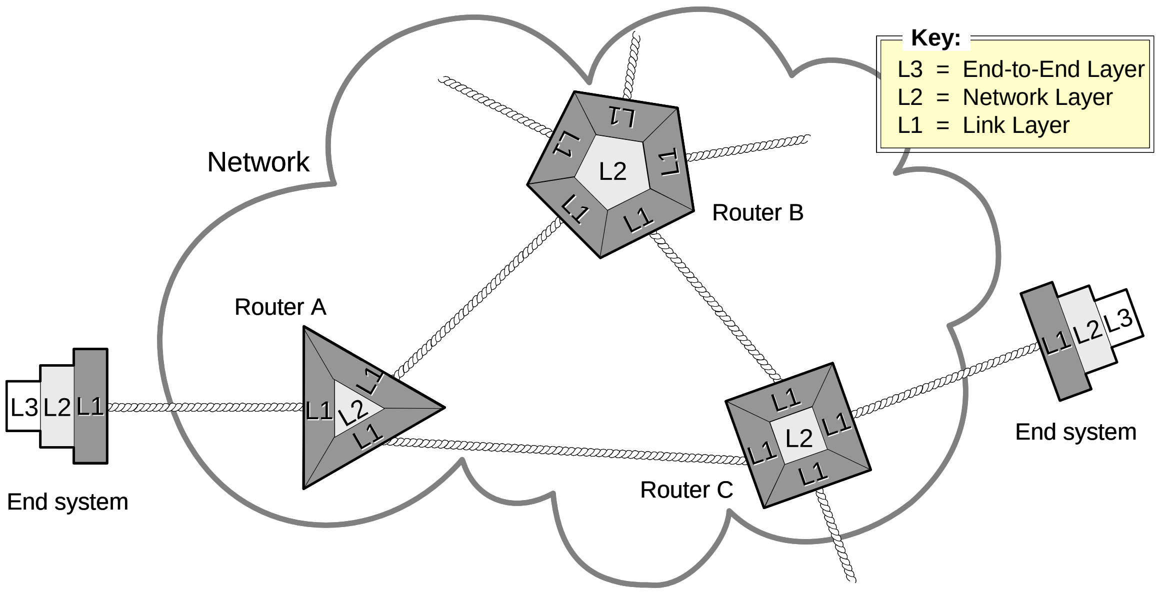

Figure 4-4: Distribution of protocol layers in routers and end systems.

The router's network layer must deal simultaneously with many parallel data streams. To achieve high performance, the network layer could be implemented in parallel hardware. The issues that make the router's network layer design of difficult include:

- *Maintaining consistent forwarding tables*: If the networking layer is distributed over parallel hardware units, each unit must maintain its own copy of the forwarding table. Because the forwarding table dynamically changes (as updated by the routing algorithm), all copies must be maintained consistent.
- *Achieving high-speed packet switching*: Once a packet's outgoing port is decided, the packet must be moved as quickly as possible from the incoming to the outgoing port.
- *Reducing queuing and blocking*: If two or more packets cross each other's way, the must be ordered in series, where they move one-by-one while others are waiting for their turn. The pathways inside the router should be designed to minimize chances for queuing and blocking to occur.

Over the years, different architectures have been used for routers. Particular architectures have been selected based on a number of factors, including cost, number of network ports, required performance, and currently available technology. The detailed implementations of individual commercial routers have generally remained proprietary, but in broad terms, all routers have evolved in similar ways. The evolution in the architecture of routers is illustrated in Figure 4-5.

First-Generation Routers: Central CPU Processor. The original routers were built around a conventional computer architecture, as shown in Figure 4-5(a): a shared central bus, with a central CPU, memory, and peripheral Line Cards (or, Network Interface Cards). Each Line Card performs the link-layer function, connecting the router to each of the communication links. The central CPU performs the network-layer function. Packets arriving from a link are transferred across the shared bus to the CPU, where a forwarding decision is made. The packet is then transferred across the bus again to its outgoing Line Card, and onto the communication link. Figure 4-6 highlights the datapath of first-generation routers.

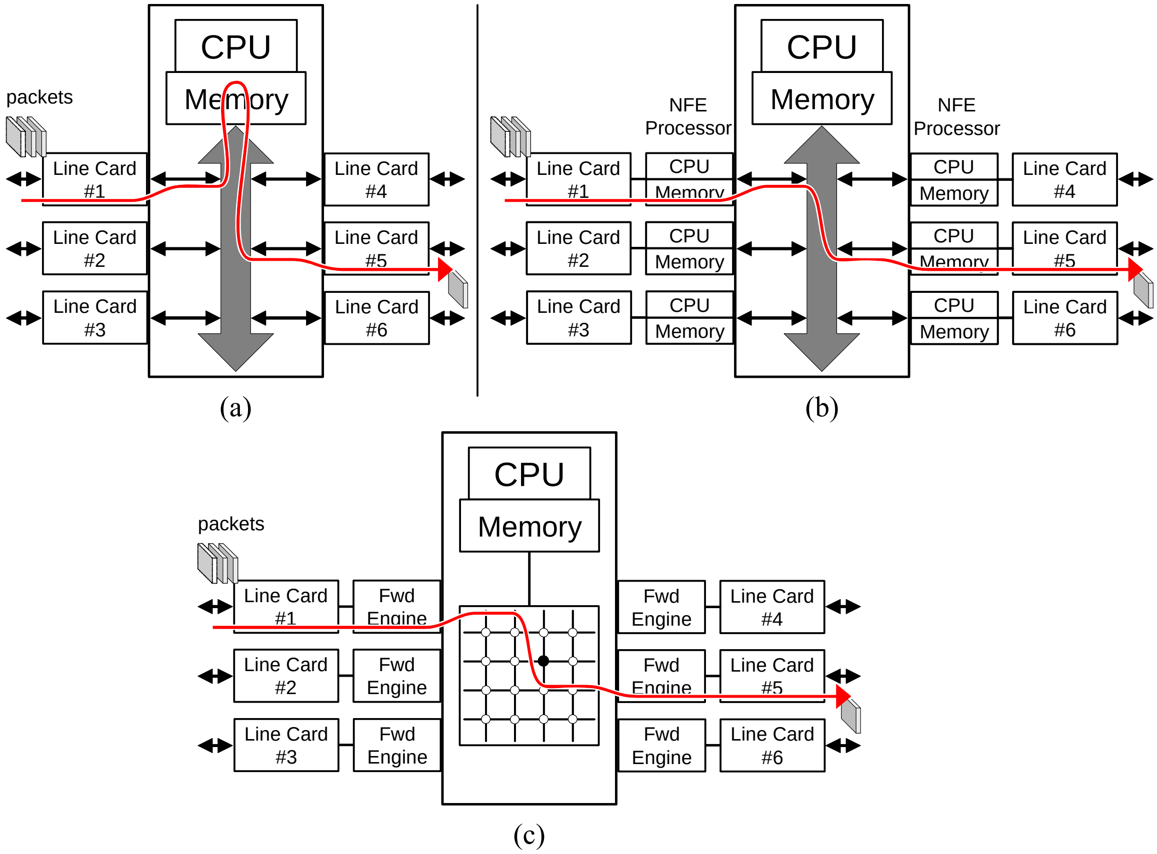

Figure 4-5: The basic architectures of packet-switching processors. (a) Central CPU processor; (b) Parallel network-front-end processors; (c) Switching fabric. The curved line indicates the packet path, from input to output port.

Second-Generation Routers: Network Front-end (NFE) Processors. The main limitation of the architecture in Figure 4-5(a) is that the central CPU must process every packet, ultimately limiting the throughput of the system. To increase the system throughput, the architecture in Figure 4-5(b) implements parallelism by placing a separate CPU at each interface. That is, the link-layer is still implemented in individual Line Cards, but the network-layer function is distributed across several dedicated CPUs, known as *network front-end* (NFE) *processor*s. A local forwarding decision is made in a NFE processor, and the packet is immediately forwarded to its outgoing interface. The central CPU is needed to run the routing algorithm and for centralized system management functions. It also computes the forwarding table and distributes it to the NFE processors.

The architecture in Figure 4-5(b) has higher performance than a first-generation design because the network-layer function is distributed over several NFE processors that run in parallel; and because each packet need only traverse the bus once, thus increasing the system throughput. Figure 4-7 highlights the datapath of second-generation routers. However, the performance is still limited by two factors. First, forwarding decisions are made in software, and so are limited by the speed of the NFE processor, which is a general purpose CPU. But general purpose CPUs are not well suited to applications in which the data (packets) flow through the system; CPUs are better suited to applications in which data is examined multiple times, thus allowing the efficient use of

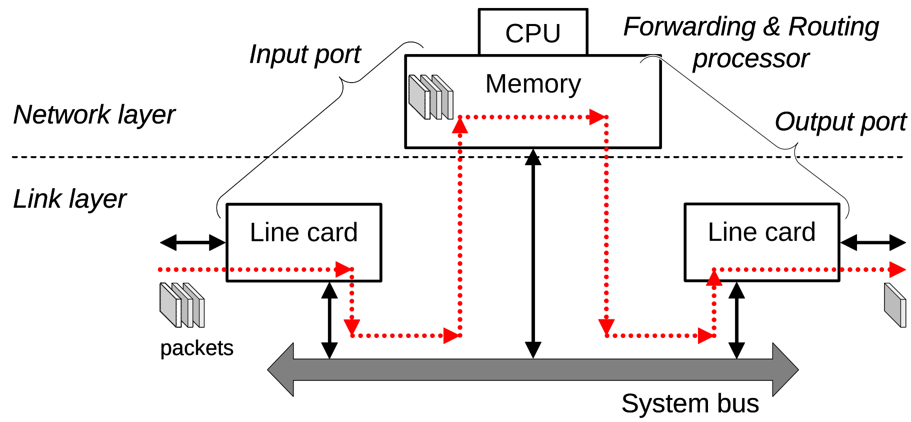

Figure 4-6: Packet datapath for switching via memory. Also shown in Figure 4-5(a).

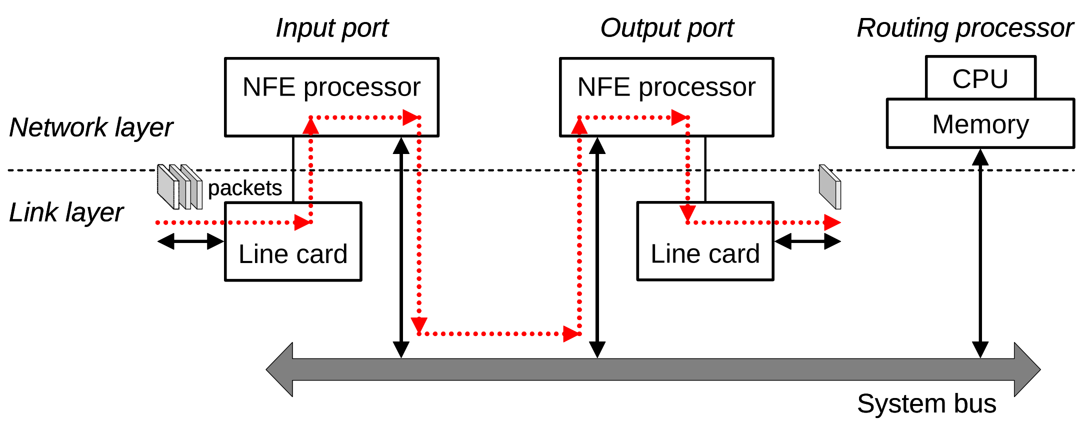

Figure 4-7: Packet datapath for switching via bus. Also shown in Figure 4-5(b).

a cache. Carefully designed, special purpose ASICs can readily outperform a CPU when making forwarding decisions, managing queues, and arbitrating access to the bus. Hence, CPUs are being replaced increasingly by specialized ASICs. The second factor that limits the performance is the use of a shared bus—only one packet may traverse the bus at a time between two Line Cards. Performance can be increased if multiple packets can be transferred across the bus simultaneously. This is the reason that a switch fabric is used in high-end routers.

Third-Generation Routers: Switching Fabric. By introducing a hardware-forwarding engine and replacing the bus with an interconnection network, we reach the architecture shown in Figure 4-5(c). In an interconnection network, multiple Line Cards can communicate with each other simultaneously greatly increasing the system throughput. Today, the highest performance routers are designed according to this architecture.

Input Ports

The key functions of input ports are to receive packets and make the forwarding decision. This spans both link and network layers of the protocol stack.

A network port is not the same as a Line Card. A Line Card supports the link-layer functionality of a network port, which is receiving and transmitting packets. A Line Card may also support the network-layer functionality, if the Network Front-End Processor is located on a Line Card as in

Visit http://en.wikipedia.org/wiki/ASIC for information about Application-Specific Integrated Circuits

Routing function	Unicast routing	Unicast routing with Types of Service	Multicast routing
Forwarding algorithm	Longest prefix match on destination address	Longest prefix match on destination address **+** exact match on Type of Service	Longest match on source address **+** exact match on source address, destination address, and incoming interface

Figure 4-8: Forwarding algorithms for different routing functions.

second-generation routers. However, it may not have any of the network-layer functionality, which is the case in first-generation routers where all network-layer functionality is supported by the central processor and memory.

Output Ports

The key function of output ports is to transmit packets on the outgoing communication links. If packets are arriving at a greater rate than the output port is able to transmit, some packets will be enqueued into waiting lines. The output port may also need to manage how different types of packets are lined up for transmission. This is known as *scheduling* and several scheduling techniques are described in Chapter 5. As with input ports, these functions span both link and network layers of the protocol stack.

4.1.3 Forwarding Table Lookup

We know from Section 1.4.4 that routers use *destination address prefixes* to identify a contiguous range of IP addresses in their routing messages. A destination prefix is a group of IP addresses that may be treated similarly for packet forwarding purposes. Based on its routing table, the router derives its forwarding table, also known as FIB (Forwarding Information Base), and uses it for making forwarding decisions for data packets. Each entry in a forwarding table/FIB represents a mapping from an IP address prefix (a range of addresses) to an outgoing link, with the property that packets from any destination with that prefix may be sent along the corresponding link. Forwarding table entries are called **routes**.

The algorithm used by the forwarding component of a router to make a forwarding decision on a packet uses two sources of information: (1) the forwarding table or FIB, and (2) the packet header. Although IP addresses are always the same length, IP prefixes are of variable length. The IP destination lookup algorithm needs to find the *longest prefix match*—the longest prefix in the FIB that matches the high-order bits in the IP address of the packet being forwarded. Longest prefix match used to be computationally expensive. The advances that have been made in longest-match algorithms in recent years have solved the problem of matching.

Packet forwarding decision depends on several parameters, depending on the routing function that needs to be supported (Figure 4-8), such as unicast routing, multicast routing, or unicast routing with Types of Service. Therefore, in addition to the information that controls where a packet is forwarded (next hop), an entry in the forwarding table may include the information about what resources the packet may use, such as a particular outgoing queue that the packet should be placed on (known as *packet classification*, to be described later). Forwarding of unicast packets requires longest prefix match based on the network-layer destination address. Unicast forwarding with Types of Service requires the longest match on the destination network-layer address, plus the exact match (fixed-length match) on the Type of Service (TOS) bits carried in the network-layer header (Figure 1-36). Forwarding of multicast packets requires longest match on the source network-layer address, plus the exact match (fixed-length match) on both source and destination addresses, where the destination address is the multicast group address.

For the purposes of multicast forwarding, some entries in the forwarding table/FIB may have multiple subentries. In multicast, a packet that arrives on one network interface needs to be sent out on multiple outgoing interfaces that are identified in subentries of a FIB record.

4.1.4 Switching Fabric Design

Problems related to this section: ?? → Problem 4.7

Switch Design Issues:

• Switch contention occurs when several packets are crossing each other's path – switch cannot support arbitrary set of transfers;

• Complex rearranging of the timetable for packet servicing (known as *scheduling*) is needed to avoid switch contention;

• High clock/transfer rate needed for bus-based design (first- and second-generation routers);

• Packet queuing (or, buffering) to avoid packet loss is needed when the component that provides service (generally known as "server") is busy;

Example switch fabrics include:

• Bus (first- and second-generation routers)

• Crossbar

• Banyan network

Banyan networks and other interconnection networks were initially developed to connect processors in a multiprocessor. They typically provide lower capacity than a complete crossbar.

Switching fabric may introduce different types of packet blocking. For example, if two or more packets at different inputs want to cross the switching fabric simultaneously towards the same output, then these packets experience *output blocking*. When one packet is heading for an idle port, but in front of it (in the same waiting line/queue) is another packet headed for a different output port that is currently busy, and the former packet must wait until the latter departs, then the former packet experiences *head-of-line blocking*. Find more information about packet blocking in Section 4.1.5.

Crossbar

The simplest switch fabric is a crossbar, which is a matrix of pathways that can be configured to connect any input port to any output port. An $N \times N$ crossbar has N input buses, N output buses, and N^2 crosspoints, which are either ON or OFF. If the (i, j) crosspoint is on, the i^{th} input port is connected to the j^{th} output port.

A crossbar needs a switching timetable, known as *switching schedule*, that tells it which inputs to connect to which outputs at a given time. If packets arrive at fixed intervals then the schedule can be computed in advance. However, in the general case, the switch has to compute the schedule while it is operating.

If packets from all N inputs are all heading towards different outputs then crossbar is N times faster than a bus-based (second-generation) switch. However, if two or more packets at different inputs want to go to the same output, then crossbar suffers from "output blocking" and as a result, it is not fully used. In the worst-case scenario, each output port must be able to accept packets from all input ports at once. To avoid output blocking, each output port would need to have a memory bandwidth equal to the total switch throughput, i.e., $N \times$ input port datarate. In reality, sophisticated designs are used to address this issue with lower memory bandwidths.

Banyan Network

Banyan network is a so-called *self-routing switch fabric*, because each switching element forwards the incoming packets based on the packet's tag that represents the output port to which this packets should go. The input port looks up the packet's outgoing port in the forwarding table based on packet's destination, and *tags* the packet with a binary representation of the output port. A Banyan switch fabric is organized in a hierarchy of switching elements. A switching element at level i checks the i^{th} bit of the tag; if the bit is 0, the packet is forwarded to the upper output, and otherwise to the lower output. Therefore, the tag can be considered a self-routing header of the packet, for routing the packet inside the switching fabric. The tag is removed at the output port before the packet leaves the router.

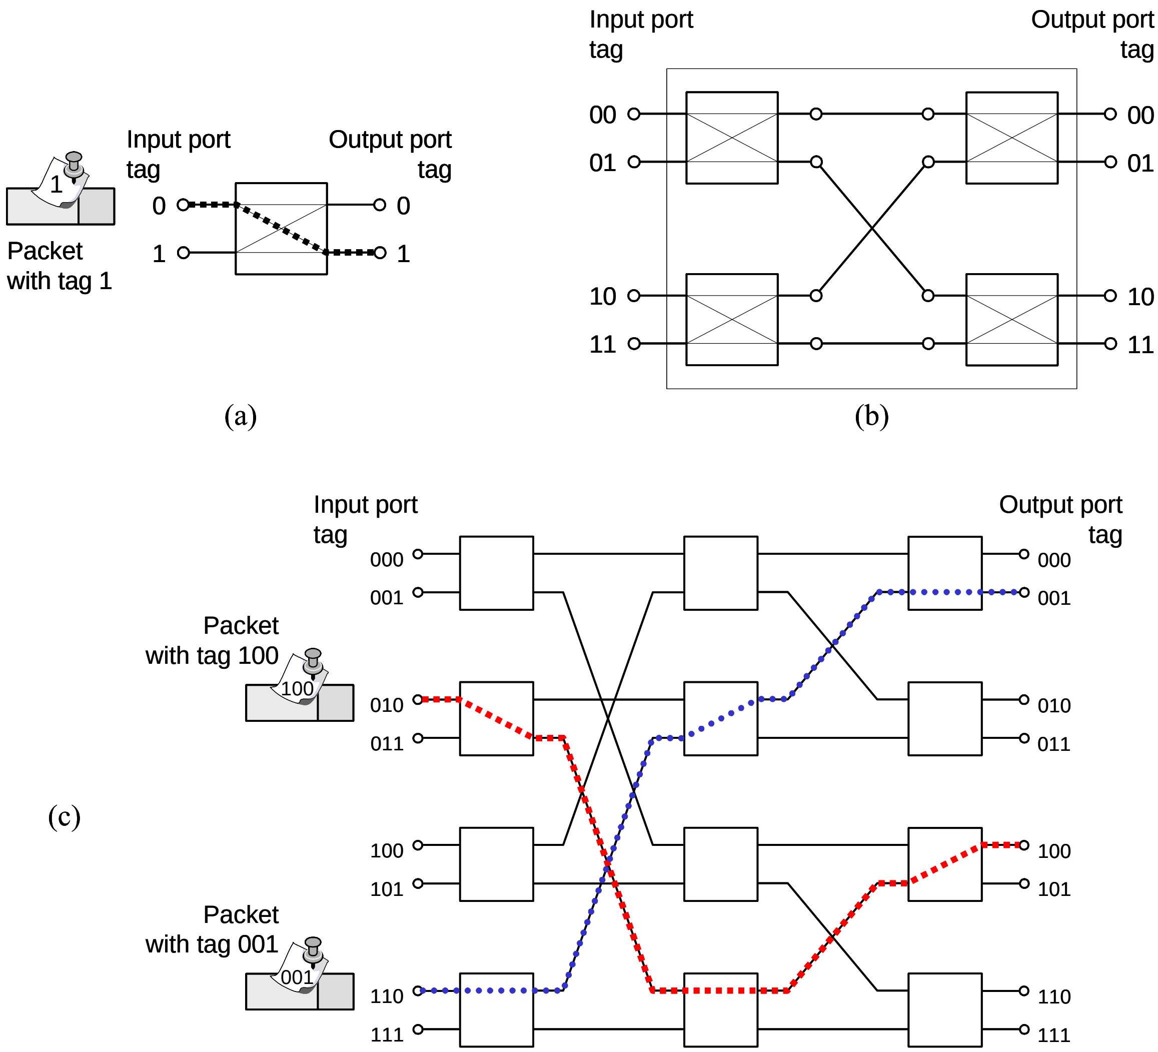

Figure 4-9: Banyan switch fabric. (a) A Banyan switching element is a 2×2 switch that moves packets either to output port 0 (upper port) or output port 1 (lower port), depending on the packet's tag. (b) A 4×4 Banyan network is composed from four 2×2 switching elements. (c) An 8×8 Banyan network is composed from twelve 2×2 switching elements.

The building block of a Banyan network is a 2×2 switch, i.e., a switch with two inputs and two outputs. The upper input and output ports are labeled with 0 and the lower input and output ports are labeled with 1. This switch moves packets based on a single-bit tag. For example, in Figure 4-9(a) a packet labeled with tag "1" arrives at input port 0 of a 2×2 switch. The switch directs the packet to the lower output port (the output port 1). To create a 4×4 switch, we need four 2×2 switching elements placed in a grid as in Figure 4-9(b). First we take two 2×2 switches and label them 0 and 1. Then we take another pair of 2×2 switches and place them before the first two. When a packet enters a switching element in the first stage, it is sent to the 2×2 switch labeled 0 if the first bit of the packet's tag is 0. Otherwise, it is sent to the switch labeled 1. To create an 8×8 switch, we need two 4×4 switches and four 2×2 switching elements placed in a grid as in Figure 4-9(c). Again, we label the two 4×4 switches as 0 and 1. The four 2×2 switching elements are placed before the 4×4 switches, and they send the packets to the corresponding 4×4 switch based on the first bit of the packet's tag. (Note that there are several equivalent 8×8 Banyan switches, only one of which is shown in Figure 4-9(c).)

If two incoming packets on any 2×2 switching element want to go to the same output port of this element, they *collide* and *block* at this element. For example, if two packets with tags "000" and "010" arrived at the input ports "000" and "001" in Figure 4-9(c), then they will collide at the first stage (in the upper-left corner 2×2 switching element), because they both need to go to the same output of this switching element. The switching element discards both of the colliding packets. Because packet loss is not desirable, we need to either prevent collisions or deal with them when they happen. In either case, instead of loss, the packet experiences delay while waiting for its turn to cross the switching fabric.

One option is to deal with collisions when they happen. This option requires a memory buffer for storing packets within the switching element. One of the colliding packets is transferred to the requested direction, while the other is stored in the buffer and sent in the subsequent cycle. This design is called *internal queuing* in Section 4.1.5. In the worst-case of input-traffic pattern, the internal buffer size must be large enough to hold several colliding packets.

Another option is to deal with collisions is to prevent them from happening. One way of preventing collisions is to *check* whether a path is available before sending a packet from an input port.

An alternative way of preventing collisions is by choosing the order in which packets appear at the input of the switching fabric. Obviously, the router cannot choose the input port at which a particular packet will arrive—packets arrive along the links depending on the upstream nodes that transmitted them. What can be done is to insert an additional network (known as *sorting network*) before a Banyan network, which rearranges the packets so that they are presented to the Banyan network in the order that avoids collisions. This is what a Batcher network does.

Batcher-Banyan Network

A Batcher network is a hardware network that takes a list of numbers and sorts them in the ascending order. Again, we assume that packets are tagged at the input ports at which they arrive. A tag is a binary representation of the output port to which the packet needs to be moved. A Batcher network sorts the packets that are currently at input ports into a non-decreasing order of their tags.

Figure 4-10 shows several Batcher sorting networks. To get an intuition why this network will correctly sort the inputs, consider Figure 4-10(b). The first four comparators will "sink" the largest value to the bottom and "lift" the smallest value to the top. The final comparator simply sorts out the middle two values.

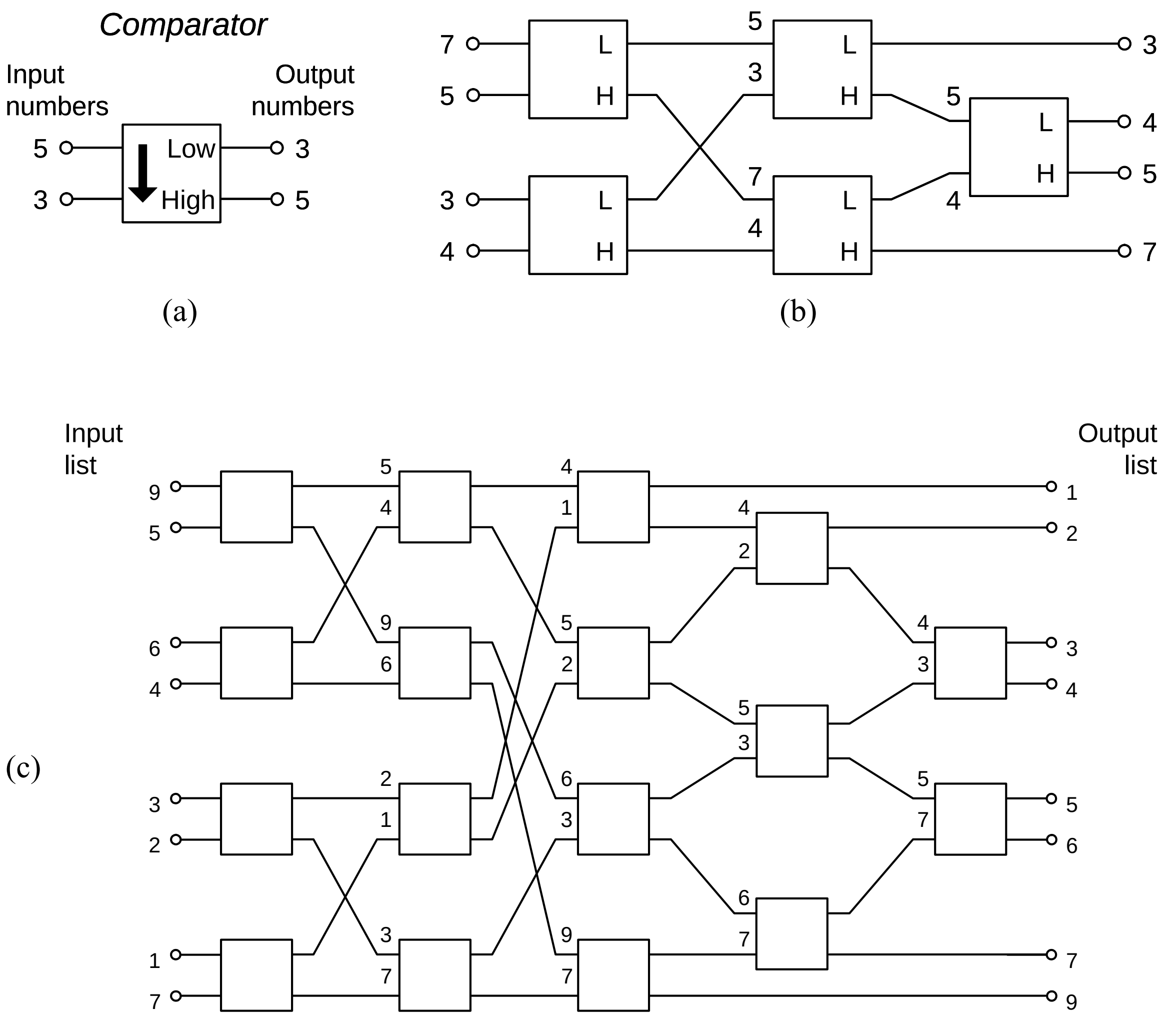

Figure 4-10: Batcher network. (a) A comparator element output the smaller input value at the upper output and the greater input value at the lower output. (b) A 4×4 Batcher network is composed five comparators. (c) An 8×8 Batcher network is composed from twelve comparators.

A combined Batcher-Banyan network is collision-free only when there are no packets heading for the same output port. Because there may be duplicates (tags with the same output port number) or gaps in the sequence, an additional network or special control sequence is required. Figure 4-11 shows a trap network and a shuffle-exchange network (or, concentrator) which serve to remove duplicates and gaps. In order to eliminate the packets for the same output, a *trap network* is required at the output of the Batcher sorter. Because packets are sorted by the Batcher sorter, the packets for the same output can be checked by comparison with neighboring packets. The trap network performs this comparison, and selects only one packet per output. Duplicates are trapped and dealt with separately.

One way to deal with the trapped duplicates is to store them and recirculate back to the entrance of the Batcher network, so in the next cycle they can compete with incoming packets. This requires that at least half the Batcher's inputs be reserved for recirculated packets (to account for the worst-case scenario).

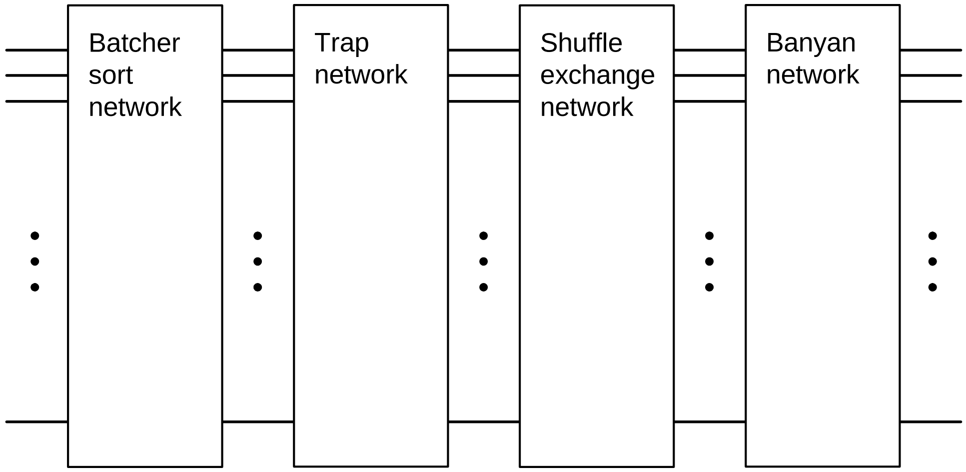

Figure 4-11: Batcher-Banyan network. Internal blocking is avoided by sorting the inputs to a Banyan network. A trap network and a shuffle-exchange network remove duplicates and gaps.

An alternative is to take the duplicates through multiple Banyan networks, where each packet that wants to go to the same output port is presented to a separate Banyan network.

Even when duplicates are removed, unselected and eliminated packets (gaps in the input sequence) generate empty inputs for the Banyan network. These gaps cause collisions in the Banyan network even if all packets are heading to different outputs. For example, consider an 8×8 Batcher-Banyan network with four packets heading to outputs 0, 0, 0, and 1, respectively. Although two of the three packets for the output 0 are trapped, the remaining two packets still collide in the second stage of the Banyan. To solve this problem, a *shuffle-exchange network* (or, *concentrator*) is required. In order to eliminate empty inputs, packets are shifted and the conflict is avoided. Although the role of the concentrator is just shifting the packets and eliminating empty inputs, the implementation is difficult because the number of packets that will be trapped cannot be predicted. Usually, special control sequence is introduced, or a Batcher sorter is used again as the concentrator.

4.1.5 Where and Why Queuing Happens

Problems related to this section: Problem 4.9 → ??

Queuing happens when customers are arriving at a rate that is higher than the server is able to service. In routers, waiting lines (queues) may be formed for any of the three services shown in illustrated in Figure 4-3, except for receiving packets at the input port. Packet queuing in routers is also known as **switch buffering**. The router needs memory allow for buffering, i.e., storing packets while they are waiting for service. By adopting different designs, the router architect can control where the buffering will occur.

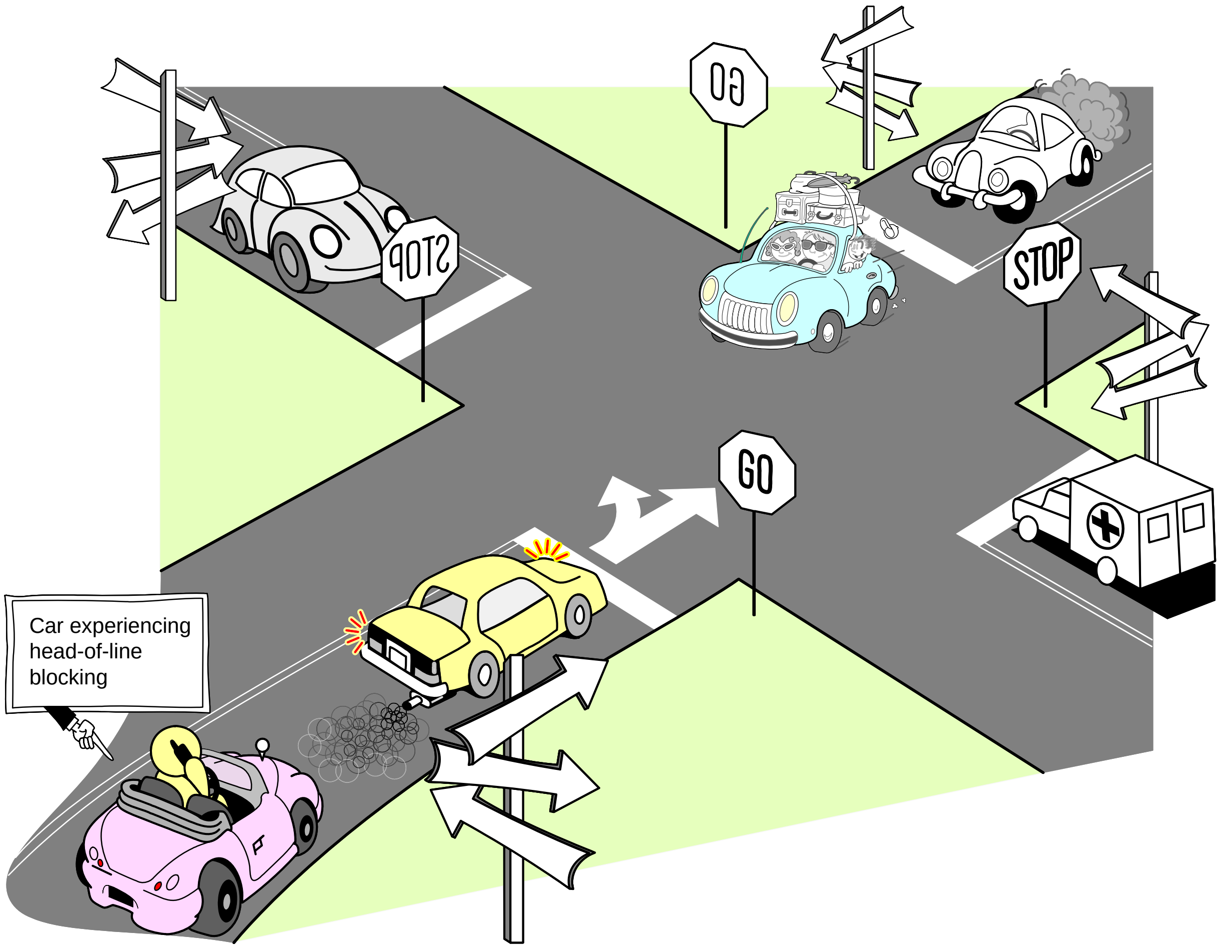

Figure 4-12: Illustration of forwarding issues by a road intersection analogy.

Before considering how queuing occurs in packet switches, let us consider the analogy with vehicular traffic, illustrated in Figure 4-12. The car in the lower left corner could go if it were not for the car in front of it that wishes to make the left turn but cannot because of the cars arriving in the parallel lane from the opposite direction. We say that the lower-left-corner car experiences **head-of-line** (HOL) **blocking**. A queue of cars will be formed at an entrance to the road intersection (or, "input port") because the front car cannot cross the intersection area (or, "switching fabric").

Another cause for queuing occurs the access to the intersection area (or, "switching fabric") has to be serialized. In Figure 4-12, the cars crossing from upper left corner to lower right corner and vice versa must wait for their turn because the intersection area is busy. Notice that we need an "arbiter" to serialize the access to the intersection area, and STOP/GO signals in Figure 4-12 serve this purpose. The corresponding queuing occurs in a router when the switching fabric has insufficient capacity to support incoming packet traffic. The queuing occurs at the input port and it is known as *input queuing* (or, *input buffering*) or in the switch fabric (*internal queuing*).

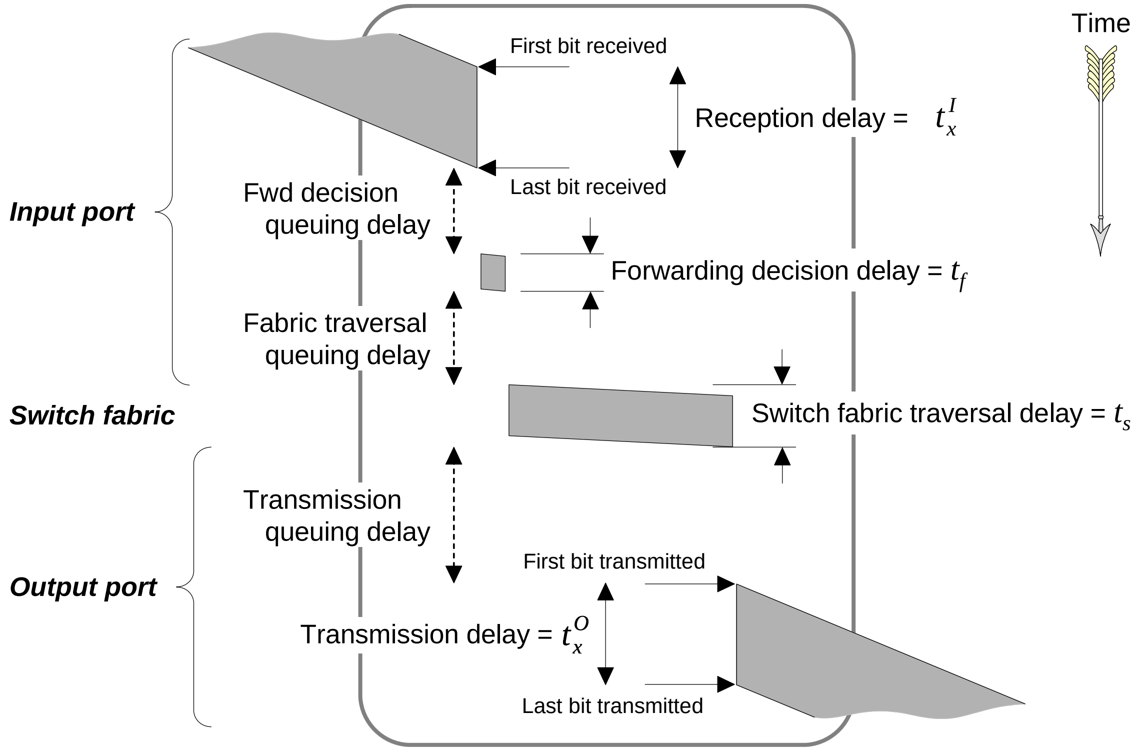

Figure 4-13: Components of delay in data packet forwarding.

Yet another reason for queuing occurs when an outgoing road cannot support the incoming traffic and the corresponding intersection exit becomes congested. In this case, even if an incoming car does not experience head-of-line blocking and it has a GO signal, it still bust wait because the exit is congested. The corresponding queuing occurs in a router when the outgoing communication line has insufficient capacity to support incoming packet traffic. The queuing may occur at the input port (*input queuing*), in the switch fabric (*internal queuing*), or at the output port (*output queuing*).

Figure 4-13 summarizes the datapath delays in a router. Some of these delays may be negligible, depending on the switch design. As already noted, queuing may also occur in the switch fabric (*internal queuing*), which is not shown in Figure 4-13.

Input Queuing

In switches with pure *input queuing* (or, *input buffering*), packets are stored at the input ports and released when they win access to both the switching fabric and the output line. An arbiter decides the timetable for accessing the fabric depending on the status of the fabric and the output lines (Figure 4-14). Because the packets leaving the input queues are guaranteed access to the fabric and the output line, there is no need for an output queue.

The key advantage of input queuing is that links in the switching fabric (and the input queues themselves) need to run at the speed of the input communication lines. For a router with N input ports and N output ports, only the arbiter needs to run N times faster than the input lines.

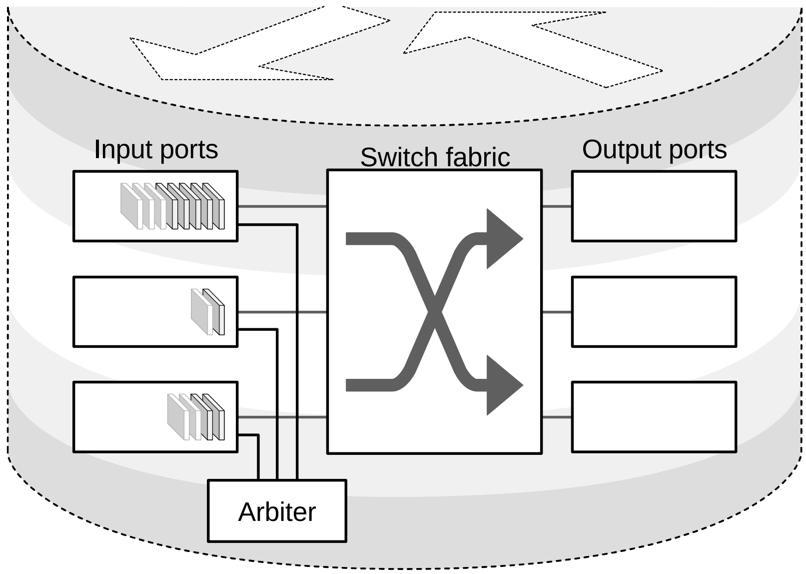

Figure 4-14: A router with input queuing. The arbiter releases a packet from an input queues when a path through the switch fabric and the output line is available.

The problem with input-queued routers is that if input queues are served in a first-come-first-served (FCFS) order, then a head-of-line packet destined for a busy output blocks the packets in the queue behind it. This is known as **head-of-line** (HOL) **blocking**. HOL blocking can be prevented if packets are served according to a timetable different from FCFS. Scheduling techniques that prevent HOL blocking are described in Chapter 5.

Output Queuing

In switches with pure *output queuing* (or, *output buffering*), packets are stored at the output ports. Incoming packets immediately proceed through the switching fabric to their corresponding output port. Because multiple packets may be simultaneously heading towards the same output port, the switch must provide a switching-fabric speedup proportional to the number of input ports. In a switch with N input ports, each output port must be able to store N packets in the time it takes a single packet to arrive at an input port. This makes the hardware for the switching fabric and output queues more expensive than in input queuing.

Notice that output queue do not suffer from head-of-line blocking—the transmitter can transmit the packets in the packets according to any desired timetable. The output port may rearrange its output queue to transmit packets according to a timetable different from their arrival order. We will study scheduling disciplines in Chapter 5.

Internal Queuing

- Head of line blocking
- What amount of buffering is needed?

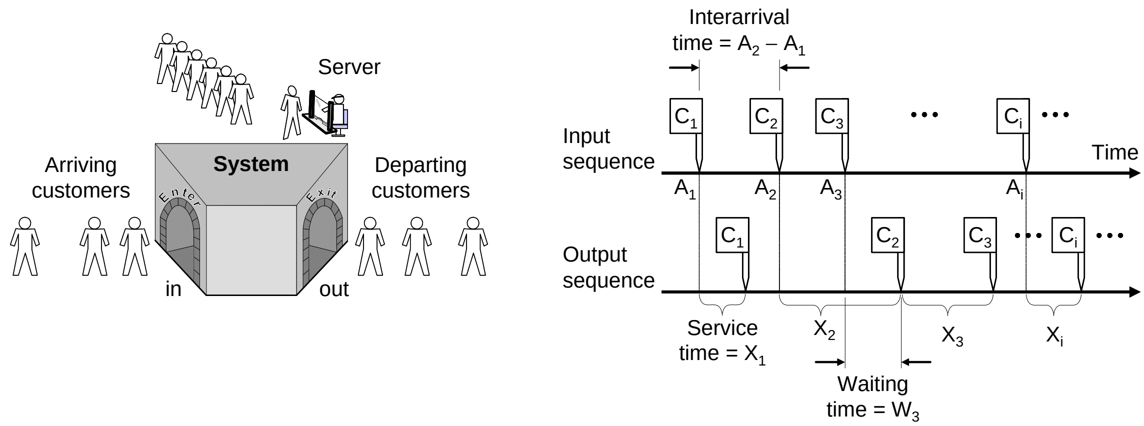

Figure 4-15: General service delay model: customers are delayed in a system for their own service time plus a possible waiting time. Customer 3 has to wait in line because a previous customer is being serviced at customer 3's arrival time.

4.2 Queuing Models

Queuing introduces latency, and the potential for packet loss if a queue overflows. When traffic patterns are bursty, the queuing-induced latency varies unpredictably from packet to packet, manifesting itself as jitter (delay variability) in the affected traffic streams. Modeling queuing processes is important for understanding the problem and designing the solutions.

General Server

A general service model is shown in Figure 4-15. Customers arrive in the system at a certain rate. It is helpful if the arrival times happened to be random and independent of the previous arrivals, because such systems can be well modeled. The server services customers in a certain order, the simplest being their order of arrival, also called *first-come-first-served* (FCFS). Every physical processing takes time, so a customer i takes a certain amount of time to service, the service time denoted as X_i.

Most commonly used performance measures are: (1) *the average number of customers in the system*; and, (2) *average delay per customer*. A successful method for calculating these parameters is based on the use of a *queuing model*. Figure 4-16 shows a simple example of a queuing model, where the system is represented by a single-server queue. The *queuing time* is the time that a customer waits before it enters the service. Figure 4-17 illustrates queuing system parameters on an example of a bank office with a single teller.

Why Queuing Happens?

Queuing occurs because of the server's inability to process the customers at the rate at which they are arriving. When a customer arrives at a busy server, it enters a waiting line (queue) and waits on its turn for processing. The critical assumption here is the following:

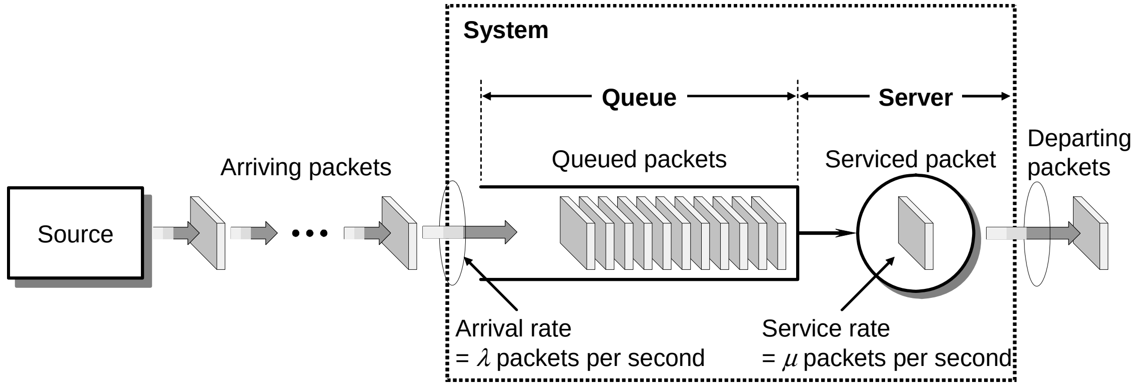

Figure 4-16: Simple queuing system with a single server.

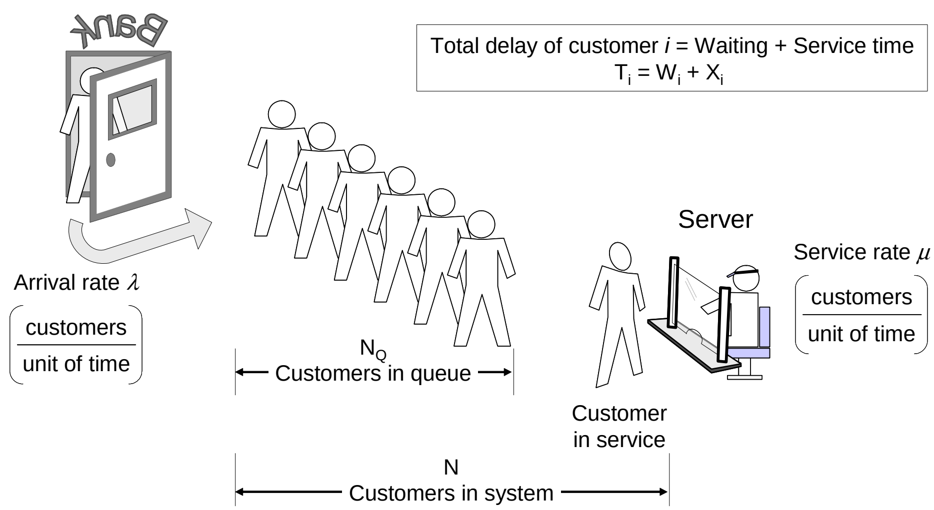

Figure 4-17: Illustration of queuing system parameters.

$$\text{Average arrival rate} \leq \text{Maximum service rate}$$

Otherwise, the queue length would grow unlimited and the system would become meaningless because some customers would have to wait infinite amount of time to be serviced. A corollary of this requirement is that queuing is an artifact of irregular customer arrival patterns, sometimes being too many, sometimes very few. Customers arriving in groups create queues. Had they been arriving "individually" (well spaced), allowing the server enough time to process the previous one, there would be no queuing. The arrival pattern where the actual arrival rate is equal to the average one would incur no queuing delays on any customer.

This is illustrated in Figure 4-18 where we consider a bank teller that can service five customers per hour, $\mu = 5\dfrac{\text{customers}}{\text{hour}}$, on average. This means, serving one customer takes 12 minutes, on average. Assume that for a stretch of time all arriving customers take 12 minutes to be served and that three customers arrive as shown in the figure. Although the server capacity is greater than the arrival rate, the second and third customers still need to wait in line before being served, because their arrivals are too closely spaced. If the customers arrived spaced according to their departure times at the same server, there would be no queuing delay for any customer. However, if this

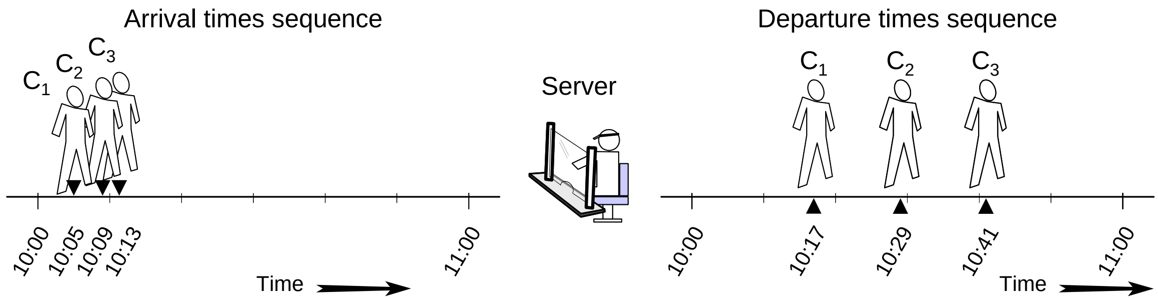

Figure 4-18: Illustration of how queues are formed. The server can serve 5 customers per hour and only 3 customers arrive during an hour period. Although the server capacity is greater than the arrival rate, some customers may still need to wait before being served, because their arrivals are too closely spaced.

sequence arrived at a server that can service only four customers per hour, again there would be queuing delays. Thus, having a server with service rate greater than the arrival rate is no guarantee that there will be no queuing delays. In summary, queuing results because packet arrivals cannot be preplanned and provisioned for—it is too costly or physically impossible to support peak arrival rates.

Note also that in the steady state, the average departure rate equals the average arrival rate. Server utilization = (arrival rate / max. service rate)

Communication Channel

Queuing delay is the time it takes to transmit the packets that arrived earlier at the network interface. Packet's service time is its transmission time, which is equal to L/C, where L is the packet length and C is the server capacity. In case of packet transmission, "server capacity" is the outgoing channel capacity. The average queuing time is typically a few transmission times, depending on the load of the network.

→()_____)→ $\quad$ delay $\propto$ capacity $^{-1}$

Another parameter that affects delay is *error rate*—errors result in retransmissions, which significantly influence the delay. Reliable vs. unreliable (if error correction is employed + Gaussian channel)

We study what are the sources of delay and try to estimate its amount. In a communication system, main delay contributors are (see Section 1.3):

- Processing (e.g., conversion of a stream of bytes to packets or packetization, compression/fidelity reduction, encryption, switching at routers, etc.)
- Queuing, due to irregular packet arrivals, sometimes too many, sometimes just few
- Transmission, converting the digital information into analog signals that travel the medium
- Propagation, signals can travel at most at the speed of light, which is finite
- Errors or loss in transmission or various other causes (e.g., insufficient buffer space in routers, recall Figure 2-11 for TCP), resulting in retransmission

Errors result in retransmission. For most links, error rates are negligible, but for multiaccess links, particularly wireless links, they are significant.

Processing may also need to be considered to form a queue if this time is not negligible.

Give example of how delay and capacity are related, see Figure from Peterson & Davie, or from [Jeremiah Hayes 1984].

Notation

Some of the symbols that will be used in this chapter are defined as follows (see also Figure 4-17):

$A(t)$ Counting process that represents the total number of tasks/customers that arrived from 0 to time t, i.e., $A(0) = 0$, and for $s < t$, $A(t) - A(s)$ equals the number of arrivals in the time interval $(s, t]$

λ Arrival rate, i.e., the average number of arrivals per unit of time, in steady state

$N(t)$ Number of tasks/customers in the system at time t

N Average number of tasks/customers in the system (this includes the tasks in the queue and the tasks currently in service) in steady state

N_Q Average number of tasks/customers waiting in queue (but not currently in service) in steady state

μ Service rate of the server (in customers per unit time) at which the server operates when busy

X_i Service time of the i^{th} arrival (depends on the particular server's service rate μ and can be different for different servers)

T_i Total time the i^{th} arrival spends in the system (includes waiting in queue plus service time)

T Average delay per task/customer (includes the time waiting in queue and the service time) in steady state

W Average queuing delay per task/customer (not including the service time) in steady state

ρ Rate of server capacity utilization (the fraction of time that the server is busy servicing a task, as opposed to idly waiting)

4.2.1 Little's Law

Imagine that you perform the following experiment. You are frequently visiting your local bank office and you always do the following:

1. As you walk into the bank, you count how many customers are in the room, including those waiting in the line and those currently being served. Let us denote the average count as N. You join the queue as the last person; there is no one behind you.

2. You will be waiting W time, on average, and then it will take X time, on average, for you to complete your job. The expected amount of time that has elapsed since you joined the queue until you are ready to leave is $T = W + X$. During this time T new customers will arrive at an arrival rate λ.
3. At the instant you are about to leave, you look over your shoulder at the customers who have arrived after you. These are all new customers that have arrived while you were waiting or being served. You will count, on average, $\lambda \cdot T$ customers in the system.

If you compare the average number of customers you counted at your arrival time (N) and the average number of customers you counted at your departure time ($\lambda \cdot T$), you will find that they are equal. This is called *Little's Law* and it relates the average number of tasks in the system, the average arrival rate of new tasks, and the average delay per task:

Average number of tasks in the system = Arrival rate × Average delay per task

$$N = \lambda \cdot T \tag{4.1a}$$

I will not present a formal proof of this result, but the reader should glean some intuition from the above experiment. For example, if customers arrive at the rate of 5 per minute and each spends 10 minutes in the system, Little's Law tells us that there will be 50 customers in the system on average.

The above observation experiment essentially states that the number of customers in the system, on average, does not depend on the time when you observe it. A stochastic process is **stationary** if *all* its statistical properties are invariant with respect to time.

Another version of Little's Law is

$$N_Q = \lambda \cdot W \tag{4.1b}$$

The argument is essentially the same, except that the customer looks over her shoulder as she enters service, rather than when completing the service. A more formal discussion is available in [Bertsekas & Gallagher, 1992].

Little's Law applies to any system in equilibrium, as long as nothing inside the system is creating new tasks or destroying them. Of course, to reach an equilibrium state we have to assume that the traffic source generates infinite number of tasks.

Using Little's Law, given any two variables, we can determine the third one. However, in practice it is not easy to get values that represent well the system under consideration. The reader should keep in mind that N, T, N_Q, and W are random variables; that is, they are not constant but have probability distributions. One way to obtain those probability distributions is to observe the system over a long period of time and acquire different statistics, much like traffic observers taking tally of people or cars passing through a certain public spot. Another option is to make certain assumptions about the statistical properties of the system. In the following, we will take the second approach, by making assumptions about statistics of customer arrivals and service times. From these statistics, we will be able to determine the expected values of other parameters needed to apply Little's Law.

Kendall's notation for queuing models specifies six factors:

Arrival Process / Service Proc. / Num. Servers / Max. Occupancy / User Population / Scheduling Discipline

1. *Arrival Process* (first symbol) indicates the statistical nature of the arrival process. The letter *M* is used to denote pure random arrivals or pure random service times. It stands for Markovian, a reference to the memoryless property of the exponential distribution of interarrival times. In other words, the arrival process is a Poisson process. Commonly used letters are:
   M – for exponential distribution of interarrival times
   G – for general independent distribution of interarrival times
   D – for deterministic (constant) interarrival times
2. *Service Process* (second symbol) indicates the nature of the probability distribution of the service times. For example, *M*, *G*, and *D* stand for exponential, general, and deterministic distributions, respectively. In all cases, successive interarrival times and service times are assumed to be statistically independent of each other.
3. *Number of Servers* (third symbol) specifies the number of servers in the system.
4. *Maximum Occupancy* (fourth symbol) is a number that specifies the waiting room capacity. Excess customers are blocked and not allowed into the system.
5. *User Population* (fifth symbol) is a number that specifies the total customer population (the "universe" of customers)
6. *Scheduling Discipline* (sixth symbol) indicates how the arriving customers are scheduled for service. Scheduling discipline is also called *Service Discipline* or *Queuing Discipline*. Commonly used service disciplines are the following:
   FCFS – first-come-first-served, also called first-in-first-out (FIFO), where the first customer that arrives in the system is the first customer to be served
   LCFS – last-come-first served (like a popup stack)
   FIRO – first-in-random-out

Service disciplines will be covered later in Chapter 5, where fair queuing (FQ) service discipline will be introduced. Only the first three symbols are commonly used in specifying a queuing model, although sometimes other symbols will be used in the rest of this chapter.

4.2.2 *M* / *M* / 1 Queuing System

Problems related to this section: Problem 4.16 → Problem 4.20

A correct notation for the system we consider is *M*/*M*/1/∞/∞/FCFS. This system can hold unlimited (infinite) number of customers, i.e., it has an unlimited waiting room size or the maximum queue length; the total customer population is unlimited; and, the customers are served in the FCFS order. It is common to omit the last three items and simply use *M*/*M*/1.

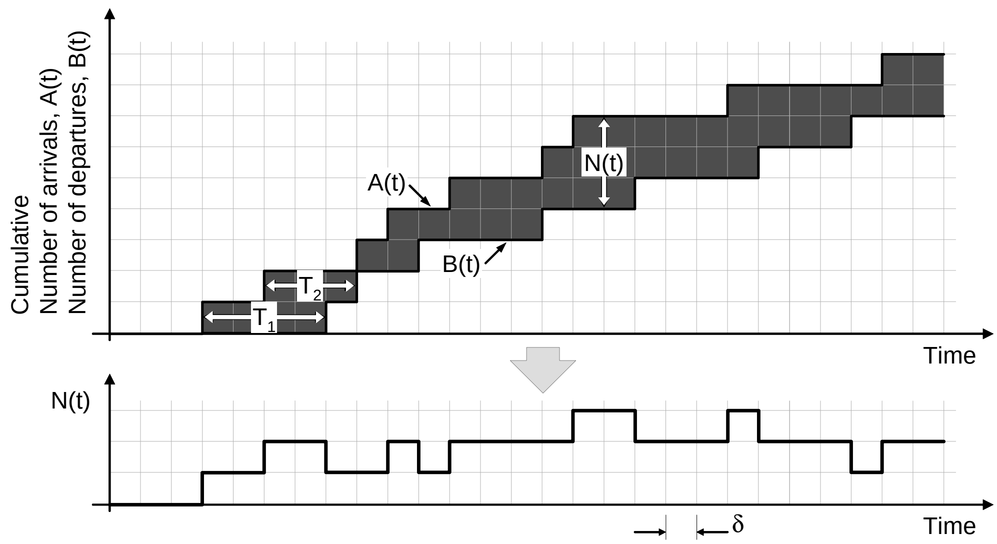

Figure 4-19: Example of birth and death processes. Top: Arrival and departure processes; Bottom: Number of customers in the system.

Figure 4-19 illustrates an $M/M/1$ queuing system, for which the process $A(t)$, total number of customers that arrived from 0 to time t, has a Poisson distribution. A Poisson process is generally considered a good model for the aggregate traffic of a large number of similar and independent customers. Then, $A(0) = 0$, and for $s < t$, $A(t) - A(s)$ equals the number of arrivals in the interval (s, t). The intervals between two arrivals (interarrival times) for a Poisson process are independent of each other and exponentially distributed with the parameter λ. If t_n denotes the time of the n^{th} arrival, the interarrival intervals $\tau_n = t_{n+1} - t_n$ have the probability distribution

$$P\{\tau_n \le s\} = 1 - e^{-\lambda \cdot s}, \qquad s \ge 0$$

It is important that we select the unit time period δ in Figure 4-19 small enough so that it is likely that at most one customer will arrive during δ. In other words, δ should be so small that it is unlikely that two or more customers will arrive during δ.

The process $A(t)$ is a pure *birth process* because it monotonically increases by one at each arrival event. So is the process $B(t)$, the number of departures up until time t. The process $N(t)$, the number of customers in the system at time t, is a *birth and death process* because it sometimes increases and at other times decreases. It increases by one at each arrival and decreases by one at each completion of service. We say that $N(t)$ represents the ***state*** of the system at time t. Notice that the state of this particular system (a birth and death process) can either increases by one or decreases by one—there are no other options. The intensity or rate at which the system state increases is λ and the intensity at which the system state decreases is μ. This means that we can represent the rate at which the system changes the state by the diagram in Figure 4-21.

Now suppose that the system has evolved to a steady-state condition. That means that the state of the system is independent of the starting state. The sequence $N(t)$ representing the number of customers in the system at different times does not converge. This is a random process taking unpredictable values. What does converge are the probabilities p_n that at any time a certain number of customers n will be observed in the system

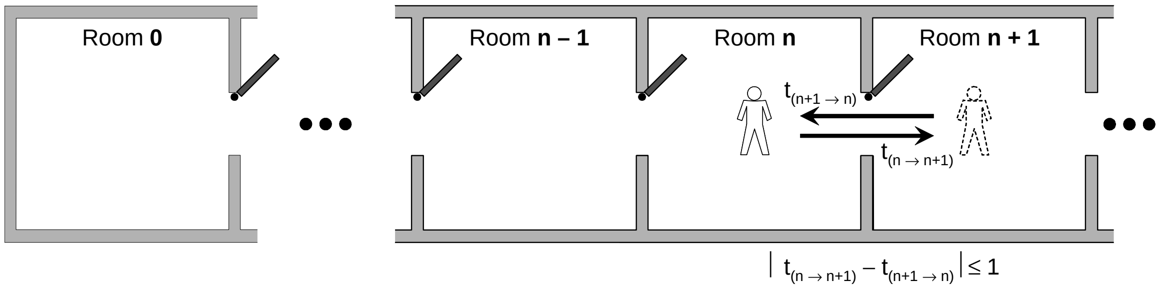

Figure 4-20: Intuition behind the balance principle for a birth and death process.

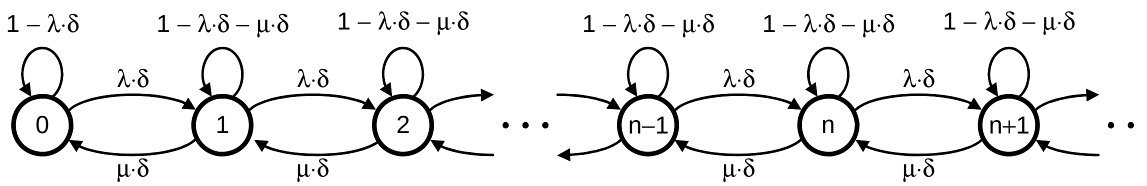

Figure 4-21: Transition probability diagram for the number of customers in the system.

$$\lim_{t \to \infty} P\{N(t) = n\} = p_n$$

Note that during any time interval, the total number of transitions from state n to $n + 1$ can differ from the total number of transitions from $n + 1$ to n by at most 1. Thus asymptotically, the frequency of transitions from n to $n + 1$ is equal to the frequency of transitions from $n + 1$ to n. This is called the *balance principle*. As an intuition, each state of this system can be imagined as a room, with doors connecting the adjacent rooms. If you keep walking from one room to the adjacent one and back, you can cross at most once more in one direction than in the other. In other words, the difference between how many times you went from $n + 1$ to n vs. from n to $n + 1$ at any time can be no more than one.

Given the stationary probabilities and the arrival and service rates, from our rate-equality principle we have the following *detailed balance equations*

$$p_n \cdot \lambda = p_{n+1} \cdot \mu, \qquad n = 0, 1, 2, \ldots \tag{4.2}$$

These equations simply state that the rate at which the process leaves state n equals the rate at which it enters that state. The ratio $\rho = \lambda/\mu$ is called the *utilization factor* of the queuing system, which is the long-run proportion of the time the server is busy. With this, we can rewrite the detailed balance equations as

$$p_{n+1} = \rho \cdot p_n = \rho^2 \cdot p_{n-1} = \ldots = \rho^{n+1} \cdot p_0 \tag{4.3}$$

If $\rho < 1$ (service rate exceeds arrival rate), the probabilities p_n are all positive and add up to unity, so

$$1 = \sum_{n=0}^{\infty} p_n = \sum_{n=0}^{\infty} \rho^n \cdot p_0 = p_0 \cdot \sum_{n=0}^{\infty} \rho^n = \frac{p_0}{1-\rho} \tag{4.4}$$

by using the well-known summation formula for the geometric series (see the derivation of Eq. (1.8) in Section 1.3.1). Combining equations (4.3) and (4.4), we obtain the probability of finding n customers in the system

$$p_n = P\{N(t) = n\} = \rho^n \cdot (1 - \rho), \qquad n = 0, 1, 2, \ldots \tag{4.5}$$

The average number of customers in the system in steady state is $N = \lim_{t\to\infty} E\{N(t)\}$. Because (4.5) is the p.m.f. for a geometric random variable, meaning that $N(t)$ has a geometric distribution, checking a probability textbook for the expected value of geometric distribution quickly yields

$$N = \lim_{t\to\infty} E\{N(t)\} = \frac{\rho}{1-\rho} = \frac{\lambda}{\mu - \lambda} \tag{4.6}$$

It turns out that for an $M/M/1$ system, by knowing only the arrival rate λ and service rate μ, we can determine the average number of customers in the system. From this, Little's Law (4.1a) gives the average delay per customer (waiting time in queue plus service time) as

$$T = \frac{N}{\lambda} = \frac{1}{\mu - \lambda} \tag{4.7}$$

The average waiting time in the queue, W, is the average delay T less the average service time $1/\mu$, like so

$$W = T - \frac{1}{\mu} = \frac{1}{\mu - \lambda} - \frac{1}{\mu} = \frac{\rho}{\mu - \lambda}$$

and by using the version (4.1b) of Little's Law, we have $N_Q = \lambda \cdot W = \rho^2 / (1 - \rho)$.

4.2.3 *M* / *M* / 1 / *m* Queuing System

Problems related to this section: Problem 4.22 → Problem 4.23

Now consider the $M/M/1/m$ system that is the same as $M/M/1$ except that the system can be occupied by up to m customers, which implies a finite waiting room or maximum queue length. The customers arriving when the queue is full are *blocked* and not allowed into the system. We have $p_n = \rho^n \cdot p_0$ for $0 \le n \le m$; otherwise $p_n = 0$. Using the relation $\sum_{n=0}^{m} p_n = 1$ we obtain

$$p_0 = \frac{1}{\sum_{n=0}^{m} \rho^n} = \frac{1-\rho}{1-\rho^{m+1}}, \quad 0 \le n \le m$$

From this, the steady-state occupancy probabilities are given by (cf. Eq. (4.5))

$$p_n = \frac{\rho^n \cdot (1-\rho)}{1-\rho^{m+1}}, \quad 0 \le n \le m \tag{4.8}$$

Assuming again that $\rho < 1$, the expected number of customers in the system is

$$N = E\{N(t)\} = \sum_{n=0}^{m} n \cdot p_n = \frac{1-\rho}{1-\rho^{m+1}} \cdot \sum_{n=0}^{m} n \cdot \rho^n = \frac{1-\rho}{1-\rho^{m+1}} \cdot \rho \cdot \sum_{n=0}^{m} n \cdot \rho^{n-1} = \frac{\rho \cdot (1-\rho)}{1-\rho^{m+1}} \cdot \frac{\partial}{\partial \rho}\left(\sum_{n=0}^{m} \rho^n\right)$$

$$= \frac{\rho \cdot (1-\rho)}{1-\rho^{m+1}} \cdot \frac{\partial}{\partial \rho}\left(\frac{1-\rho^{m+1}}{1-\rho}\right) = \frac{\rho}{1-\rho} - \frac{(m+1)\cdot \rho^{m+1}}{1-\rho^{m+1}} \tag{4.9}$$

Thus, the expected number of customers in the system is always less than for the unlimited queue length case, Eq. (4.6).

It is also of interest to know the probability of a customer arriving to a full waiting room, also called *blocking probability* p_B. Generally, the probability that a customer arrives when there are n customers in the queue is (using Bayes' formula)

$$P\{N(t)=n \mid \text{a customer arrives in } (t, t+\delta)\} = \frac{P\{\text{a customer arrives in}(t,t+\delta) \mid N(t)=n\} \cdot P\{N(t)=n\}}{P\{\text{a customer arrives in}(t,t+\delta)\}}$$

$$= \frac{(\lambda \cdot \delta)\cdot p_n}{\lambda \cdot \delta} = p_n$$

because of the memoryless assumption about the system. Thus, the blocking probability is the probability that an arrival will find m customers in the system, which is (using Eq. (4.8))

$$p_B = P\{N(t)=m\} = p_m = \frac{\rho^m \cdot (1-\rho)}{1-\rho^{m+1}} \tag{4.10}$$

4.2.4 *M* / *G* / 1 Queuing System

We now consider a class of systems where arrival process is still memoryless with rate λ. However, the service times have a general distribution—not necessarily exponential as in the $M/M/1$ system—meaning that we do not know anything about the distribution of service times. Suppose again that the customers are served in the order they arrive (FCFS) and that X_i is the service time of the i^{th} arrival. We assume that the random variables (X_1, X_2, …) are independent of each other and of the arrival process, and identically distributed according to an unspecified distribution function.

The class of $M/G/1$ systems is a superset of $M/M/1$ systems. The key difference is that in general there may be an additional component of memory. In such a case, one cannot say as for $M/M/1$ that the future of the process depends only on the present length of the queue. To calculate the average delay per customer, it is also necessary to account for the customer that has been in service for some time. Similar to $M/M/1$, we could define the state of the system as the number of customers in the system and use the so called moment generating functions to derive the system parameters. Instead, a simpler method from [Bertsekas & Gallagher, 1992] is used.

Assume that upon arrival the i^{th} customer finds N_i customers waiting in queue and one currently in service. The time the i^{th} customer will wait in the queue is given as

$$W_i = \sum_{j=i-N_i}^{i-1} X_j + R_i \tag{4.11}$$

where R_i is the *residual service time* seen by the i^{th} customer. By this we mean that if customer j is currently being served when i arrives, R_i is the remaining time until customer j's service is

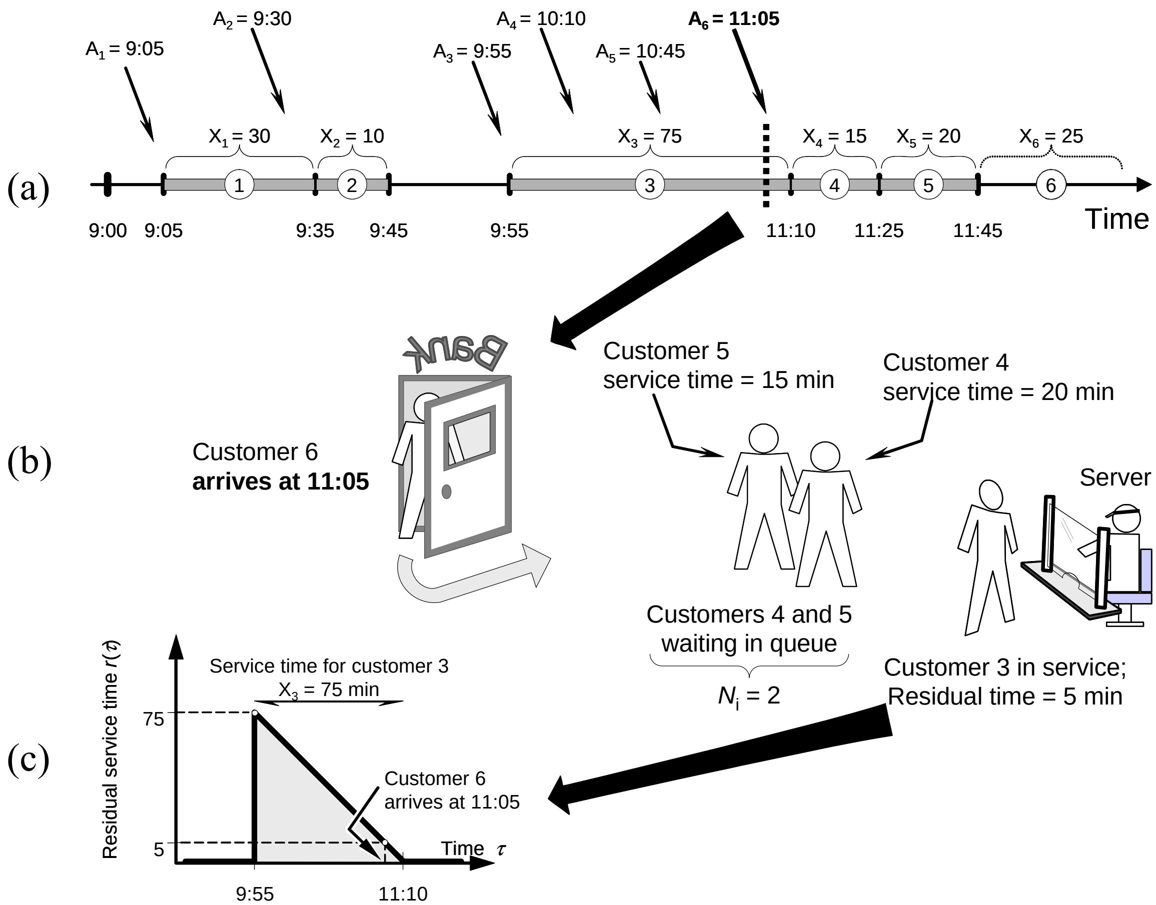

Figure 4-22: (a) Example of customer arrivals and service times; see Example 4.1 for details. (b) Detail of the situation found by customer 6 at his/her arrival. (c) Residual service time for customer 3 at the arrival of customer 6.

completed. The residual time's index is i (not j) because this time depends on i's arrival time and is not inherent to the served customer. If no customer is served at the time of i's arrival, then R_i is zero.

Example 4.1 Delay in a Bank Teller Service

An example pattern of customer arrivals to the bank from Figure 4-17 is shown in Figure 4-22. Assume that nothing is known about the distribution of service times. In this case, customer $k = 6$ will find customer 3 in service and customers 4 and 5 waiting in queue, i.e., $N_6 = 2$. The residual service time for 3 at the time of 6's arrival is 5 min. Thus, customer 6 will experience the following queuing delay:

$$W_6 = \sum_{j=4}^{5} X_j + R_6 = (15 + 20) + 5 = 40\,\text{min}$$

This formula simply adds up all the times shown in Figure 4-22(b). Notice that the residual time depends on the arrival time of customer $i = 6$ and not on how long the service time of customer $(i - N_i - 1) = 3$ is.

The total time that 6 will spend in the system (the bank) is $T_6 = W_6 + X_6 = 40 + 25 = 65$ min.

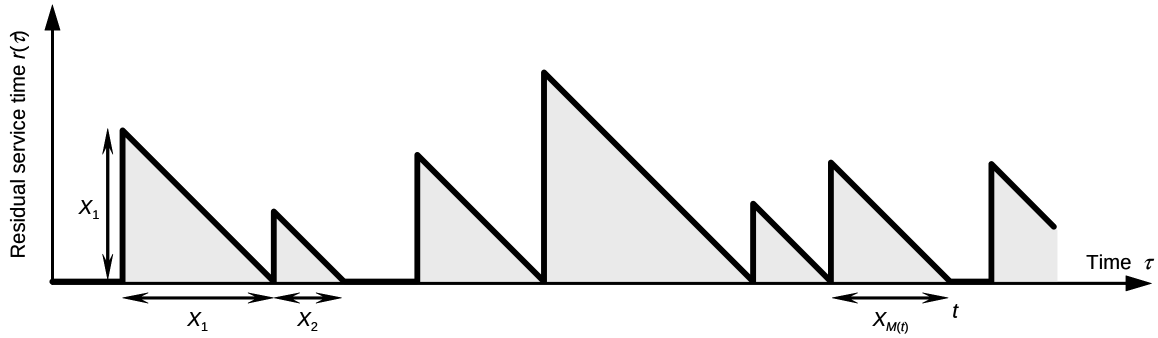

Figure 4-23: Expected residual service time computation. The time average of $r(\tau)$ is computed as the sum of areas of the isosceles triangles over the given period t.

By taking expectations of Eq. (4.11) and using the independence of the random variables N_i and $X_{i-1}, X_{i-2}, \ldots, X_{i-N_i}$ (which means that how many customers are found in the queue is independent of what business they came for), we have

$$E\{W_i\} = E\left\{\sum_{j=i-N_i}^{i-1} E\{X_j \mid N_i\}\right\} + E\{R_i\} = E\{X\} \cdot E\{N_i\} + E\{R_i\} \;=\; \frac{1}{\mu} \cdot N_Q + E\{R_i\} \qquad (4.12)$$

Throughout this section all long-term average quantities should be viewed as limits when time or customer index converges to infinity. We assume that these limits exist, which is true for most systems of interest provided that the utilization $\rho < 1$.The second term in the above equation is the *mean residual time*, $R = \lim_{t\to\infty} E\{R_i\}$, and it will be determined by a graphical argument. The residual service time $r(\tau)$ can be plotted as in Figure 4-22(c). The general case is shown in Figure 4-23. Every time a new customer enters the service, the residual time equals that customer's service time. Then it decays linearly until the customer's service is completed. The time average of $r(\tau)$ in the interval $[0, t]$ is

$$\frac{1}{t}\int_0^t r(\tau)d\tau = \frac{1}{t}\sum_{i=1}^{M(t)} \frac{1}{2} X_i^2$$

where $M(t)$ is the number of service completions within $[0, t]$. Hence, we obtain

$$R = \lim_{i\to\infty} E\{R_i\} = \lim_{i\to\infty}\left(\frac{1}{t}\int_0^t r(\tau)d\tau\right) = \frac{1}{2}\lim_{i\to\infty}\left(\frac{M(t)}{M(t)} \cdot \frac{1}{t}\sum_{i=1}^{M(t)} X_i^2\right) = \frac{1}{2}\lim_{i\to\infty}\left(\frac{M(t)}{t}\right) \cdot \lim_{i\to\infty}\left(\frac{\sum_{i=1}^{M(t)} X_i^2}{M(t)}\right) = \frac{1}{2}\lambda \cdot \overline{X^2}$$

where $\overline{X^2}$ is the second moment of service time, computed as

$$E\{X^n\} = \begin{cases} \sum_{X\,:\,p_i>0} p_i \cdot (X_i)^n & \text{if } X \text{ is discrete r.v.} \\ \int_{-\infty}^{\infty} f(x) \cdot x^n \cdot dx & \text{if } X \text{ is continuous r.v.} \end{cases}$$

By substituting this expression in the queue waiting time, Eq. (4.12), we obtain the so called Pollaczek-Khinchin (P-K) formula

$$W = \frac{\lambda \cdot \overline{X^2}}{2 \cdot (1 - \rho)} \tag{4.13}$$

The P-K formula holds for any distribution of service times as long as the variance of the service times is finite.

Example 4.2 Queuing Delays of an Go-Back-*N* ARQ

Consider a Go-Back-*N* ARQ such as described earlier in Section 1.3.2. Assume that packets arrive at the sender according to a Poisson process with rate λ. Assume also that errors affect only the data packets, from the sender to the receiver, and not the acknowledgment packets. What is the expected queuing delay per packet in this system?

Notice that the expected service time per packet equals the expected delay per packet transmission, which is determined in the solution of Problem 1.11 at the back of this text as follows

$$\overline{X} = E\{T_{\text{total}}\} = t_{\text{succ}} + \frac{p_{\text{fail}}}{1 - p_{\text{fail}}} \cdot t_{\text{fail}}$$

The second moment of the service time, $\overline{X^2}$, is determined similarly as:

Finally, Eq. (4.13) yields

$$W = \frac{\lambda \cdot \overline{X^2}}{2(1-\rho)} = \frac{\lambda \cdot \overline{X^2}}{2(1-\lambda/\mu)} = \frac{\lambda \cdot \overline{X^2}}{2(1-\lambda \cdot \overline{X})}$$

4.3 Networks of Queues

4.4 Summary and Bibliographical Notes

Section 4.1 focuses on one function of routers—forwarding packets—but this is just one of its many jobs. Section 1.4 describes another key function: building and maintaining the routing tables. In addition, more and more applications, such as firewalls, VPN concentration, voice gateways and video monitoring, are being implemented in routers. Cisco's Integrated Services Router (ISR), for example, even includes an optional application server blade for running various Linux and open source packages.

[Keshav & Sharma, 1998]

Kumar, *et al*. [1998] provide a good overview of router architectures and mechanisms.

James Aweya, "IP router architectures: An overview"

The material presented in this chapter requires basic understanding of probability and random processes. [Yates & Goodman, 2004] provides an excellent introduction and [Papoulis & Pillai, 2001] is a more advanced and comprehensive text.

[Bertsekas & Gallagher, 1992] provides a classic treatment of queuing delays in data networks. Most of the material in Sections 4.2 and 4.3 is derived from this reference.

Problems

Problem 4.1

Problem 4.2

Problem 4.3

Consider a regular PC that is used as a router, i.e., this is first-generation router architecture. The router has 4 network ports, each on its own line card. All four links have the same data rate of R bits/sec. The system bus operates at a four times higher data rate, i.e., $4\times R$ bps. Consider a scenario where steady traffic is arriving on all four ports and all packets are of the same length L.

(a) What is the worst-case delay that a packet can experience in this router?
(b) Will there be any head-of-line or output blocking observed?

Problem 4.4

Problem 4.5

Consider a router X that uses Banyan switching fabric. The figure below shows the router's connectivity to the adjacent routers, as well as the forwarding table of router X.

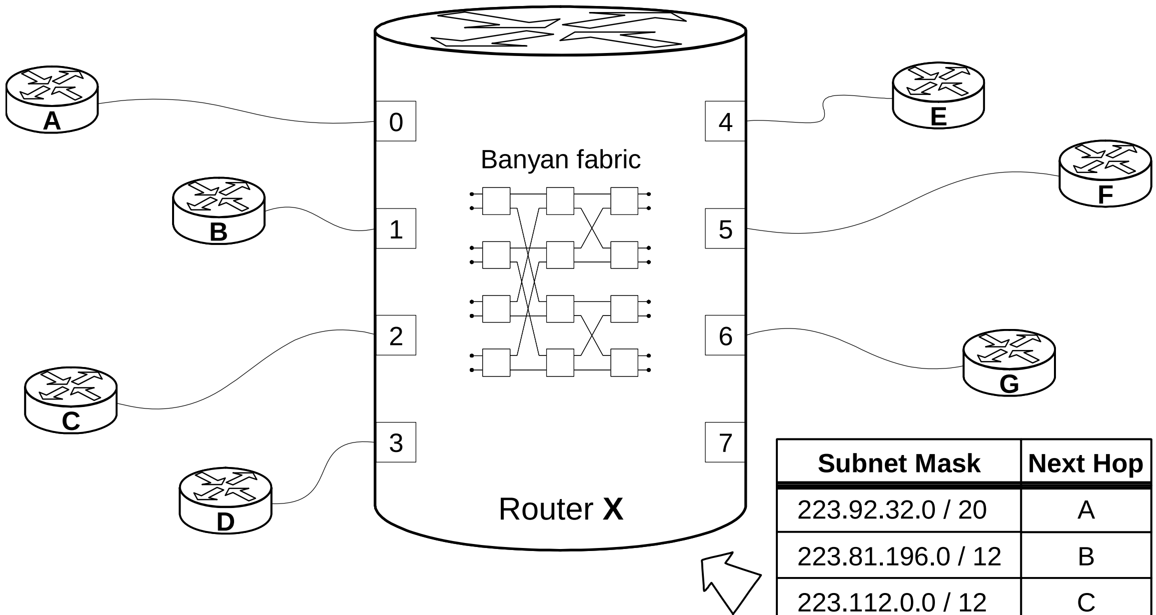

Subnet Mask	Next Hop
223.92.32.0 / 20	A
223.81.196.0 / 12	B
223.112.0.0 / 12	C
223.120.0.0 / 14	D
128.0.0.0 / 1	E
64.0.0.0 / 2	F
32.0.0.0 / 3	G

Assume that the following packets arrived simultaneously on router *X*:

Packet arrived from	Packet destination IP address
B	63.67.145.18
C	223.123.59.47
G	223.125.49.47

Draw and explain a diagram that shows how these packets will traverse the switching fabric.

Note: Check Problem 1.29 in Chapter 1 to see how the next hop for a packet is decided.

Problem 4.6

Problem 4.7

Using the switch icons shown below, sketch a simple 4×4 Batcher-Banyan network (without a trap network and a shuffle-exchange network). Label the input ports and the output ports of the fabric, from top to bottom, as 0, 1, 2 and 3.

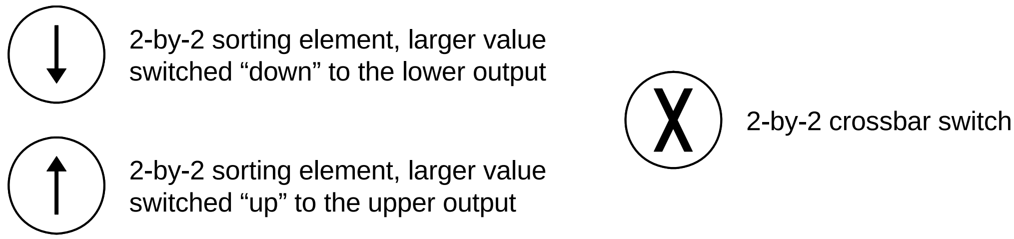

Suppose four packets are presented to the input ports of the Batcher-Banyan fabric that you sketched. Suppose further that the incoming packets are heading to the following output ports:

- Packet at input port 0 is heading to output port 1
- Packet at input port 1 to output port 0
- Packet at input 2 to output port 0
- Packet at input port 3 to output port 2

Show on the diagram of the fabric the switching of these packets through the fabric from the input ports to the output ports. Will any collisions and/or idle output lines occur?

Problem 4.8

Problem 4.9

Consider the router shown below where the data rates are the same for all the communication lines. The switch fabric is a crossbar so that at most one packet at a time can be transferred to a given output port, but different output ports can simultaneously receive packets from different input ports. Assume that the fabric moves packets two times faster than the data rate of the communication lines. Packets at each port are moved in a first-come-first-served (FCFS) order. If two or more packets arrive simultaneously at different input ports and are heading towards the same output port, their order of transfer is decided so that the lower index port wins (e.g., if ports 2 and 3 contend at the same time to the same output port, port 2 goes first). If a packet at a lower-index input port arrives and goes to the same output port for which a packet at a higher-index port is already waiting, then the higher-index port wins.

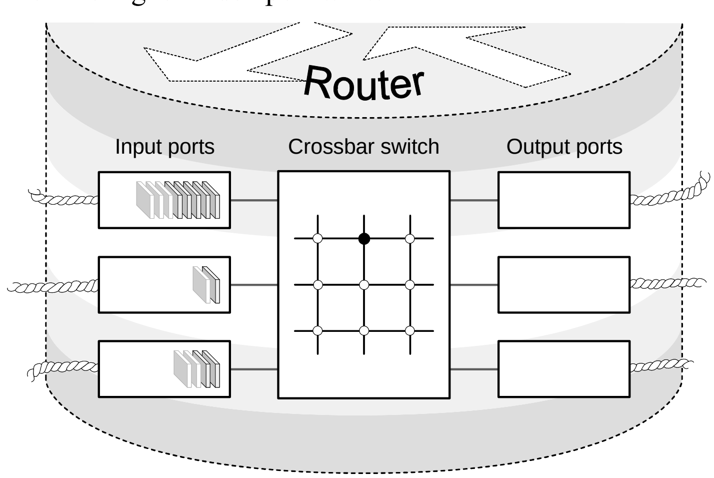

Consider the following traffic arrival pattern:

Input port 1: packet of length 2 received at time $t = 0$ heading to output port 2; packet of length 4 at time 4 to output 2

Input port 2: packet of length 8 at time 0 to output 2; packet of length 2 at time 2 to output 1

Input port 3: packet of length 8 at time 2 to output 2; packet of length 2 at time 4 to output 1

Draw the timing diagram for transfers of packets across the switching fabric. Will there be any head-of-line or output blocking observed? Explain your answer.

Problem 4.10

Problem 4.11

Problem 4.12

Problem 4.13

[Little's Law] Consider a system with a single server. The arrival and service times for the first 10 customers are as follows: ($A_1 = 0$, $X_1 = 3$); (2, 4); (3, 5); (4, 2); (6, 5); (7, 2); (10, 4); (11, 3); (12, 5); and ($A_{10} = 13$, $X_{10} = 3$).

(a) Draw the arrivals as a birth-death process similar to Figure 4-19.

(b) What is the average number of customers in the system N and the average delay T per customer in this system during the observed period? Assuming that the arrival rate is $\lambda =$ 1 customer/unit-of-time, does the system satisfy the Little's Law over the observed period?

Problem 4.14

Problem 4.15

Problem 4.16

Consider a router that can process 1,000,000 packets per second. Assume that the load offered to it is 950,000 packets per second. Also assume that the interarrival times and service durations are exponentially distributed.

(a) How much time will a packet, on average, spend being queued before being serviced?

(b) Compare the waiting time to the time that an average packet would spend in the router if no other packets arrived.

(c) How many packets, on average, can a packet expect to find in the router upon its arrival?

Problem 4.17

Consider an $M/G/1$ queue with the arrival and service rates λ and μ, respectively. What is the probability that an arriving customer will find the server busy (i.e., serving another customer)?

Problem 4.18

Messages arrive at random to be sent across a communications link with a data rate of 9600 bps. The link is 70% utilized, and the average message length is 1000 bytes. Determine the average waiting time for exponentially distributed length messages and for constant-length messages.

Problem 4.19

A facility of m identical machines is sharing a single repairperson. The time to repair a failed machine is exponentially distributed with mean $1/\lambda$. A machine, once operational, fails after a time that is exponentially distributed with mean $1/\mu$. All failure and repair times are independent. What is the steady-state proportion of time where there is no operational machine?

Problem 4.20

Imagine that K users share a link (e.g., Ethernet or Wi-Fi) with throughput rate R bps (i.e., R represents the actual number of file bits that can be transferred per second, after accounting for overheads and retransmissions). User's behavior is random and we model it as follows. Each user requests a file and waits for it to arrive. After receiving the file, the user sleeps for a random time, and then repeats the procedure. Each file has an exponential length, with mean $A \times R$ bits. Sleeping times between a user's subsequent requests are also exponentially distributed but with a mean of B seconds. All these random variables are independent. Write a formula that estimates the average time it takes a user to get a file since completion of his previous file transfer.

Problem 4.21

Problem 4.22

Consider a single queue with a constant service time of 4 seconds and a Poisson input with mean rate of 0.20 items per second.

(a) Find the mean and standard deviation of queue size

(b) Find the mean and standard deviation of the time a customer spends in system.

Problem 4.23

Consider the Go-back-N protocol used for communication over a noisy link with the probability of packet error equal to p_e. Assume that the link is memoryless, i.e., the packet error events are independent from transmission to transmission. Also assume that the following parameters are

given: the round-trip time (RTT), packet size L, and transmission rate R. What is the average number of successfully transmitted packets per unit of time (also called *throughput*), assuming that the sender always has a packet ready for transmission?
Hint: Recall that the average queuing delay per packet for the Go-back-N protocol is derived in Example 4.2 (Section 4.2.4).

Problem 4.24

Problem 4.25

Chapter 5
Mechanisms for Quality-of-Service

Contents

This chapter reviews mechanisms used in network routers to provide quality-of-service (QoS). Section 3.3 reviewed some end-to-ed mechanisms for providing quality of service, and hinted at mechanisms used in routers. This chapter details the router-based QoS mechanisms.

End-to-end QoS is built from the concatenation of edge-to-edge QoS from each network domain (or Autonomous System) through which traffic passes, and ultimately depends on the QoS characteristics of the individual hops along any given route. Networking solutions for end-to-end QoS are usually broken into three parts: per-hop QoS, traffic engineering, and signaling/provisioning. This chapter starts with mechanisms used to provide per-hop QoS, and describes traffic-engineering solutions in Section 5.4.3. Signaling/provisioning was already considered in Section 3.3.4 and will be considered further here.

The goal of per-hop QoS is to enable congestion-point routers and switches to provide predictable differentiated loss, latency, and jitter characteristics to traffic classes of interest to the service provider or its customers.

5.1 Scheduling

The queuing models in Chapter 4 considered delays and blocking probabilities under the assumption that tasks/packets are served on a first-come-first-served (FCFS) basis and that a task is blocked if it arrives at a full queue (if the waiting room capacity is limited). The property of a queue that decides the order of servicing of packets is called ***scheduling discipline*** (also called *service discipline* or *queuing discipline*, see Section 4.2). The property of a queue that decides which task is blocked from entering the system or which packet is dropped for a full queue is called *blocking policy* or *packet-discarding policy* or ***drop policy***. The simplest combination is

FCFS with tail drop, i.e., always service head of the line and, if necessary, drop the last arriving packet and this is what we considered in Section 4.2.3.

Scheduling has direct impact on a packet's queuing delay and hence on its total delay. Dropping decides whether the packet will arrive to destination at all. FCFS does not make any distinction between packets. A single FCFS queue cannot simultaneously support QoS-sensitive and QoS-insensitive traffic. While a long queue is less likely to overflow during a traffic burst (thus reducing packet loss probability), it potentially increases the queuing delay for non-dropped packets. A short queue reduces this delay, but conversely increases the probability of packet loss for bursty traffic.

Additional concerns may compel the network designer to consider making distinction between packets and design more complex scheduling disciplines and dropping policies. Such concerns include:

- *Prioritization*, where different tasks/packets can have assigned different priorities, so that the delay time for certain packets is reduced (at the expense of other packets)
- *Fairness*, so that different flows (identified by source-destination pairs) are offered equitable access to system resources
- *Protection*, so that misbehavior of some flows (by sending packets at a rate faster than their fair share) should not affect the performance achieved by other flows

Prioritization and fairness are complementary, rather than mutually exclusive. Fairness ensures that traffic flows of equal priority receive equitable service and that flows of lower priority are not excluded from receiving any service because all of it is consumed by higher priority flows. Fairness and protection are related so that ensuring fairness automatically provides protection, because it limits a misbehaving flow to its fair share. However, the converse need not be true. For example, if flows are *policed* at the entrance to the network, so that they are forced to confirm to a predeclared traffic pattern, they are protected from each other, but their resource shares may not be fair. Policing will be considered later in Section 5.2.

5.1.1 Scheduling Disciplines

We already mentioned that a single FCFS queue cannot simultaneously support QoS-sensitive and QoS-insensitive traffic. The solution is to split traffic across multiple queues at each congestion point, assigning different classes of traffic to queues sized for each class's desired loss, latency, and jitter characteristics. Access to the resource (e.g., outbound link) is mediated by a scheduler, which empties each queue in proportion to its allocated resource share or priority. Therefore, the system that wishes to make distinction between packets (QoS-enabled router or switch) must (1) classify packets, (2) differentially queue packets per class, and (3) provide controllable and predictable scheduling of packet transmissions from each class (queue) onto the outbound link. This approach is often referred to as a **classify, queue, and schedule (CQS) architecture** and it comprises two components (Figure 5-1):

1. ***Classifier****—Forms different waiting lines for different packet types*. The criteria for sorting packets into different lines include: priority, source and/or destination network address, application port number, etc.

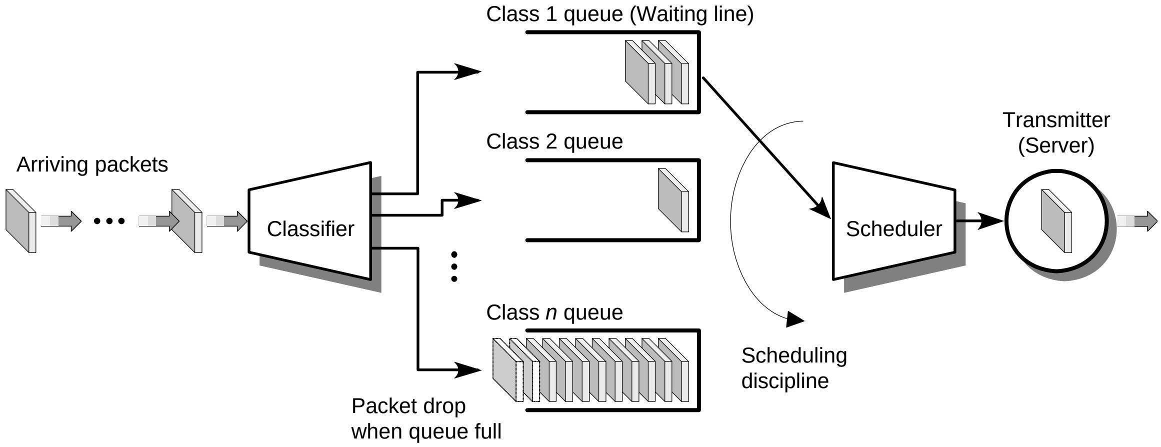

Figure 5-1: Components of a scheduler. Classifier sorts the arriving packets into different waiting lines based on one or more criteria, such as priority or source identity. Scheduler then places the packets into service based on the scheduling discipline. A single server serves all waiting lines.

2. ***Scheduler****—Calls packets from waiting lines for service*. Options for the rules of calling the packets for service (*scheduling discipline*) include: (*i*) first serve all the packets waiting in the high-priority line, if any; then go to the next lower priority class, etc.; (*ii*) serve the lines in round-robin manner by serving one or more packets from one line (but not necessarily all that are currently waiting), then go to serve few from the next waiting line, etc., then repeat the cycle.

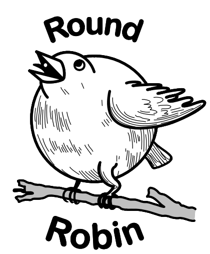

FCFS places indiscriminately all the arriving packets at the tail of a single queue. The idea with *prioritization* is that the packets with highest priority, upon arrival to the system, are placed at the head-of-the line, so they bypass waiting in the line. They may still need to wait if there is another packet (perhaps even of a lower priority) currently being transmitted. *Non-preemptive scheduling* is the discipline under which the ongoing transmission of lower-priority packets is not interrupted upon the arrival of a higher-priority packet. Conversely, *preemptive scheduling* is the discipline under which lower-priority packet is bumped out of service (back into the waiting line or dropped from the system) if a higher-priority packet arrives at the time a lower-priority packet is being transmitted.

Packet *priority* may be assigned simply based on the packet type, or it may be result of applying a complex set of *policies*. For example, the policies may specify that a certain packet type of a certain user type has high priority at a certain time of the day and low priority at other times.

Although priority scheduler does provide different performance characteristics to different classes, it still has shortcomings. For example, it does not deal with fairness and protection. An aggressive or misbehaving high-priority source may take over the communication line and elbow out all other sources. Not only the flows of the lower priority will suffer, but also the flows of the same priority are not protected from misbehaving flows.

A *round robin scheduler* alternates the service among different flows or classes of packets. In the simplest form of round robin scheduling the head of each queue is called, in turn, for service. That is, a class-1 packet is transmitted, followed by a class-2 packet, and so on until a class-n packet is transmitted. The whole round is repeated forever or until there are no more packets to transmit. If a particular queue is empty, because no packets of such type arrived in the meantime, the scheduler has two options:

1. Keep unused the portion of service or work allocated for that particular class and let the server stay idle (*non-work-conserving scheduler*)
2. Let the packet from another queue, if any, use this service (*work-conserving scheduler*)

A work-conserving scheduler will never allow the link (server) to remain idle if there are packets (of any class or flow) queued for transmission. When such scheduler looks for a packet of a given class but finds none, it will immediately check the next class in the round robin sequence.

One way to achieve control of channel conditions (hence, performance bounds) is to employ time division multiplexing (TDM) or frequency division multiplexing (FDM). TDM/FDM maintains a separate channel for each traffic flow and never mixes packets from different flows, so they never interfere with each other. TDM and FDM are non-work-conserving. Statistical multiplexing is work-conserving and that is what we consider in the rest of this section.

5.1.2 Fair Queuing

Problems related to this section: Problem 5.2 → Problem 5.8

Suppose that a system, such as transmission link, has insufficient resource to satisfy the demands of all users, each of whom has an equal right to the resource, but some essentially demand fewer resources than others. How, then, should we divide the resource? A sharing technique widely used in practice is called *max-min fair share*. Intuitively, a fair share first fulfils the demand of users who need less than they are entitled to, and then evenly distributes unused resources among the "big" users (Figure 5-2). Formally, we define max-min fair share allocation to be as follows:

- Resources are allocated in order of increasing demand
- No source obtains a resource share larger than its demand
- Sources with unsatisfied demands obtain an equal share of the resource

This formal definition corresponds to the following operational definition. Consider a set of sources 1, ..., n that have resource demands $r_1, r_2, ..., r_n$. Without loss of generality, order the source demands so that $r_1 \le r_2 \le \dots \le r_n$. Let the server have capacity C and all sources are equally entitled to the resource, although they may need more or less than they are entitled to. Then, we initially assign C/n of the resource to the source with the smallest demand, r_1. This may be more than what source 1 wants, perhaps, so we can continue the process. The process ends when each source receives no more than what it asks for, and, if its demand was not satisfied, no less than what any other source with a higher index (i.e., demand) received. We call such an allocation a *max-min fair allocation*, because it maximizes the minimum share of a source whose demand is not fully satisfied.

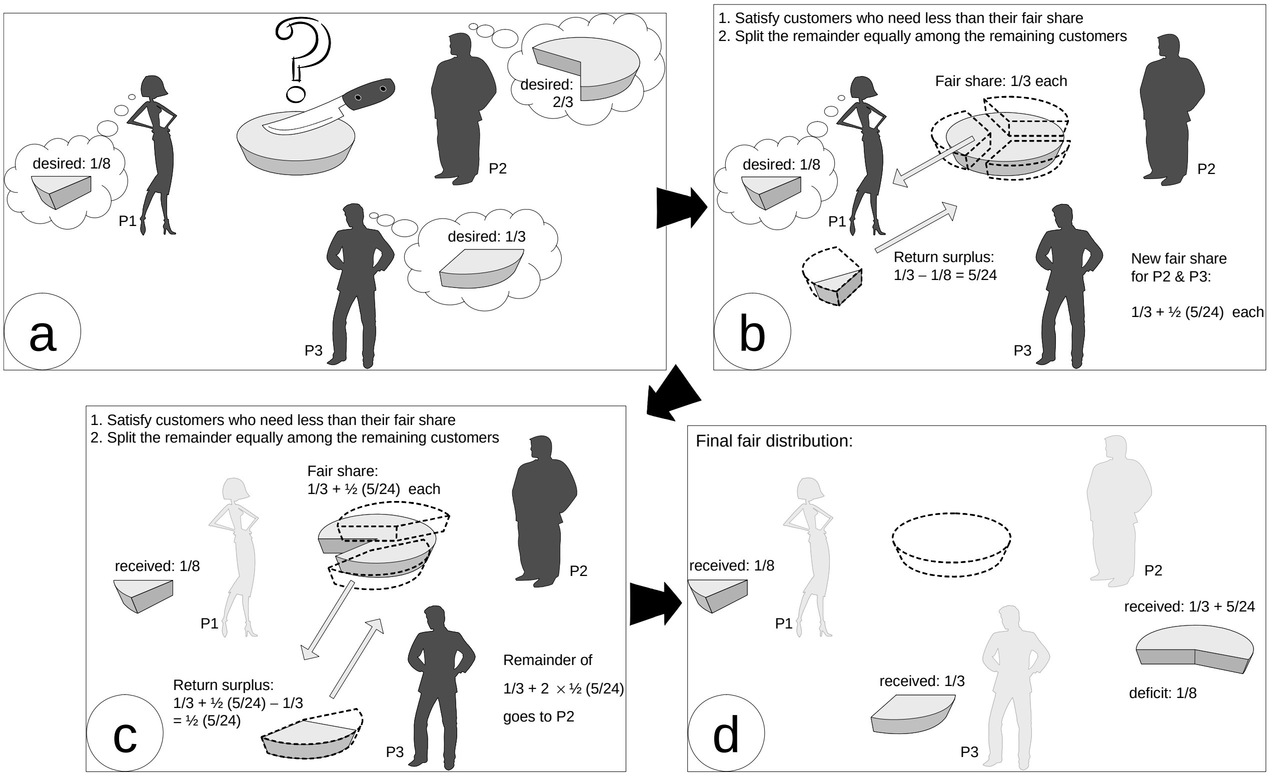

Figure 5-2: Illustration of max-min fair share algorithm; see text for details.

Example 5.1 Max-Min Fair Share

Consider the server in Figure 5-3 where packets are arriving from $n = 4$ sources of equal priority to be transmitted over a wireless link. (Assume that the full link bandwidth is available and ignore the link-layer overhead due to interfame spaces, backoff, collisions, etc.) The total required link capacity is:

$$\underbrace{8 \times 2048}_{\text{Appl. A}} + \underbrace{25 \times 2048}_{\text{Appl. B}} + \underbrace{50 \times 512}_{\text{Appl. C}} + \underbrace{40 \times 1024}_{\text{Appl. D}} = \underbrace{134{,}144 \text{ bytes/sec}}_{\text{Total demand}} = 1{,}073{,}152 \text{ bits/sec}$$

but the available capacity of the link is $C = 1$ Mbps $= 1{,}000{,}000$ bits/sec. By the notion of fairness and given that all sources are equally "important," each source is entitled to $C/n = ¼$ of the total capacity = 250 Kbps. Some sources may not need this much and the surplus is equally divided among the sources that need more than their fair share. The following table shows the max-min fair allocation procedure.

Sources	Demands [bps]	Balances after 1st round	Allocation #2 [bps]	Balances after 2nd round	Allocation #3 (Final) [bps]	Final balances
Application 1	131,072 bps	+118,928 bps	131,072	0	131,072 bps	0
Application 2	409,600 bps	–159,600 bps	332,064	–77,536 bps	336,448 bps	–73,152 bps
Application 3	204,800 bps	+45,200 bps	204,800	0	204,800 bps	0
Application 4	327,680 bps	–77,680 bps	332,064	+4,384 bps	327,680 bps	0

After the first round in which each source receives $¼C$, sources 1 and 3 have excess capacity, because they are entitled to more than what they need. The surplus of $C' = 118{,}928 + 45{,}200 = 164{,}128$ bps is equally distributed between the sources in deficit, that is sources 2 and 4. After the second round of allocations, source 4 has excess of $C'' = 4{,}384$ bps and this is allocated to the only remaining source in

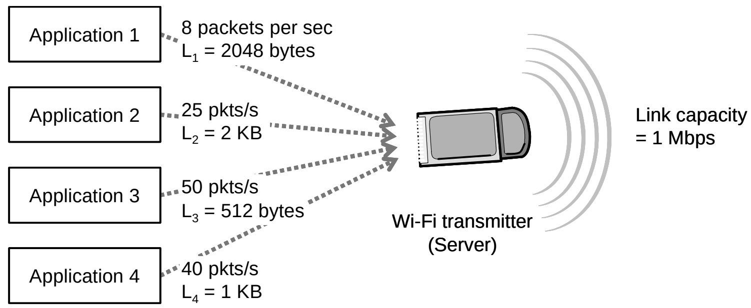

Figure 5-3: Example of a server (Wi-Fi transmitter) transmitting packets from four sources (applications) over a wireless link; see text for details.

deficit, which is source 2. Finally, under the fair resource allocation, sources 1, 3, and 4 have fulfilled their needs, but source 2 remains short of 73.152 Kbps.

Thus far, we have assumed that all sources are equally entitled to the resources. Sometimes, we may want to assign some sources a greater share than others. In particular, we may want to associate ***weights*** $w_1, w_2, ..., w_n$ with sources 1, 2, …, n, to reflect their relative entitlements to the resource. We extend the concept of max-min fair share to include such weights by defining the *max-min weighted fair share allocation* as follows:

- Resources are allocated in order of increasing demand, normalized by the weight
- No source obtains a resource share larger than its demand
- Sources with unsatisfied demands obtain resource shares in proportion to their weights

The following example illustrates the procedure.

Example 5.2 Weighted Max-Min Fair Share

Consider the same scenario as in Example 5.1, but now assume that the sources are weighted as follows: $w_1 = 0.5$, $w_2 = 2$, $w_3 = 1.75$, and $w_4 = 0.75$. The first step is to normalize the weights so they are all integers, which yields: $w'_1 = 2$, $w'_2 = 8$, $w'_3 = 7$, and $w'_4 = 3$. A source i is entitled to $w'_i \cdot \frac{1}{\sum w'_j}$ of the total capacity, which yields 2/20, 8/20, 7/20, and 3/20, for the respective four sources. The following table shows the results of the *weighted* max-min fair allocation procedure.

Src	Demands	Allocation #1 [bps]	Balances after 1st round	Allocation #2 [bps]	Balances after 2nd round	Allocation #3 (Final) [bps]	Final balances
1	131,072 bps	100,000	–31,072	122,338	–8,734 bps	131,072 bps	0
2	409,600 bps	400,000	–9,600	489,354	+79,754 bps	409,600 bps	0
3	204,800 bps	350,000	+145,200	204,800	0	204,800 bps	0
4	327,680 bps	150,000	–177,680	183,508	–144,172 bps	254,528 bps	–73,152 bps

This time around, source 3 in the first round is allocated more than it needs, while all other sources are in deficit. The excess amount of $C' = 145{,}200$ bps is distributed as follows. Source 1 receives $\frac{2}{2+8+3} \cdot 145{,}200 = 22{,}338$ bps, source 2 receives $\frac{8}{2+8+3} \cdot 145{,}200 = 89{,}354$ bps, and source 4 receives

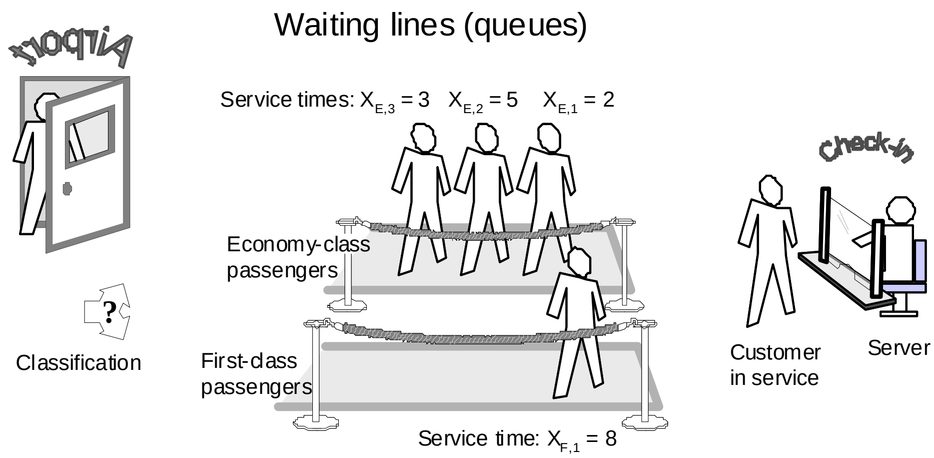

Figure 5-4: Dynamic fair-share problem: in what order should the currently waiting customers be called in service?

$\frac{3}{2+8+3} \cdot 145{,}200 = 33{,}508$ bps. Notice that in the denominators is always the sum of weights for the currently considered sources. After the second round of allocations, source 2 has excess of $C'' = 79{,}754$ bps and this is distributed among sources 1 and 4. Source 1 receives $\frac{2}{2+3} \cdot 79{,}754 = 31{,}902$, which along with 122,338 it already has yields more than it needs. The excess of $C''' = 23{,}168$ is given to source 4, which still remains short of 73.152 Kbps.

Min-max fair share (MMFS) defines the ideal fair distribution of a shared scarce resource. Given the resource capacity C and n customers, under MMFS a customer i is guaranteed to obtain at least $C_i = C/n$ of the resource. If some customers need less than what they are entitled to, then other customers can receive more than C/n. Under weighted MMFS (WMMFS), a customer i is guaranteed to obtain at least $C_i = \dfrac{w_i}{\sum_{j=1}^{n} w_j} C$ of the resource.

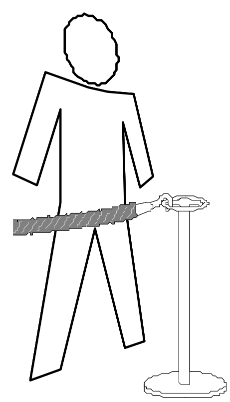

However, MMFS does not specify how to achieve this in a *dynamic system* where the demands for resource vary over time. To better understand the problem, consider the airport check-in scenario illustrated in Figure 5-4. Assume there is a single window (server) and both first-class and economy passengers are given the same weight. The question is, in which order the waiting customers should be called for service so that both queues obtain equitable access to the server resource? Based on the specified service times (Figure 5-4), the reader may have an intuition that, in order to maintain fairness on average, it is appropriate to call the first two economy-class passengers before the first-class passenger and finally the last economy-class passenger. The rest of this section reviews practical schemes that achieve just this and, therefore, guarantee (weighted) min-max fair share resource allocation when averaged over a long run.

Generalized Processor Sharing (GPS)

Min-max fair share cannot be directly applied in network scenarios because packets are transmitted as atomic units and they can be of different length, thus requiring different

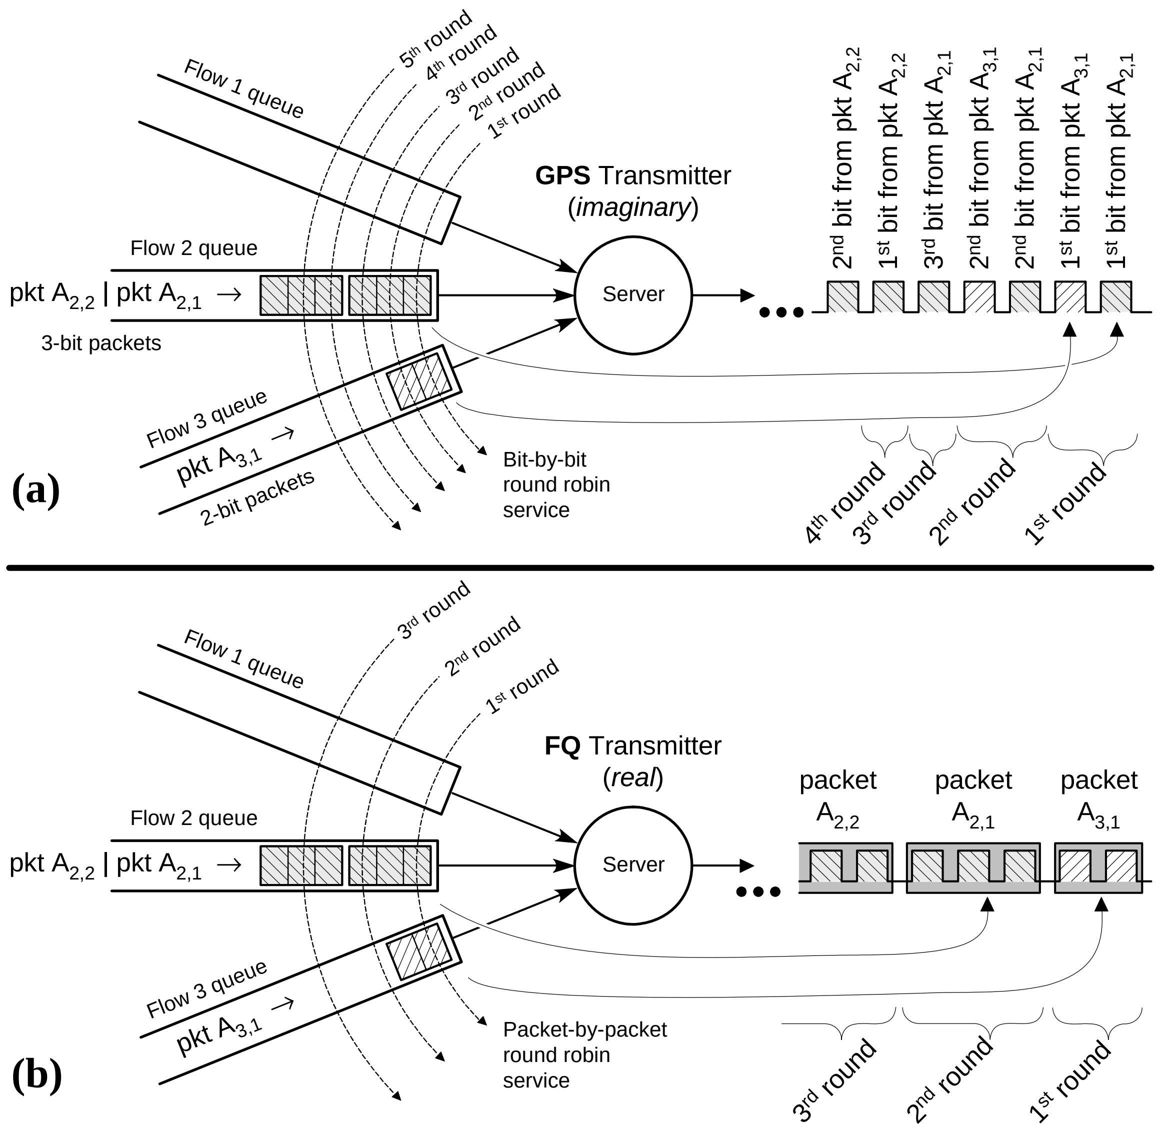

Figure 5-5: Imaginary bit-by-bit GPS (generalized processor sharing) (a) is used to derive the schedule for actual FQ (fair queuing) scheduling (b) used in real routers.

transmission times. In Example 5.1, packets from sources 1 and 2 are twice longer than those from source 4 and four times than from source 3. It is, therefore, difficult to keep track of whether each source receives its fair share of the server capacity. To arrive at a practical technique, we start by considering an idealized technique called *Generalized Processor Sharing* (GPS).

GPS maintains different waiting lines for packets belonging to different flows. There are two restrictions that apply:

- A packet cannot jump its waiting line, i.e., scheduling within individual queues is FCFS
- Service is *non-preemptive*, meaning that an arriving packet never bumps out a packet that is currently being transmitted (in service)

GPS works in a bit-by-bit round robin fashion, as illustrated in Figure 5-5(a). That is, the router transmits a bit from queue 1 (if there is any), then a bit from queue 2, and so on, for all queues

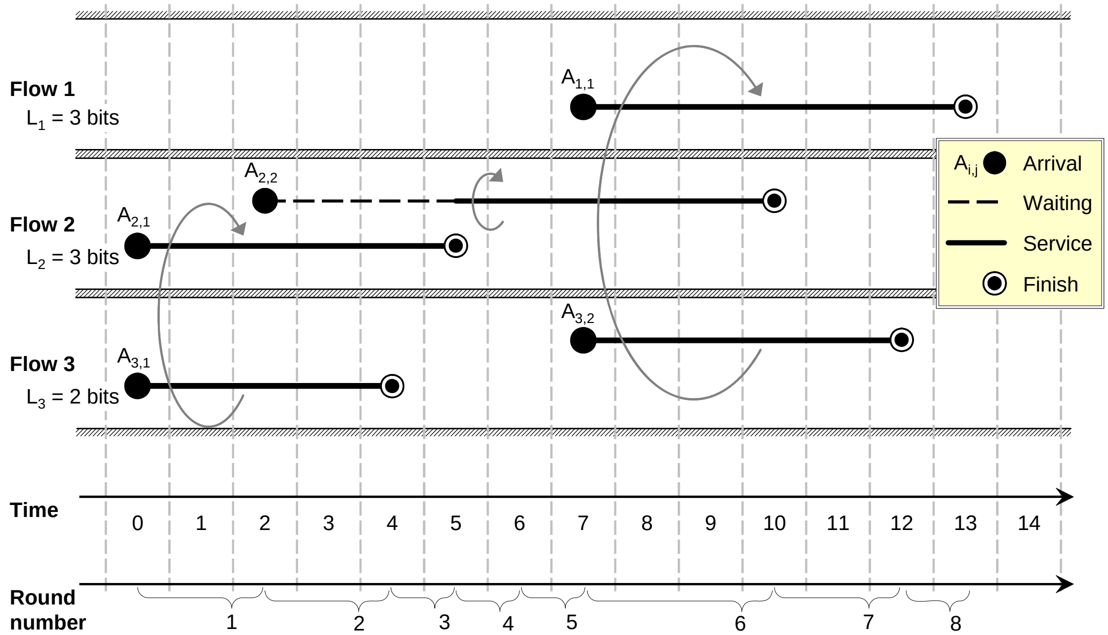

Figure 5-6. Example of bit-by-bit GPS. The output link service rate is *C* = 1 bit/s.

that have packets ready for transmission. Let $A_{i,j}$ denote the time that j^{th} packet arrives from i^{th} flow at the server (transmitter). Let us, for the sake of illustration, consider the example with packets only 2 or 3 bits long. Packets from flows 1 and 2 are 3 bits long, and from flow 3 are 2 bits long. At time zero, packet $A_{2,1}$ arrives from flow 2 and packet $A_{3,1}$ arrives from flow 3. Both packets find no other packets in their respective waiting lines, so their transmission starts immediately, one bit per round. Notice that one *round* of transmission now takes *two* time units, because two flows must be served per round. Because the packet from flow 3 is shorter than that from flow 2 (and their transmission started simultaneously), the transmission of packet $A_{3,1}$ will be finished sooner. After this moment, one *round* of transmission now takes *one* time unit, because only one flow (flow 2) must be served per round.

The key idea of Fair Queuing (FQ) is to run imaginary GPS transmissions, as in Figure 5-5(a), determined packet finish-round numbers under GPS, and then line up the packets for actual transmission under FQ in the ascending order of their finish numbers, as in Figure 5-5(b). This process will be elaborated in the rest of this section.

The GPS example from in Figure 5-5(a) is continued in Figure 5-6. At time t = 2 there is one more arrival on flow 2: $A_{2,2}$. Because in flow 2 $A_{2,2}$ finds $A_{2,1}$ in front of it, it must wait until the transmission of $A_{2,1}$ is completed. At time t = 7 there are two arrivals: $A_{1,1}$ and $A_{3,2}$. The transmission of both packets starts immediately because currently they are the only packets in their flows. (The bits should be transmitted atomically, a bit from each flow per unit of time as shown in Figure 5-5(a), rather than continuously as shown in Figure 5-6, but because this is an abstraction anyway, we leave it as is.)

As seen, a *k*-bit packet takes always *k* rounds to transmit, but the actual time duration can vary, depending on the current number of active flows—the more flows served in a round, the longer the round takes. (A flow is *active* if it has packets enqueued for transmission.) For example, in

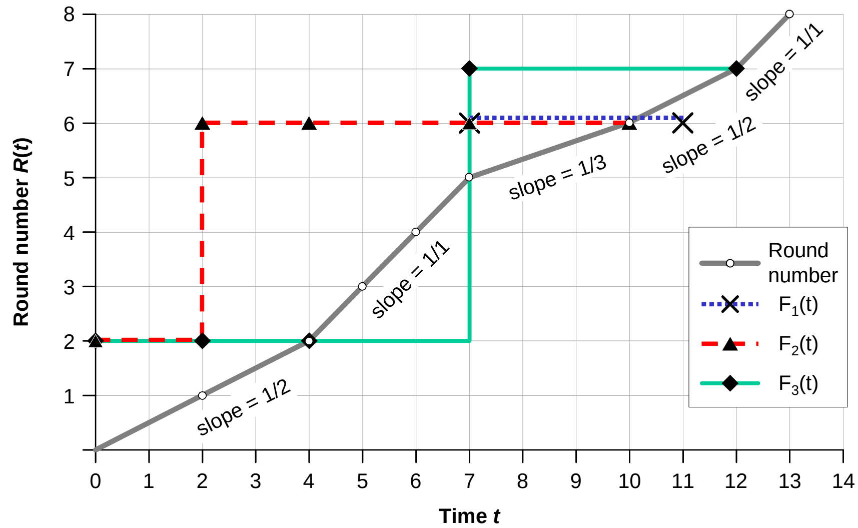

Figure 5-7. Piecewise linear relationship between round number and time for the example from Figure 5-6. Also shown are finish numbers $F_i(t)$ for the different flows.

Figure 5-6 it takes 4 s to transmit the first packet of flow 3, and 8 s for the second packet of the same flow although both have equal length!

The piecewise linear relationship between the time and round number is illustrated in Figure 5-7 for the example from Figure 5-6. The slope of each linear segment is computed as one round divided by the number of bits that need to be transmitted in one round. In other words, the slope is inversely proportional to the current number of active flows. Looking back at Figure 5-5(a), we see that in the beginning, the GPS transmitter needs to transmit two bits per round during the first two rounds (bits from flows 2 and 3). Hence, the slope of the round-number curve is 1/2. At time 4, transmission of the packet from flow 3 is finished, so the slope rises to 1/1 because there remains a single active flow. Similarly, at time 7 the slope falls to 1/3. In general, the function $R(t)$ increases at a rate

$$\frac{dR(t)}{dt} = \frac{C}{n_{\text{active}}(t)} \tag{5.1}$$

where C is the transmission capacity of the output line. Obviously, if $n_{\text{active}}(t)$ is *constant* then $R(t) = t \cdot C / n_{\text{active}}$, but this need not be the case because packets in different flows arrive randomly. In general, the round number is determined in a piecewise manner $R(t_i) = \sum_{j=1}^{i} \frac{(t_j - t_{j-1}) \cdot C}{n_{\text{active}}(t_j)} + \frac{t_1 \cdot C}{n_{\text{active}}(t_1)}$ as will be seen later in Example 5.3. Each time $R(t)$ reaches a new integer value marks an instant at which all the queues have been given an equal number of opportunities to transmit a bit (of course, an empty queue does not utilize its given opportunity).

GPS provides max-min fair resource allocation. Unfortunately, GPS cannot be implemented because it is not feasible to interleave the bits from different flows. A practical solution is the fair queuing mechanism that approximates this behavior on a packet-by-packet basis, which is presented next.

Fair Queuing

Similar to GPS, a router using FQ maintains different waiting lines for packets belonging to different flows at each output port. FQ determines when a given packet would finish being transmitted if it were being sent using bit-by-bit round robin (GPS) and then uses this finishing tag to rank order the packets for transmission.

The service round in which a packet $A_{i,j}$ would finish service under GPS is called the packet's **finish number**, denoted $F_{i,j}$. For example, in Figure 5-5(a) and Figure 5-6 packet $A_{3,1}$ has finish number $F_{3,1}$ = 2, packet $A_{2,1}$ has finish number $F_{2,1}$ = 3, and so on. Obviously, packet's finish number depends on the packet size and the current round number at the start of packet's service. It is important to recall that the finish number is, generally, different from the actual time at which the packet is served. For example, packet $A_{3,1}$ is serviced by $t = 4$, and $A_{2,1}$ is serviced by $t = 5$, because time is different from the round number (Figure 5-7).

Let $L_{i,j}$ denote the size (in bits) of packet $A_{i,j}$. Under bit-by-bit GPS it takes $L_{i,j}$ rounds of service to transmit this packet. Let $F_{i,j}$ denote the time when the transmitter finishes transmitting j^{th} packet from i^{th} flow. Suppose that a packet arrives at time t_a on a server which previously cycled through $R(t_a)$ rounds. Under GPS, the packet would have to wait for service *only if* there are currently packets from *this* flow either under service or enqueued for service or both—packets from other flows would not affect the start of service for this packet. Therefore, the start round number for servicing packet $A_{i,j}$ is the *highest* of these two

- The current round $R(t_a)$ at the packet's arrival time t_a
- The finishing round of the last packet, if any, from the *same* flow

or in short, the start round number of the packet $A_{i,j}$ is $\max\{F_{i,j-1}, R(t_a)\}$. The finish number of this packet is computed as

$$F_{i,j} = \max\{F_{i,j-1}, R(t_a)\} + L_{i,j} \tag{5.2}$$

Once assigned, the finish number remains constant and does *not* depend on future packet arrivals and departures. FQ scheduler performs the following procedure every time a new packet arrives:

1. Calculate the finish number for the newly arrived packet using Eq. (5.2)
2. For all the packets *currently waiting* for service (in any queue), sort them in the ascending order of their finish numbers
3. When the packet currently in transmission, if any, is finished, call the packet with the *smallest finish number* in service

Note that the sorting in step 2 does not include the packet currently being transmitted, if any, because FQ uses non-preemptive scheduling. Also, it is possible that a packet can be scheduled ahead of a packet waiting in a different line because the former is shorter than the latter and its finish number happens to be smaller than the finish number of the already waiting (longer) packet. The fact that FQ uses non-preemptive scheduling makes it an approximation of the bit-by-bit round robin GPS, rather than an exact simulation.

For example, Figure 5-7 also shows the curves for the finish numbers $F_i(t)$ of the three flows from Figure 5-6. At time 0, packets $A_{2,1}$ and $A_{3,1}$ arrive simultaneously, but because $F_{3,1}$ is smaller than $F_{2,1}$, packet $A_{3,1}$ goes into service first.

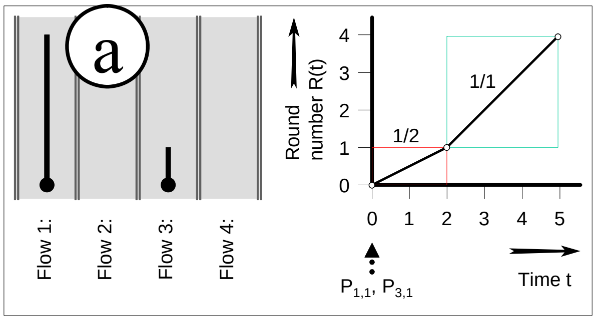

Arrival time	Flow ID	Packet size
t = 0	@Flow1	16,384 bytes
t = 0	@Flow3	4,096 bytes
t = 3	@Flow2	16,384 bytes
t = 3	@Flow4	8,192 bytes
t = 6	@Flow3	4,096 bytes
t = 12	@Flow3	4,096 bytes

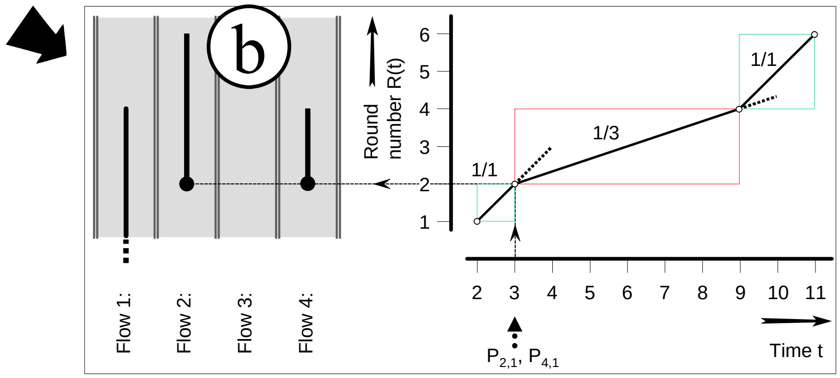

Figure 5-8: Determining the round numbers under bit-by-bit GPS for Example 5.3. (a) Initial round numbers as determined at the arrival of packets $P_{1,1}$ and $P_{3,1}$. (b) Round numbers recomputed at the arrival of packets $P_{2,1}$ and $P_{4,1}$. (Continued in Figure 5-9.)

Example 5.3 Packet-by-Packet Fair Queuing

Consider the system from Figure 5-3 and Example 5.1 and assume for the sake of illustration that the time is quantized to the units of the transmission time of the smallest packets. The smallest packets are 512 bytes from flow 3 and on a 1 Mbps link it takes 4.096 ms to transmit such a packet. For the sake of illustration, assume that a packet arrives on flows 1 and 3 each at time zero, then a packet arrives on flows 2 and 4 each at time 3, and packets arrive on flow 3 at times 6 and 12. Show the corresponding packet-by-packet FQ scheduling.

The first step is to determine the round numbers for the arriving packets, given their arrival times. The process is illustrated in Figure 5-8. The round numbers are shown also in the units of the smallest packet's number of bits, so these numbers must be multiplied by 4096 to obtain the actual round number. Bit-by-bit GPS would transmit bits from two packets ($A_{1,1}$ and $A_{3,1}$) in the first round, so the round takes two time units and the slope is 1/2. In the second round, only one packet is being transmitted ($A_{1,1}$), the round duration is one time unit and the slope is 1/1. The GPS server completes two rounds by time 3, $R(3) = 2$, at which point two new packets arrive ($A_{2,1}$ and $A_{4,1}$). The next arrival is at time 6 (the actual time is $t_1 = 24.576$ ms) and the round number is determined as

$$R(t_3) = \frac{(t_3 - t_2)\cdot C}{n_{\text{active}}(t_3)} + R(t_2) = \frac{(0.024576 - 0.012288)\times 1000000}{3} + 8192 = 4094 + 8192 = 12288$$

which in our simplified units is $R(6) = 3$. The left side of each diagram in Figure 5-8 also shows how the packet arrival times are mapped into the round number units. Figure 5-9 summarizes the process of determining the round numbers for all the arriving packets.

The actual order of transmissions under packet-by-packet FQ is shown in Figure 5-10. At time 0 the finish numbers are: $F_{1,1} = 4$ and $F_{3,1} = 1$, so packet $A_{3,1}$ is transmitted first and packet $A_{1,1}$ goes second.

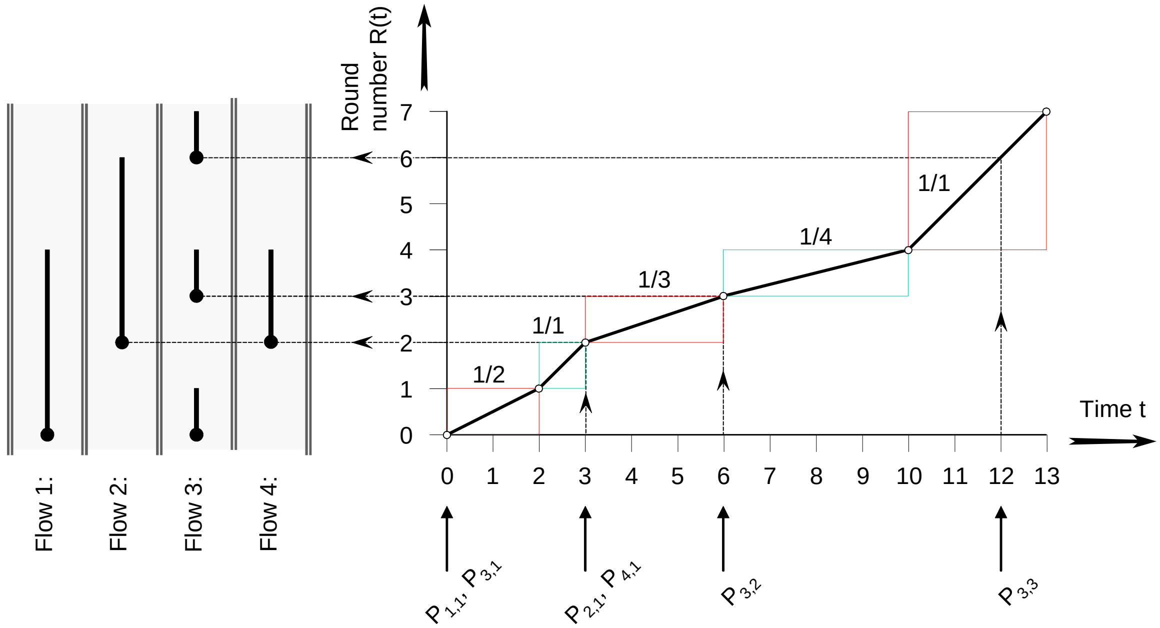

Figure 5-9: Determining the round numbers under bit-by-bit GPS for Example 5.3, completed from Figure 5-8.

At time 3 the finish numbers for the newly arrived packets are: $F_{2,1} = \max\{0, R(3)\} + L_{2,1} = 2 + 4 = 6$ and $F_{4,1} = \max\{0, R(3)\} + L_{4,1} = 2 + 2 = 4$, so $F_{4,1} < F_{2,1}$. The ongoing transmission of packet $A_{1,1}$ is *not* preempted and will be completed at time 5, at which point packet $A_{4,1}$ will enter the service. At time 6 the finish number for packet $A_{3,2}$ is $F_{3,2} = \max\{0, R(6)\} + L_{3,2} = 3 + 1 = 4$. The current finish numbers are $F_{3,2} < F_{2,1}$ so $A_{3,2}$ enters the service at time 7, followed by $A_{2,1}$ which enters the service at time 8. Finally, at time 12 the finish number for the new packet $A_{3,3}$ is $F_{3,3} = \max\{0, R(12)\} + L_{3,3} = 6 + 1 = 7$ and it is transmitted at 12.

In summary, the order of arrivals is $\{A_{1,1}, A_{3,1}\}$, $\{A_{2,1}, A_{4,1}\}$, $A_{3,2}$, $A_{3,3}$ where simultaneously arriving packets are delimited by curly braces. The order of transmissions under packet-by-packet FQ is: $A_{3,1}$, $A_{1,1}$, $A_{4,1}$, $A_{3,2}$, $A_{2,1}$, $A_{3,3}$.

There is a problem with the above algorithm of fair queuing which the reader may have noticed besides that computing the finish numbers is no fun at all! At the time of a packet's arrival we know only the current time, not the current round number. As suggested above, one could try using the round number slope, Eq. (5.1), to compute the current round number from the current time, but the problem with this approach is that the round number slope is not necessarily constant. An FQ scheduler computes the current round number on every packet arrival, to assign the finish number to the new packet. Because the computation is fairly complex, this poses a major problem with implementing fair queuing in high-speed networks. Some techniques for overcoming this problem have been proposed, and the interested reader should consult [Keshav 1997].

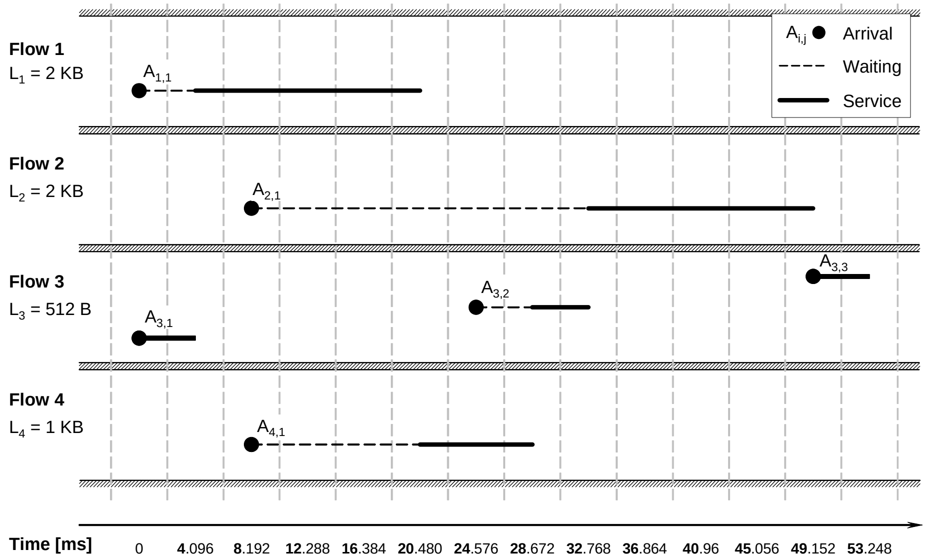

Figure 5-10: Time diagram of packet-by-packet FQ for Example 5.3.

5.1.3 Weighted Fair Queuing

Now assume that weights w_1, w_2, ..., w_n are associated with sources (flows) 1, 2, …, n, to reflect their relative entitlements to transmission bandwidth. As before, a queue is maintained for each source flow. Under weighted min-max fair share, flow i is guaranteed to obtain at least $C_i = \frac{w_i}{\sum w_j} C$ of the total bandwidth C. The bit-by-bit approximation of weighted fair queuing (WFQ) would operate by allotting each queue a different number of bits per round. The number of bits per round allotted to a queue should be proportional to its weight, so the queue with twice higher weight should receive two times more bits/round.

Packet-by-packet WFQ can be generalized from bit-by-bit WFQ as follows. For a packet of length $L_{i,j}$ (in bits) that arrives at t_a the finish number under WFQ is computed as

$$F_{i,j} = \max\left(F_{i,j-1},\, R(t_a)\right) + \frac{L_{i,j}}{w_i} \tag{5.3}$$

From the second term in the formula, we see that if a packet arrives on each of the flows i and k and $w_i = 2 \cdot w_k$, then the finish number for a packet of flow i is calculated assuming a bit-by-bit depletion rate that is twice that of a packet from flow k.

All queues are set to an equal maximum size, so the flows with the highest traffic load will suffer packet discards more often, allowing lower-traffic ones a fair share of capacity. Hence, there is no advantage of being greedy. A greedy flow finds that its queues become long, because its packets in excess of fair share linger for extended periods of time in the queue. The result is increased delays and/or lost packets, whereas other flows are unaffected by this behavior. In many cases delayed packets can be considered lost because delay sensitive applications ignore late packets.

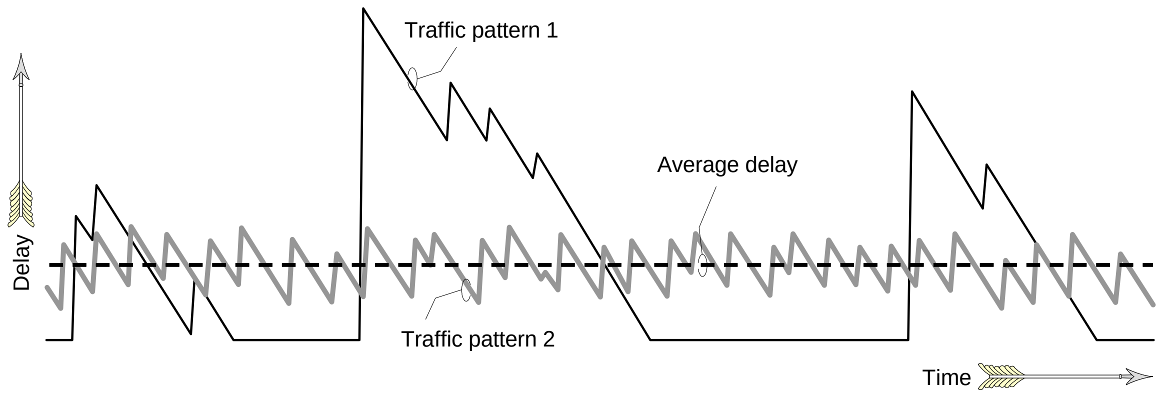

Figure 5-11: Different traffic patterns yield the same average delay.

The problem is created not only for the greedy source, but lost packets represent wasting network resources upstream the point at which they are delayed or lost. Therefore, they should not be allowed into the network at the first place. This is a task for policing.

5.2 Policing

So far, we saw how to distribute fairly the transmission bandwidth or other network resources using WFQ scheduler. However, this does not guarantee delay bounds and low losses to traffic flows. A packet-by-packet FQ scheduler guarantees a fair distribution of the resource, which results in a certain average delay per flow. However, even an acceptable average delay may have great variability for individual packets. This point is illustrated in Figure 5-11. Multimedia applications are particularly sensitive to delay variability (known as "jitter").

One idea is to regulate the number of packets that a particular flow can pass through the router per unit of time by using the "leaky bucket" abstraction (Figure 5-12). Imagine that we install a turnstile (ticket barrier) inside the router for monitoring the entry of packets into the service. To pass the turnstile, each packet must drop a token into the slot and proceed. Tokens are dispensed from a "leaky bucket" that can hold up to *b* tokens. If the bucket is currently empty, the packets must wait for a token. If a packet arrives at a fully occupied waiting area, the packet is dropped.[18]

Each flow has a quota that is characterized by simple traffic descriptors (Section 3.2.1):

Peak rate: this parameter constrains the number of packets that a flow can send over a very short period of time.

[18] There are many variations of the leaky bucket algorithm and different books introduce this abstraction in different way. In some variations, there are no tokens and the packets themselves arrive to the bucket and drain through a hole in the bucket. Because this is an abstraction anyway, I present it here in a way that I feel is the most intuitive. This does not affect the results of the algorithm.

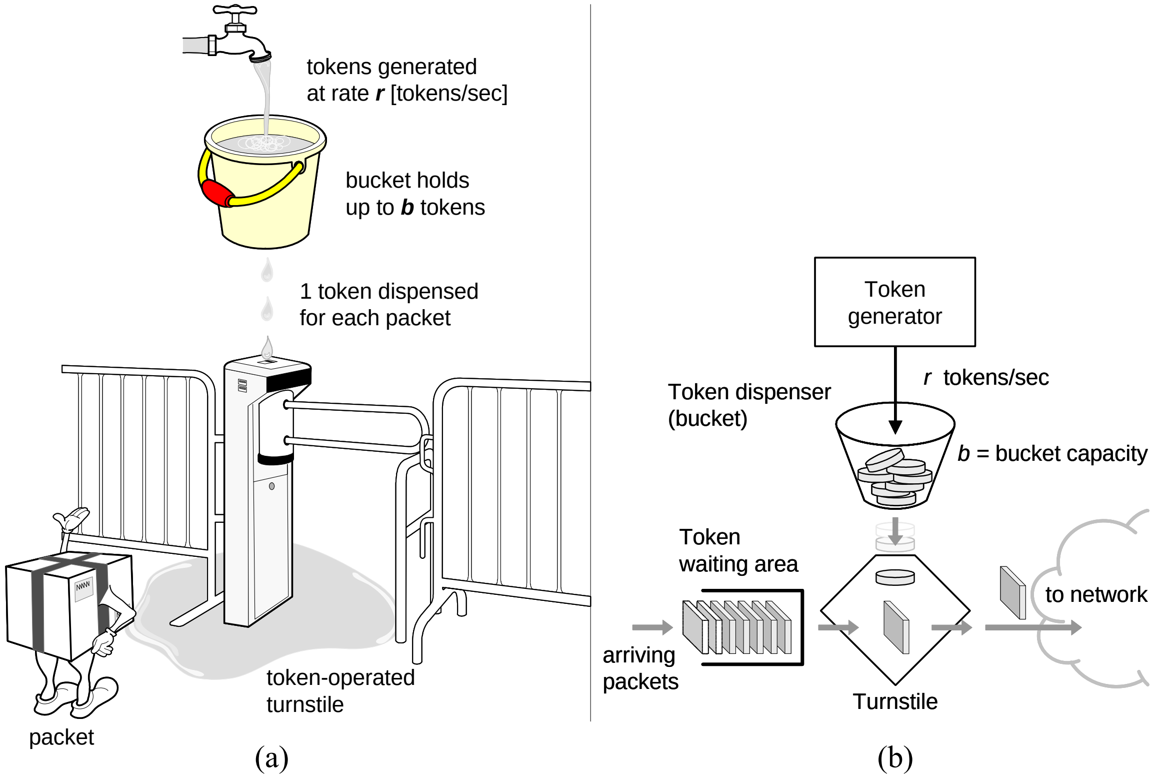

Figure 5-12: Leaky bucket.

Average rate: this parameter specifies the average number of packets that a particular flow is allowed to send over a time window. As discussed in Section 3.2.1, a key issue here is to decide the interval of time over which the average rate will be regulated.

Burst size: this parameter constrains the total number of packets (the "burst" of packets) that can be sent by a particular flow into the network over a short interval of time.

When a packet arriver at a router, it withdraws one token from the bucket before it is allowed to proceed. If the bucket is empty, the packet must wait for a token to appear.

When the packets are all the same size, this algorithm can be used as described. However, when variable-sized packets are being used, it is often better to allow a fixed number of bytes per token, rather than just one packet.

5.3 Active Queue Management

Packet-dropping strategies deal with the case when there is not enough memory to buffer an incoming packet. The simplest policy is to drop the arriving packet, known as *drop-tail* policy.

Active Queue Management (AQM) algorithms employ more sophisticated approaches. Routers with AQM detect congestion before the queue overflows and notify the source that the congestion is about to happen. One of the most widely studied and implemented AQM algorithms is the *Random Early Detection* (RED) algorithm that uses *implicit feedback* by dropping packets. A recent approach, called *Explicit Congestion Notification* (ECN) uses *explicit feedback* by marking packets instead of dropping them.

5.3.1 Random Early Detection (RED)

A router that implements RED uses two threshold values to mark positions in the queue: $Threshold_{min}$ and $Threshold_{max}$. A simplified description of RED operation follows. When a new packet arrives, its disposition is decided by these three rules:

1. If the queue currently contains fewer than $Threshold_{min}$ packets, the new packet is enqueued.
2. If the queue contains between $Threshold_{min}$ and $Threshold_{max}$ packets, the new packet is considered for enqueuing or dropping by generating a random number and evaluating its value.
3. If the queue currently contains more than $Threshold_{max}$ packets, the new packet is dropped.

where $0 \le Threshold_{min} < Threshold_{max} \le$ BufferSize, and the value 0 represents the head of the queue. Therefore, instead of waiting until the queue overflows, RED starts randomly dropping packets as congestion increases. The process is somewhat more complex, as described next.

We know that TCP often sends segments in bursts, depending on the congestion window size (Section 2.2). A burst represents a spike in traffic intensity that may last only temporarily while most of the time traffic is low-intensity. In other words, the queue may be most of the time empty, and a temporary spike does not represent congestion. Therefore, should like avoid dropping packets from a burst when queue length is greater than $Threshold_{min}$. (Of course, packets are always dropped if the queue capacity is exceeded.)

To better capture the notion of congestion and accommodate for bursty traffic, the router does *not* consider the *instantaneous* length of the queue. Instead, the router considers the *average* length of the queue when applying the three rules described above (Figure 5-13(a)). The average queue length is computed continuously using Exponential Weighted Moving Average (EWMA). This is the same method used by TCP for RTT estimation (Section 2.1.2) and by jitter buffer (Section 3.3.1). That is, at any time *t* when a new packet arrives and tries to join this queue, we compute

$$AverageQLen(t) = (1-\gamma)\cdot AverageQLen(t-1) + \gamma\cdot MeasuredQLen(t) \tag{5.4}$$

where γ denotes a value between 0 and 1. If γ is small, the average stays close to long-term trend and does not fluctuate for short bursts. This parameter should be determined empirically, and a recommended value is $\gamma = 0.002$.

When the average queue length is between $Threshold_{min}$ and $Threshold_{max}$, RED drops an arriving packet with an increasing probability as the average queue length increases (Figure 5-13(b)). For a given *AverageQLen*, we calculate the probability of packet drop as

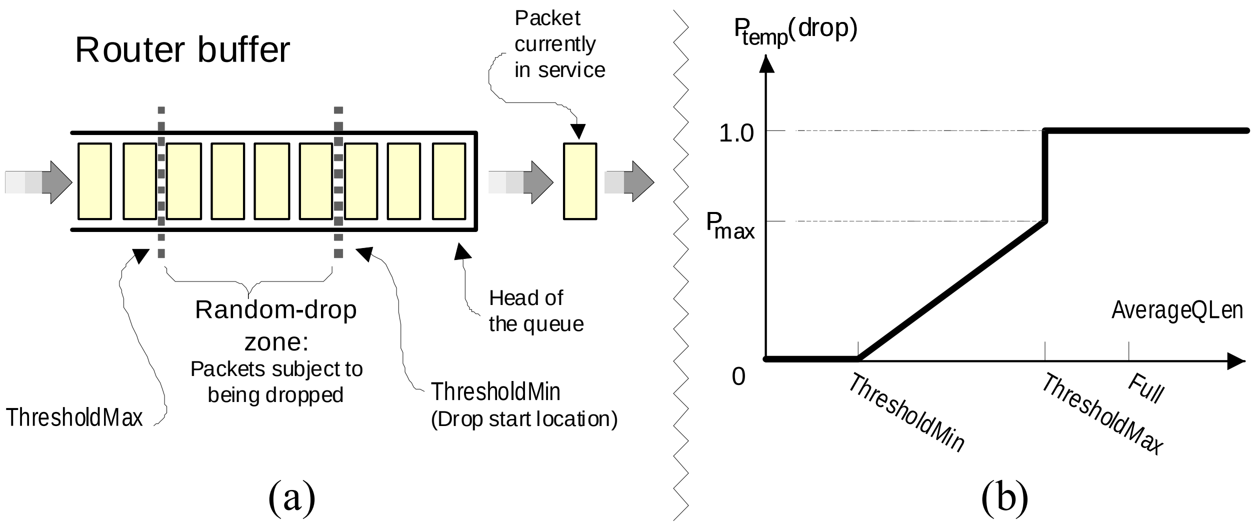

Figure 5-13: (a) RED thresholds on a FCFS (or, FIFO) queue. (b) Packet drop probability function for RED.

$$P_{\text{temp}}(AverageQLen) = P_{\max} \times \frac{(AverageQLen - Threshold_{\min})}{(Threshold_{\max} - Threshold_{\min})} \tag{5.5}$$

Research has suggested that RED works best when the probability function transitions smoothly at $Threshold_{max}$, i.e., for $P_{\max} = 1$. In addition to the average queue length, the drop probability also depends on the time elapsed since the last packet was dropped. Instead of actual time, the algorithm uses a proxy in terms of the number of newly arriving packets (variable: *count*) that have been queued (not dropped) while *AverageQLen* has been between the two thresholds. Therefore, given an *AverageQLen* the actual probability of a packet being dropped is computed as

$$P(AverageQLen) = \frac{P_{\text{temp}}(AverageQLen)}{1 - count \times P_{\text{temp}}(AverageQLen)} \tag{5.6}$$

We can observe that *P* increases as *count* increases. This helps make packet drops more evenly distributed and avoid bias against bursty traffic.

RED is intended to work primarily with TCP sources. RED is a queue management technique that attempts to provide equitable access to an FCFS system. The source that transmits at the highest rate will suffer from higher packet-dropping rate than others will. As a result, this source will reduce its transmission rate more, which yields more uniform transmission rates across the sources and more equitable access to the buffer resource. Under RED, TCP connections will experience randomly dropped packets (rather than synchronously when the buffer becomes full). Therefore, TCP senders will back off at different times. This behavior avoids the *global synchronization effect* of all connections and maintains high throughput in the routers. Both analysis and simulations shows that RED works as intended. It handles congestion, avoids the synchronization that results from drop-tail policy, and allows short bursts without dropping packets unnecessarily. It is also important that RED can control the queue length irrespective of endpoint sender cooperation. The IETF now recommends that routers implement RED.

5.3.2 Explicit Congestion Notification (ECN)

Explicit Congestion Notification (ECN) allows end-to-end notification of network congestion without dropping packets. It requires support by the underlying network (i.e., routers). ECN is an optional feature that is only used when both endpoints support it and have it activated.

We have seen that Random Early Detection (RED) mechanism drops packets when it senses potential congestion (Section 5.3.1). Unlike RED, instead of dropping a packet, an ECN-aware router may set a mark in the packet's IP header in order to signal that congestion is about to happen. The receiver of the packet is the first to learn about the potential congestion and echoes the congestion indication to the sender. The receiver uses the next acknowledgement packet to inform the sender about the impending congestion, which must react as if a packet was dropped, i.e., by reducing its congestion window (Section 2.2).

ECN uses two bits in the IP header of packets from ECN-capable transports to allow routers to record congestion, and uses two bits in the TCP header to allow the TCP sender and receiver to communicate. The two bits in the IP header are taken from the 8-bit type of service (TOS) field (Figure 1-36) and are called the *Congestion Experienced* (CE) codepoint. These are bits 6 and 7 (rightmost) of the TOS field, and a router can choose to set either bit to indicate congestion. The reason for using two bits is to increase the robustness of the mechanism. The four different codepoints are as follows:

00: Transport not ECN-capable - Non-ECT

10: ECN capable transport - ECT(0)

01: ECN capable transport - ECT(1)

11: Congestion encountered - CE

The two bits in the TCP header are taken from the 6-bit unused field (Figure 2-2).

ECN uses the same mechanism as RED (Section 5.3.1) to detect an impending congestion: it monitors the average queue size.

5.4 Multiprotocol Label Switching (MPLS)

Problems related to this section: Problem 5.9 → ?

Multiprotocol Label Switching (MPLS) is essentially a mechanism for creating and using special paths, known as "tunnels," in IP networks. We know that IP forwards data packets on a packet-by-packet basis, where each packet is forwarded independently of any other packet. It is said to be connectionless packet forwarding. MPLS allows forwarding packets on a flow-by-flow basis, so that all packets belonging to a given traffic flow are forwarded in the same manner and along the same network path (or, tunnel). In this sense, MPLS supports connection-oriented packet forwarding and the fixed forwarding paths (tunnels) represent "virtual circuits" in the network.

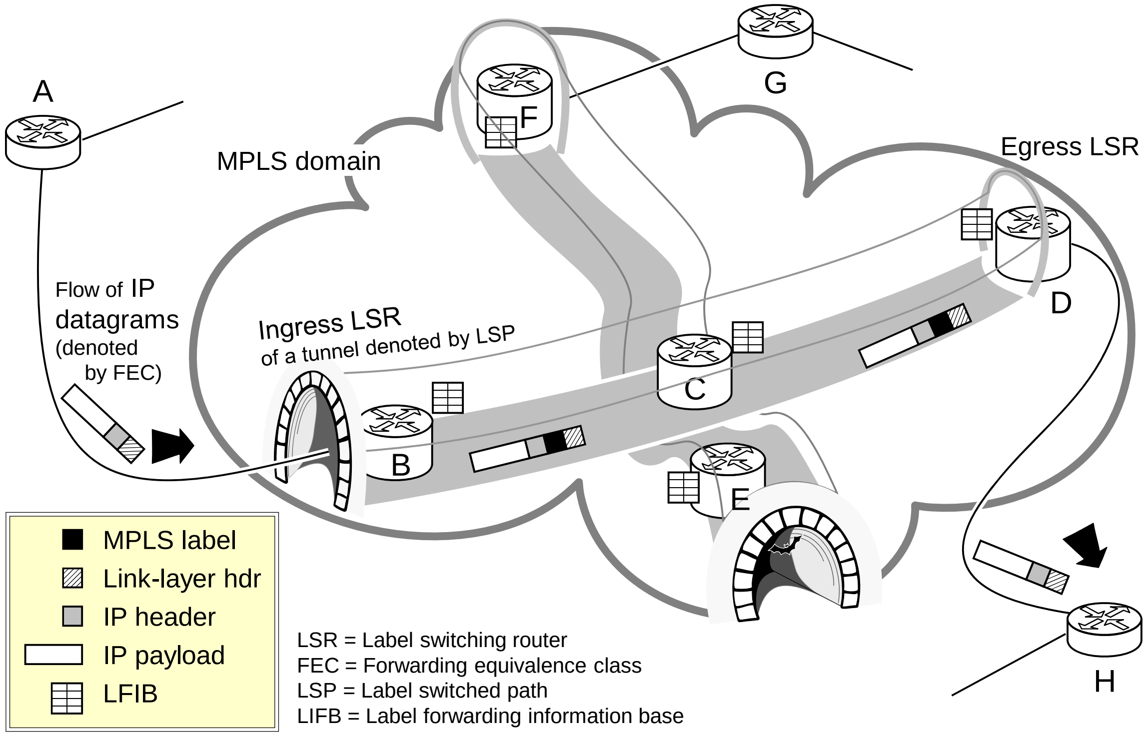

Figure 5-14: MPLS operation.

MPLS relies on IP addresses and IP routing protocols to set up the paths/tunnels. An MPLS-enabled router is known as a **Label Switching Router (LSR)**. A set of LSRs where each LSR is reachable from any other LSR via some other LSRs in the same set is called an **MPLS domain**. In other words, an MPLS domain is formed by a contiguous network. A label switching router forwards packets by examining a short, fixed-length MPLS label (Figure 5-14). The label represents a given traffic flow, and all the packets belonging to this flow should be forwarded along the path/route/tunnel associated with this label. A traffic flow is called a **Forwarding Equivalence Class (FEC)**, also a Functional Equivalence Class, and this is a group of IP packets that are forwarded in the same manner (i.e., along the same path, with the same forwarding treatment). In other words, An FEC is a set of packet flows with common cross-core forwarding-path requirements. A sequence of routers that form a path along which a given FEC flow is forwarded forms a tunnel, which is known as a **Label Switched Path (LSP)**. For example, one such path is formed by routers *B*, *C*, and *D* in Figure 5-14. Each LSP tunnel is unidirectional (one-way), starting with an ingress LSR, going through intermediate LSRs, if any, and ending with an egress LSR. If data needs to travel in the opposite direction as well, which is usually true, then a separate one-way tunnel must be built in the opposite direction. To summarize, an FEC is uniquely associated with an LSP. Each pair of routers agrees on a label independently of other router pairs along an LSP. Therefore, each label has *local scope* and different segments of an LSP tunnel may be represented by different MPLS labels. Why this is so, will be explained later.

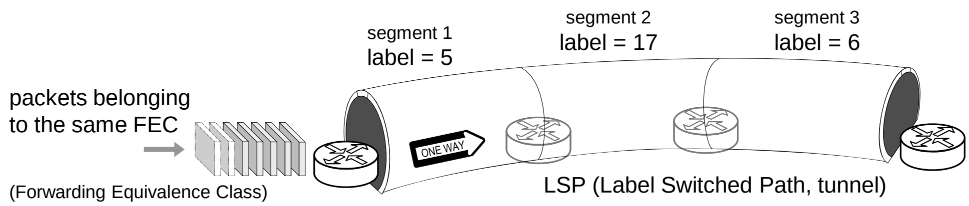

Why MPLS:

• Use switching instead of routing

• IP Traffic Engineering (TE): ability to specify routes based on resource constraints, rather than on distance (shortest path) only (Section 5.4.3). MPLS adds the ability to forward packets over arbitrary non-shortest paths, using constraint-based routing.

• Virtual Private Networks (VPNs, Section 5.4.4): Controllable tunneling mechanism emulates high-speed "tunnels" between IP-only domains.

• Route protection and restoration

A key reason for initial MPLS development was the promise of ultra-fast packet forwarding: Longest prefix match is (was) computationally expensive (Section 4.1.3); Label matching was seen as much less computationally expensive. However, with the emergence of gigabit IP routers capable of IP forwarding as fast as any MPLS-capable router performs label switching the speed advantage has diminished (although not disappeared, because label switching still can be implemented in a much simpler hardware). Currently, MPLS key strengths are seen in Traffic Engineering and Virtual Private Networks—capabilities critical to network service providers who need to better manage resources around their backbones or need to offer VPN services.

Notice that MPLS is a set of protocols (specified by IETF), rather than a single protocol. The key protocols are Label Distribution Protocol and Link Management Protocol. Several existing IP protocols are also adapted to work with MPLS.

5.4.1 MPLS Architecture and Operation

MPLS is located between the link-layer and network-layer protocols (Figure 5-15), so it is referred to as a layer 2.5 protocol (in the OSI reference architecture, where Link is layer 2 and Network is layer 3). MPLS can run over different link-layer technologies, such as Ethernet or PPP (Section 1.5). The protocol-identifier field of the link-layer header should identify the payload as an MPLS frame. For example, unique PPP code points (carried in the Protocol field, Figure 1-56) identify the PPP frame's contents as an MPLS frame. The value of the PPP Protocol field for MPLS unicast is hexadecimal 0x0281. A similar encapsulation scheme is used when transmitting over Ethernet, where the payload is identified as an MPLS frame with unique Ether-Types (Figure 1-59(a)) or LLC frame's DSAP addresses (Figure 1-59(b)). The value of Ether-Type for MPLS unicast is hexadecimal 0x8847.

Different **Forwarding Equivalence Classes (FECs)** designate different classes of service or service priorities. Each MPLS-capable router (label switching router, LSR) keeps a list of labels that correspond to different FECs on each outgoing link. All packets belonging to the same FEC have the same MPLS label value. However, not all packets that have the same label value belong

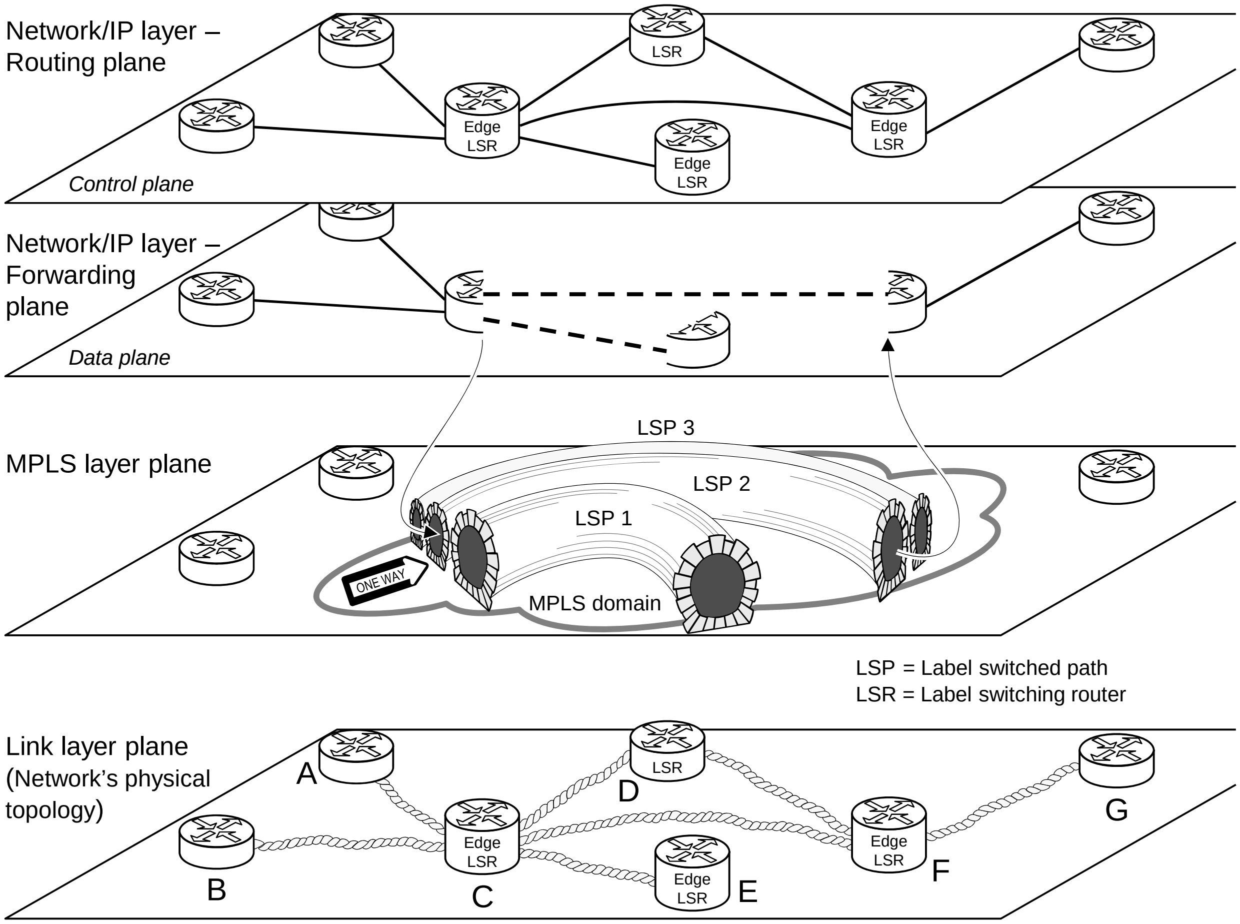

Figure 5-15: Protocol layering of an MPLS network.

to the same FEC. This fact will become clear later, as we see that FEC is determined by the label value and experimental bits (Exp) of the MPLS header. By default, packet's FEC is determined by its destination IP address. Other classification parameters include source IP address, IP protocol type, TOS field of the IP header (Figure 1-36), and TCP/UDP port numbers. The packet's arrival port may also be considered a classification parameter. Multicast packets belonging to a particular multicast group also form a separate FEC.

An edge-LSR terminates and/or originates LSPs (label switched paths) and performs both label-based forwarding and conventional IP forwarding functions. The edge-LSR converts IP packets into MPLS packets, and MPLS packets into IP packets. On ingress to an MPLS domain, an LSR accepts unlabelled IP packets and creates an initial MPLS frame by pushing a shim header between link-layer and network-layer headers (Figure 5-16). A special table in the ingress LSR, known as **Label Forwarding Information Base (LFIB)**, matches the FEC to the label. LFIB is an MPLS equivalent for the forwarding table of the IP protocol. On egress, the edge LSR terminates an LSP by popping the top MPLS stack entry, and forwarding the remaining packet based on rules indicated by the popped label (e.g., that the payload represents an IPv4 packet and should be processed according to IP forwarding rules).

Edge LSRs provide the interface between external networks and the internal label-switched paths, and core/intermediate LSRs provide transit service in the middle of the network. An intermediate

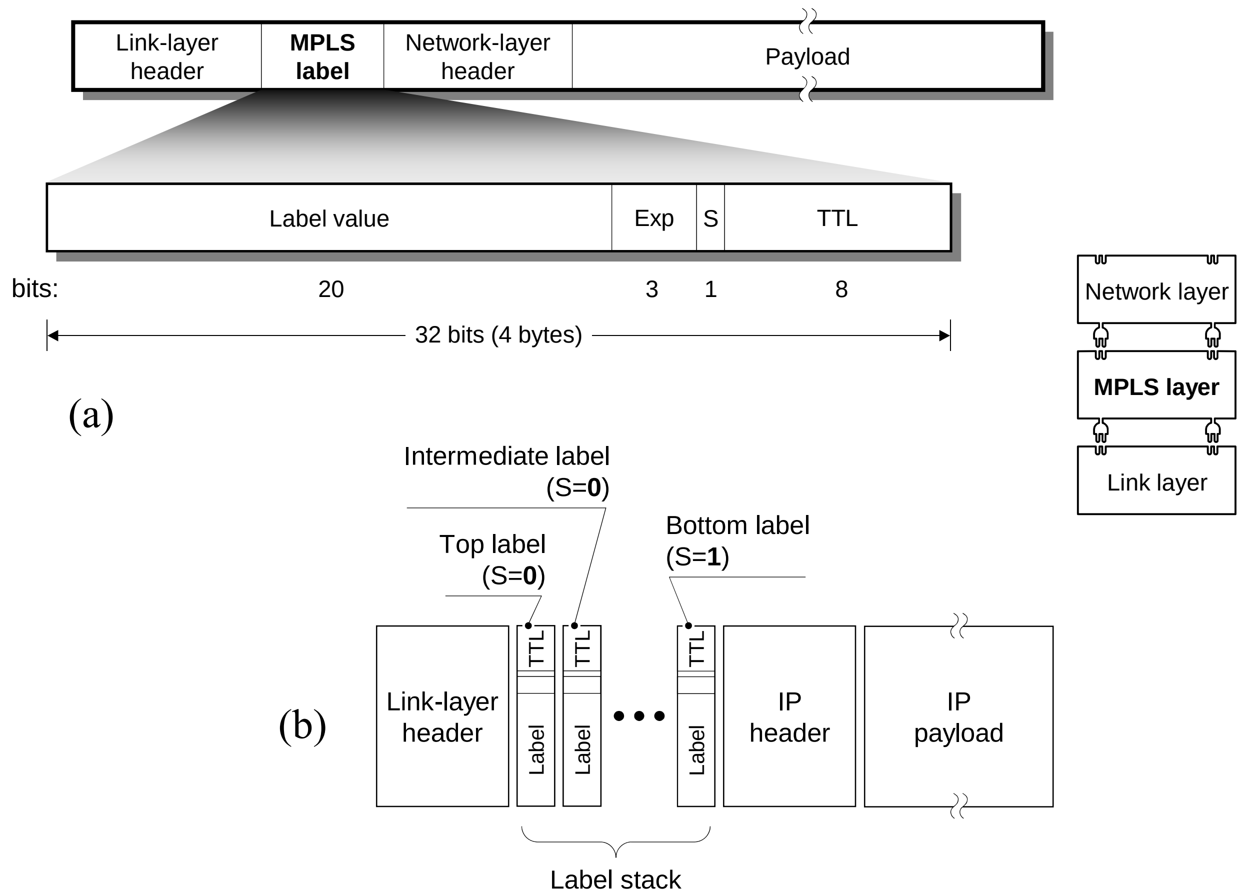

Figure 5-16: (a) MPLS label ("shim" header) format and its relative placement in IP packets. (b) The placement of the label stack; notice that the S bit is set only for the bottommost label.

LSR examines incoming MPLS packets, looks up and follows the packet's label instructions, and then forwards the packet according to the instructions. In general, the LSR performs a label swapping function.

Paths or routes are established between the edge LSRs via intermediate LSRs. These paths are called **Label Switched Paths (LSPs)**. The LSPs are designed for their traffic characteristics. The traffic-handling capability of each path is calculated. These characteristics can include peak traffic load, inter-packet variation, and dropped packet percentage calculation.

MPLS Labels

Figure 5-16(a) shows the MPLS label for frame-based packets (e.g., Ethernet, PPP). MPLS label is also known as "shim" header. The meaning of the fields is as follows:

Label value: A number representing a forwarding equivalence class (FEC) on a given outgoing link. The label has a local scope limited to a network link, which means that a link may support up to one million (2^{20} = 1,048,576) distinct labels.

Exp: experimental bits identify the class of service (or QoS).

Visit http://en.wikipedia.org/wiki/Optical_switch for information about optical switches

Bottom-of-stack bit (S): value "1" indicates that this label header is the bottom label in the stack; otherwise, it is set to zero. The stack is a collection of labels that represent a hierarchy of tunnels created over a particular outgoing link. The stack can have unlimited depth, although it is rare to see a stack of four or more labels.

TTL: time-to-live counter that has the same function as the TTL found in the IP header (Figure 1-36), which is to prevent packets from being stuck in a routing loop. The TTL counted is decreased by 1 at each hop, and if the value reaches 0, the packet is discarded. Special processing rules are used to support IP TTL semantics, as described below.

The label value at each hop is a local key representing the next-hop and QoS requirements for packets belonging to each FEC. In conventional routing, a packet is assigned to an FEC at each hop (i.e., forwarding table look-up, Section 4.1.3). Conversely, in MPLS it is only done once at the ingress of the MPLS domain. At the ingress LSR, a packet is classified and assigned an FEC/label. Packet forwarding in the MPLS domain is performed by swapping the label.

The label stack entries appear after the link-layer header and before the network-layer header (Figure 5-16(b)). The label that was last pushed on the stack (newest) is called the *top label*, and it is closest to the link-layer header. The label that was first pushed on the stack (oldest) is called the *bottom label*, and it is closest to the network-layer header. The edge LSR that pops up the bottommost label is left with a regular IP packet, which it passes up to the network layer and IP forwarding is used to move the packet onwards.

Label Bindings, LIB, and LFIB

The ordinary IP control plane builds and maintains the routing table, also known as RIB (Routing Information Base). Routing table is just built here, but forwarding decisions are made in the forwarding or data plane (top of Figure 5-15). Forwarding decisions are made based on the forwarding table, also known as FIB (Forwarding Information Base), which is derived from the routing table. In case of MPLS, the equivalent data structures are LIB (label information base) and LFIB (label forwarding information base). The prefixes-to-label bindings are built and stored in the LIB, control plane, which is then used to create the LFIB data or forwarding plane. The lookups are actually done in the LFIB, not the LIB (as for IP, in the FIB and not the RIB). Their relationship is illustrated in Figure 5-17.

Label Forwarding Information Base (LFIB) is a data structure and way of managing forwarding in which destinations and incoming labels are associated with outgoing interfaces/ports and labels. The LFIB resides in the data plane and contains a local-label-to-next-hop label mapping along with the outgoing port, which is used to forward labeled packets.

In summary, the routing table is built by the routing protocols in the IP control plane. Similarly, LIB (label information base) is built by a label distribution protocol in the MPLS control plane. LIB contains only labels, no routes (i.e., LSP tunnels). The IP forwarding table is derived in the IP data plane from the routing table. Correspondingly, the MPLS LFIB is derived from LIB in the MPLS data plane. LFIB contains bindings between labels and LSPs. A Labeled Packet is always looked up in LFIB (not in LIB!) and an IP Packet is always looked up in forwarding table (not in routing table!). However, the process is somewhat more complex for edge LSRs. On the ingress LSR, the lookup is performed against the combined IP forwarding table and LFIB, as described in the next section. In the core (intermediary LSRs), the lookup is performed only against the LFIB.

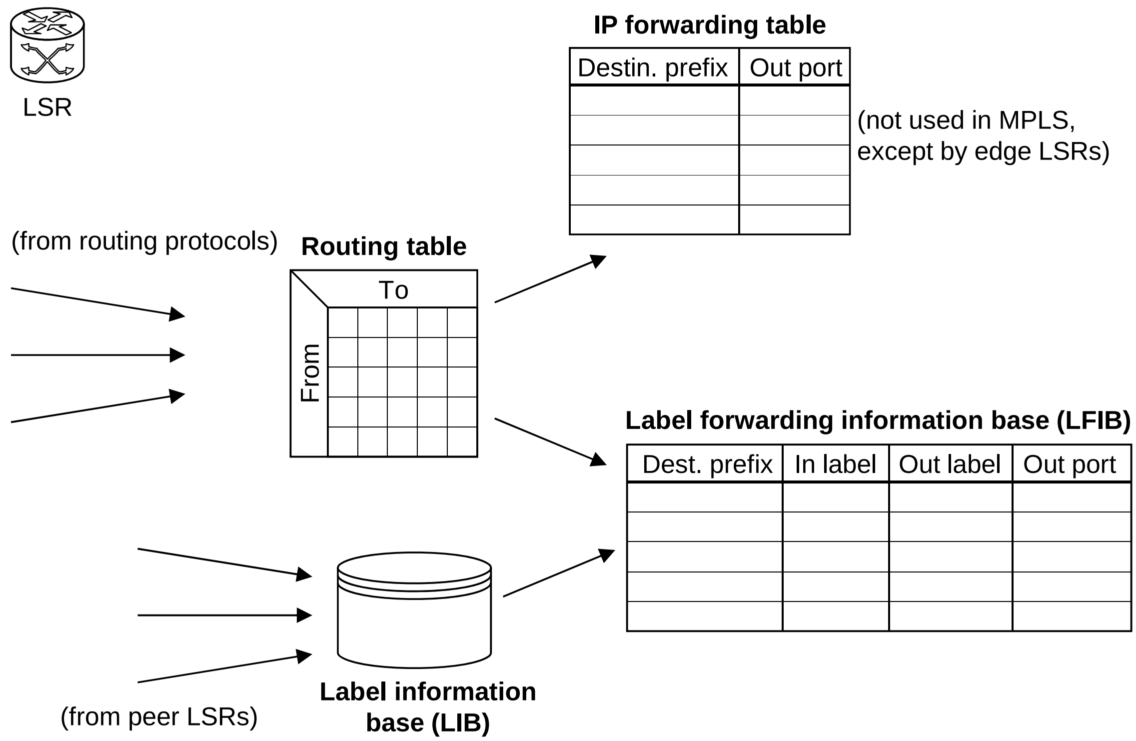

Figure 5-17: Relationship of LIB (label information base) and LFIB (label forwarding information base).

On the egress LSR, the lookup is performed against the IP forwarding table if there was only a single label in the stack and this label was popped by the penultimate hop; otherwise, the LFIB is looked up.

More precisely, table to lookup into is determined by the link-layer header Ether-Type or PPP Protocol field. The protocol identifier in the link-layer header tells the router what type of packet is coming in and therefore which table to look in:

- 0x0800 → IPv4: Lookup in the IP forwarding table
- 0x8847 → MPLS Unicast: Lookup in the MPLS LFIB (label forwarding information base)
- 0x8848 → MPLS Multicast: Lookup in the MPLS LFIB (label forwarding information base)

Next, we consider how MPLS routers (LSRs) build and utilize label-forwarding tables.

Forwarding Labeled Packets

Initially, all routers start with empty routing and forwarding tables and label-bindings. We assume that regular IP routing protocols run first and build regular IP routing tables, or RIBs (routing information bases). MPLS builds label-binding tables based on regular IP routing tables, using label distribution protocols (described in Section 5.4.2). Label bindings can also be configured manually, particularly for the purposes of Traffic Engineering, but this is a tedious and error-prone task, so even here it is preferred to use label distribution protocols combined with constraint-based routing protocols. To illustrate how MPLS-capable routers (LSRs) forward

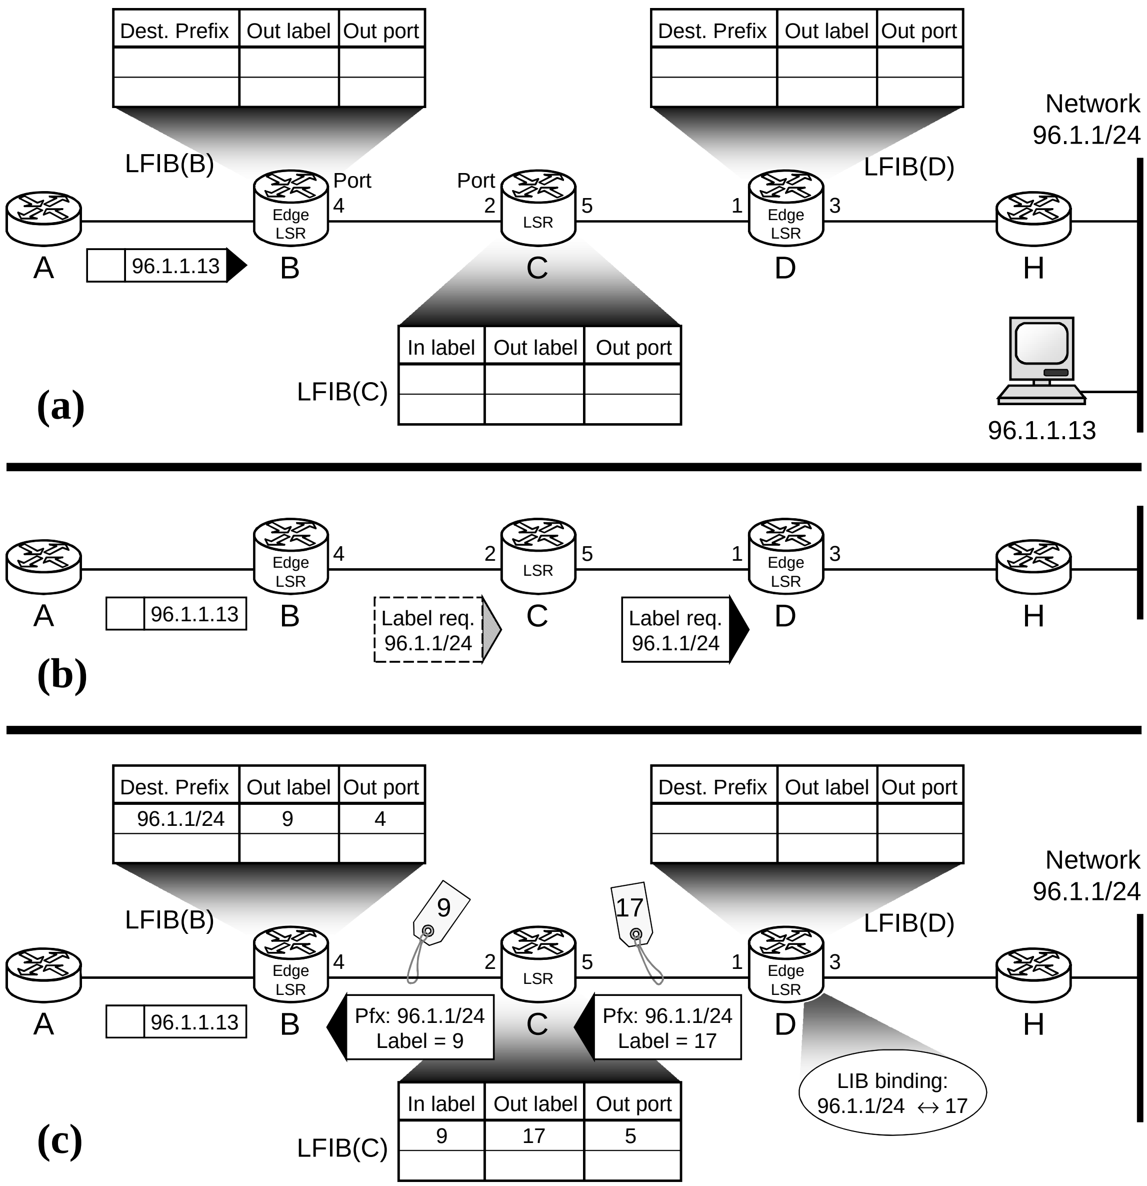

Figure 5-18: LSP path setup. (a) Packet arrives from router *A* to *B* towards a host in *H*'s network. (b) LSR *B* sends label request towards the destination. (c) LSR *D* is the edge router, so it replies with label 17, then LSR *C* selects its own label as 9 and replies to *B*.

labeled packets, let us assume the simplest scenario where label bindings are derived from the hop-by-hop information in IP routing tables.

Consider the example in Figure 5-18, which illustrates one of the tunnels from Figure 5-14. Here, edge LSR *B* receives a packet with destination IP address `96.1.1.13`. LSR *B* has the corresponding network prefix `96.1.1/24` in its routing table, but does not have the label binding in its LFIB (label forwarding information base). To obtain a label for this prefix, *B* uses a label distribution protocol and sends a message downstream (to *C* and on to *D*) requesting a label for prefix `96.1.1/24`. When edge LSR *D* receives the request, it knows that itself is the egress LSR of the new tunnel (LSP, label switched path). Therefore, *D* selects a label value not used for any other LSP and sends a response message using the label distribution protocol. In our

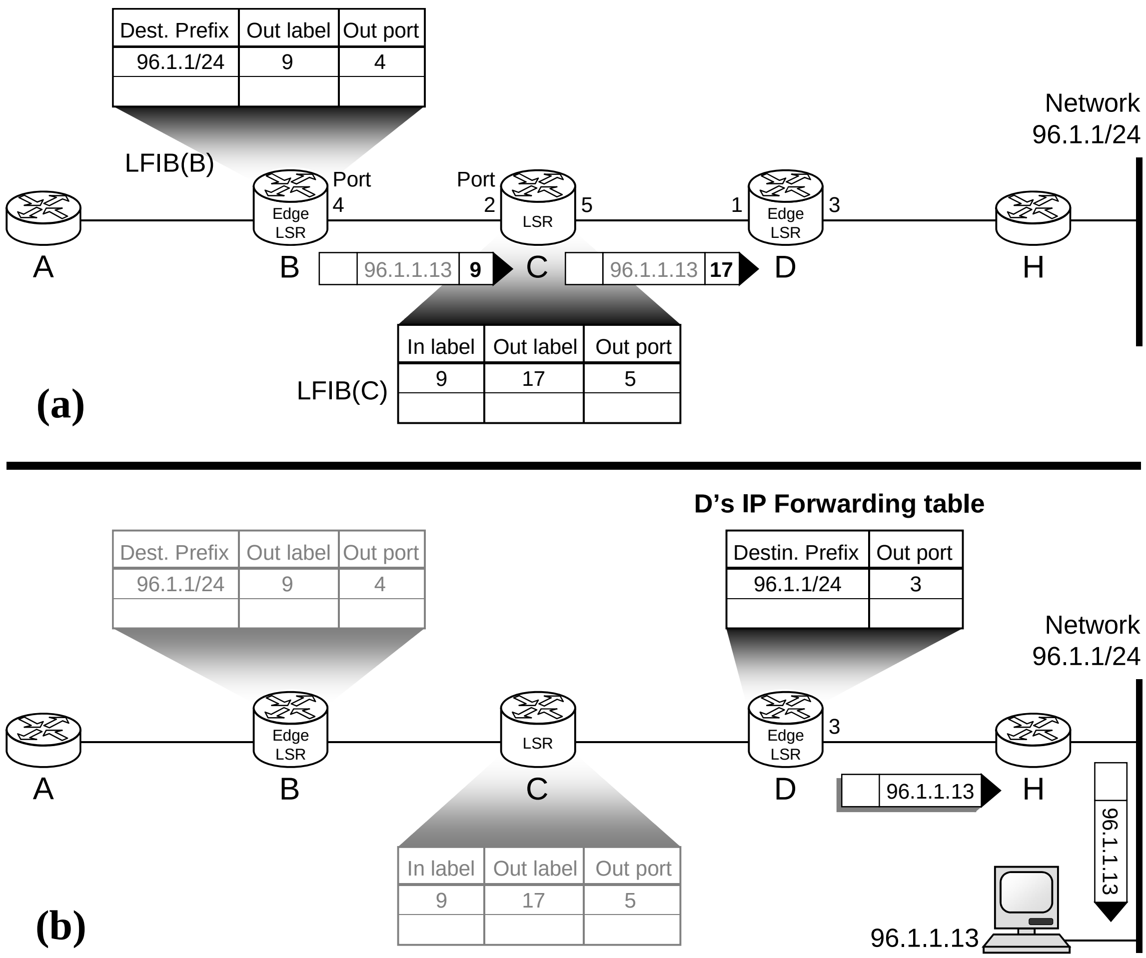

Figure 5-19: Forwarding labeled packets (*continued from* Figure 5-18). (a) Within MPLS domain, data packet is forwarded based on its MPLS label. (b) Outside MPLS domain, data packet is forwarded based on its destination IP address.

example, *D* selected the label value 17 and stored the label 17 binding for prefix `96.1.1/24` in its LIB (label information base), *not* LFIB (label forwarding information base)! *D*'s LFIB remains empty because LSR *D* does not use LFIB to forward packets from this tunnel. *D* is the egress of the tunnel and to forward packets towards *H* (which is not an LSR and is not MPLS capable), *D* will use conventional IP forwarding.

When *C* receives the label response from *D*, in general it will need to assign a different label for the next segment of the same LSP tunnel, because it may be already using label 17 for another LSP. Remember, routers are at crossroads of many different paths and these paths are established at unpredictable times. In our example (Figure 5-18), *C* selects the label value 9 for the upstream segment of LSP. Because *C* is an intermediate LSR, it does not store prefixes in its LFIB; *C* might even not be IP-capable. Rather, *C* needs just the incoming and outgoing labels. The incoming label value is 9 (will be received in MPLS packets from *B*) and the outgoing label value is 17 (will be sent in MPLS packets to *D*). In other words, the intermediate LSR *C* performs *label swapping*. Unlike intermediate LSRs, edge LSRs do not perform label swapping. Each edge LSR must understand both IP and MPLS, and its LFIB (label forwarding information base) may have

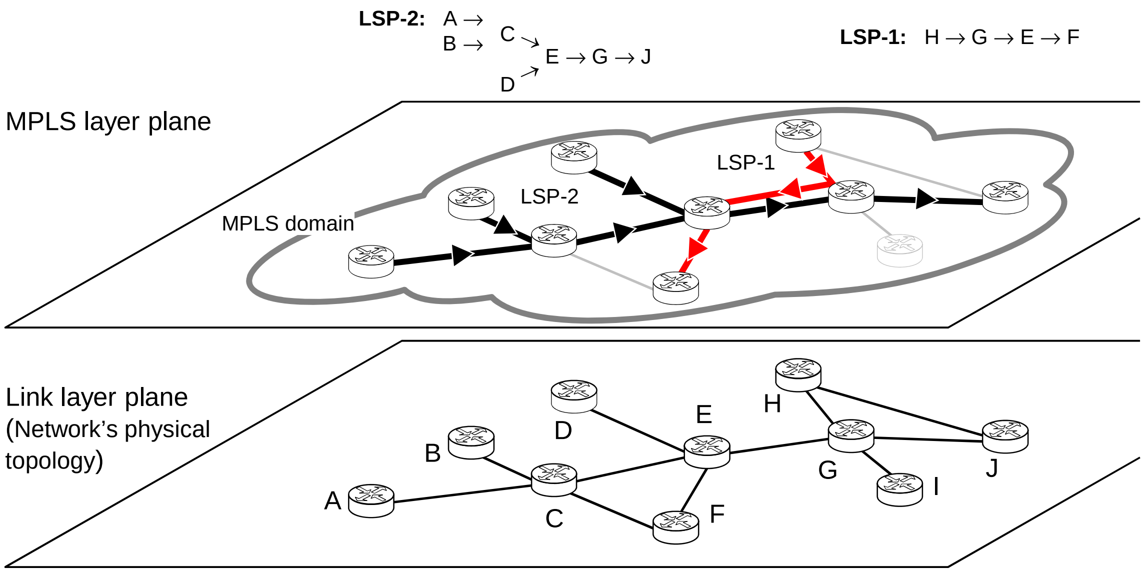

Figure 5-20: LSP topologies (or Point-of-Presence/PoP designs). LSP-1: Unique ingress and egress LSR. LSP-2: Multiple ingress LSRs, unique egress LSR.

different format. Notice also that some LSRs may play both roles: edge and intermediate, for different LSP tunnels.

Continuing with the example in Figure 5-18, the data packet has been sitting in LSR *B* while the LSP setup process took place. Once liable bindings become available, *B* will forward all data packets from this flow using label switching. An example for the first packet is shown in Figure 5-19. LSR *B*, as the ingress router of this LSP, inserts an MPLS label with value 9 (which it obtained from *C*), and sends this packet on output port 4 towards *C* (Figure 5-19(a)). When *C* receives the packet, it performs *label swapping*: For the given input label 9, *C* looks up its LFIB and finds that the outgoing label is 17 (which it received from *D*), swaps the packet's label to 17 and sends it on output port 5 towards *D*. When *D* receives the packet, it looks up the incoming label (17) and recognizes that itself is the egress LSR for this LSP. Therefore, *D* strips off the MPLS label and forwards the packet towards the next hop *H* using conventional IP forwarding (Figure 5-19(b)).

Topology of LSPs

The design of PoPs (Points-of-Presence) for all backbone IP networks, including MPLS networks, is constrained by the choice of access link type(s) to be supported for the customers of the network and the choice of core link type(s) to be used in the backbone network. Based on PoP designs, LSP (Label Switched Path) trees can be classified as these topology types (Figure 5-20):

- **Unique ingress and egress LSR:** In this case, a point-to-point path through the MPLS domain is set up. An example is LSP-1 in Figure 5-20, from the ingress LSR *H*, through the intermediate LSRs *G* and *E* towards the egress LSR *F*.

- **Multiple ingress LSRs, unique egress LSR:** In this case, LSP forms a multipoint-to-point tree topology. This happens when traffic assigned to a single FEC arises from different sources. An example is LSP-2 in Figure 5-20, where traffic assigned to a single FEC enters at three

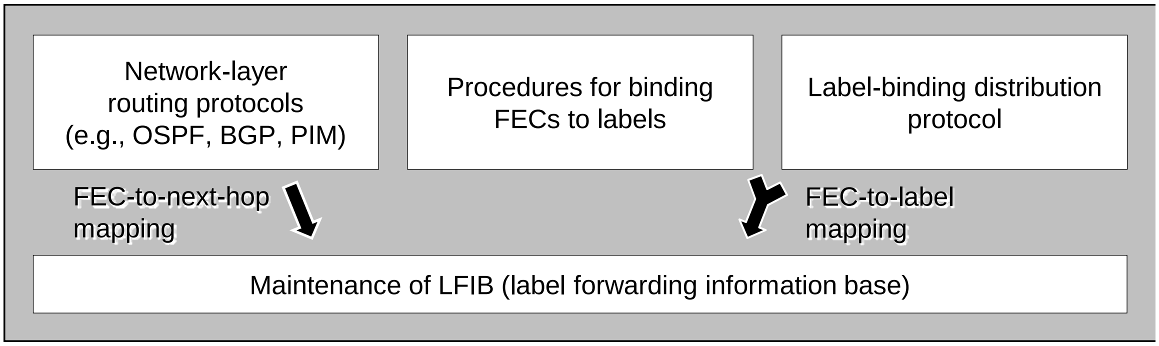

Figure 5-21: The components of the control plane of an LSR perform LFIB construction.

different ingress LSRs: *A*, *B*, and *D*. The branches from *A* and *B* join at LSR *C*, then this branch joins with *D* at *E*, and the final two hops through *G* to the egress LSR *J* are shared.

- **Multicast:** In this case, multicast traffic is carried over the MPLS domain from a single ingress LSR to multiple egress LSRs. The multicast LSP is determined by the multicast tree constructed by the multicast routing protocol (Section 3.3.2).

In principle, an ISP backbone network could configure a separate LSP to carry each class of traffic (FEC) between each pair of edge LSRs. A more practical solution is to merge LSPs of the same traffic class to obtain multipoint-to-point flows that are rooted at an egress LSR. An example is LSP-2 in Figure 5-20. The LSRs serving each of these flows would be configured to provide the desired levels of performance to each traffic class.

5.4.2 Label Distribution Protocols

Setup of LSPs (Label Switched Paths) is done by a process of *label distribution*. Label distribution may be based on information obtained from conventional *hop-by-hop routing* protocols, or it may use *explicit routing* over non-shortest paths. Label distribution protocol dynamically establishes an LSP tree between all the edge LSRs for each identifiable FEC. Requirements for a label distribution protocol include per-hop traffic differentiation capabilities, the ability to route traffic over non-shortest paths, and the ability to dynamically signal (or provision) QoS and path information across a network of routers or switches. There are many similarities between conventional routing protocols and label distribution protocols for MPLS. A key difference is the MPLS capability for explicit non-shortest-path routing.

The control plane of an LSR performs the following functions (Figure 5-21):

1. Create bindings between FECs and labels.
2. Inform the adjacent LSRs of the bindings it created (using a label distribution protocol).
3. Use information received from the adjacent LSRs to construct and maintain the forwarding table (LFIB) used by the MPLS label switching.

Label Distribution

In general, label bindings between two LSRs can be distributed by either a downstream LSR or an upstream LSR. MPLS architecture requires *downstream label distribution*: label-bindings must

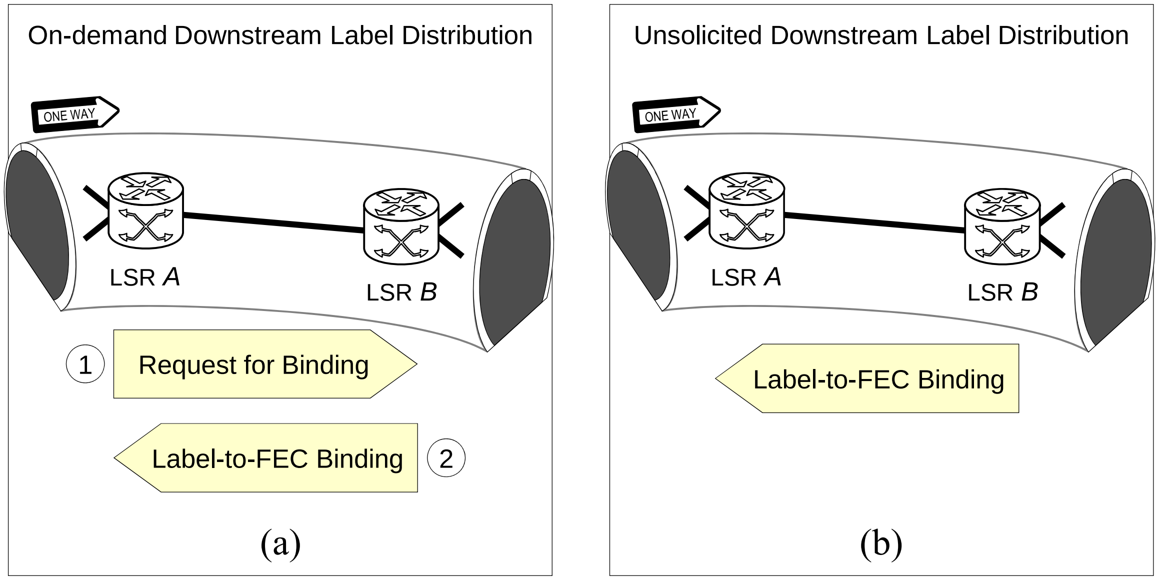

Figure 5-22: Methods for MPLS downstream label distribution.

be distributed in the direction from a downstream LSR to an upstream LSR. There are two methods for downstream label distribution:

- **On-demand Downstream Label Distribution:** In this case, a downstream LSR distributes a label binding in response to an explicit request from an upstream LSR (Figure 5-22(a)). An upstream LSR *A* recognizes a downstream LSR *B* as its next-hop for an FEC and sends a request to LSR *B* for a binding between the FEC and a label. If LSR *B* recognizes the FEC and has a next hop for it, LSR *B* creates a binding and replies to LSR *A*. Both LSRs then have a common understanding. This process is also illustrated in Figure 5-18.

- **Unsolicited Downstream Label Distribution:** In this case, a downstream LSR distributes a label binding in response to an explicit request from an upstream LSR (Figure 5-22(b)). A downstream LSR *B* discovers a "next hop" for a particular FEC, generates a label for this FEC, and communicates the binding to an upstream LSR *A*. LSR *A* inserts the binding into its LIB (label information base) and checks if it need to update the corresponding entry in its LFIB (label forwarding information base). If LSR *B* is the next hop for the FEC, LSR *A* can use that label as an outgoing label, knowing that its meaning is understood.

Each FEC is specified as a set of one or more FEC elements. Each FEC element identifies a set of packets that may be mapped to the corresponding LSP. When an LSP is shared by multiple FEC elements, the shared LSP is terminated at (or before) the node where the FEC elements can no longer share the same path. Following are the currently defined types of FEC elements:

1. *Address Prefix*. This element is an address prefix of any length from 0 to a full address, inclusive.

2. *Host Address*. This element is a full host address.

New element types may be added as needed.

Distribution Control

Independent LSP Control

Each LSR makes independent decision on when to generate labels and communicate them to upstream peers

Communicate label-FEC binding to peers once next-hop has been recognized

LSP is formed as incoming and outgoing labels are spliced together

Characteristics:

Labels can be exchanged with less delay

Does not depend on availability of egress node

Granularity may not be consistent across the nodes at the start

May require separate loop detection/mitigation method

Ordered LSP Control

Label-FEC binding is communicated to peers if:

- LSR is the 'egress' LSR to particular FEC

- label binding has been received from an upstream LSR

LSP formation 'flows' from egress to ingress

Characteristics:

Requires more delay before packets can be forwarded along the LSP

Depends on availability of egress node

Mechanism for consistent granularity and freedom from loops

Used for explicit routing and multicast

Both methods are supported in the standard and can be fully interoperable.

LDP: Label Distribution Protocol

Label Distribution Protocol (LDP) provides LSR discovery mechanisms to enable LSR peers to find each other and establish communication. The LDP protocol is defined in RFC-5036. It defines four types of messages:

- DISCOVERY: used to find neighboring LSRs. Each LSR announces and maintains its presence in a network. LSRs indicate their presence by sending Hello messages periodically. Hello messages are transmitted as UDP packets to the LDP port at the group multicast address for all routers on the subnet.

- SESSION ADJACENCY: used to initialize, keep alive, and shutdown LDP sessions. If two LSRs have discovered each other by means of the LDP Hello messages, they then establish sessions and become *LDP peers*. For this purpose, routers use LDP initialization procedure over TCP transport. After the initialization procedure is completed, the two routers are LDP peers and can exchange Advertisement messages.

- LABEL ADVERTISEMENT: used for label-binding advertisements, request, withdrawal, and release. This is the main purpose of LDP. Advertisement messages are used to maintain label mappings for FECs (Figure 5-21). In general, an LSR requests a label mapping from an LDP peer when it needs one, and advertises a label mapping to an LDP peer when it wants that peer to use the advertised label.
- NOTIFICATION: used to distribute advisory information and to signal error information.

LDP depends on a routing protocol, such as OSPF (Section 8.2.2), to initially establish reachability between the LSRs. The LDP runs over TCP for reliable delivery of messages, except for discovery of LSR peers, which uses UDP and IP multicast. It is designed to be extensible, using messages specified as TLVs (type, value, length) encoded objects.

The IP routing protocol can also be used to define the route for LSP tunnels (hop-by-hop routing). Alternatively, traffic-engineering considerations can determine the explicit route of the LSP (Section 5.4.3). Once a route is determined for an LSP, LDP is used to set up the LSP and assign the labels. Because each LSP is unidirectional, label assignment propagates back from the egress LSR to the originating point (ingress LSR), as illustrated in Figure 5-18.

RSVP-TE

RFC-3209

Explicit Routing

The exchange of PATH and RESV messages between any two LSRs establishes a label association with specific forwarding requirements. The concatenation of these label associations creates the desired edge-to-edge LSP.

5.4.3 Traffic Engineering

Service providers and enterprise operators face the challenge of providing acceptable service levels, or QoS, to their customers and users while simultaneously running an efficient and reliable network. Conventional IP routing aims to find and follow the shortest path between a packet's current location and its destination. This can lead to "hot spots" in the network—routers and links on the intersection of shortest paths to many destinations subject to high traffic load. As the average load on a router rises, packets experience increased loss rates, latency, and jitter. Two solutions exist (and may be deployed in parallel): introducing faster routers and links, or distributing (load balancing) the packet forwarding across alternate (potentially non-shortest-path) routes. The latter solution is called **Traffic Engineering (TE)**.

Constraint-based Routing

One type of constraint would be the ability to find a route (path) that has certain performance characteristics, such as minimum available bandwidth. In this case, the constraint imposed on the routing algorithm is that the computed path must have at least the specified amount of available bandwidth on all links along the path. Different paths (defined by source-destination endpoints) may have different demands for the minimum available bandwidth.

Another type of constraint would be administrative. For example, a network administrator may want to exclude certain traffic from traversing certain links in the network, where such links would be identified by a link attribute. In this case, the constraint imposed on the routing algorithm is that the computed path must not traverse through any of the specified links. On the other hand, the network administrator may want to require certain traffic to traverse only the specified links. Similar to performance constraints, different paths may have different administrative constraints.

Constraint-based routing cannot be supported by conventional IP routing protocols. The key reason is that constraint-based routing requires route (path) calculation at the source router. This requirement is because different sources may have different constraints for a path to the same destination, and the constraints associated with a particular source router are known only to this router, but not to any other router in the network. Unlike this, in conventional IP routing, a route is computed in a distributed manner by every router in the network.

Constrained Shortest Path First (CSPF) is an enhanced version of the shortest-path first (SPF) algorithm used in OSPF (Section 8.2.2). CSPF computes paths taking into account the constraints. When computing paths for LSP tunnels, CSPF considers the physical topology of the network, the attributes of the individual links between LSRs, and the attributes of existing LSPs. CSPF attempts to satisfy the requirements for a new LSP while minimizing congestion by balancing the network load.

5.4.4 Virtual Private Networks

Virtual private networks (VPNs) provide relative or absolute protection for a given traffic flow from other traffic on any particular network segment. VPNs are also used to support tiered services for traffic flows. In general, a VPN provides wide area connectivity to an organization located in multiple sites. MPLS can provide connectivity among VPN sites through LSPs that are dedicated to the given VPN. The LSPs can be used to exchange routing information between the various VPN sites, transparently to other users of the MPLS network. This behavior gives the appearance of a dedicated wide-area network.

Layer-2 VPNs, Layer-3 VPNs

It is possible to build VPNs using a pure IP solution. Although gigabit IP routers are capable of IP forwarding as fast as any MPLS-capable router performs label switching, MPLS VPNs are significantly more efficient than IP VPNs.

5.4.5 MPLS and Quality of Service

Route Protection and Restoration

- End-to-end protection
- Fast node and link reroute

MPLS Protection Types:

1+1: Backup LSP established in advance, resources dedicated, data simultaneously sent on both primary and backup

Switchover performed only by egress LSR

Fastest, but most resource intensive

1:1 : Same as 1+1 with the difference that data is not sent on the backup

Requires failure notification to the ingress LSR to start transmitting on backup

Notification may be send to egress also

Resources in the backup may be used by other traffic

Low priority traffic (e.g., plain IP traffic), shared by other backup paths.

5.5 Summary and Bibliographical Notes

Section 5.1: Scheduling

If a server (router, switch, etc.) is handling multiple flows, there is a danger that aggressive flows will grab too much of its capacity and starve all the other flows. Simple processing of packets in the order of their arrival is not appropriate in such cases, if it is desired to provide equitable access to transmission bandwidth. Scheduling algorithms have been devised to address such issues. The best known is *fair queuing* (FQ) algorithm, originally proposed by [Nagle, 1987], which has many known variations. A simple approach is to form separate waiting lines (queues) for different flows and have the server scan the queues *round robin*, taking the first packet (head-of-the-line) from each queue (unless a queue is empty). In this way, with n hosts competing for a given transmission line, each host gets to send one out of every n packets. Aggressive behavior does not pay off, because sending more packets will not improve this fraction.

A problem with the simple round robin is that it gives more bandwidth to hosts that send large packets at the expense of the hosts that send small packets. *Packet-by-packet FQ* tackles this problem by transmitting packets from different flows so that the packet completion times approximate those of a bit-by-bit fair queuing system. Every time a packet arrives, its completion time under bit-by-bit FQ is computed as its finish number. The next packet to be transmitted is the one with the smallest finish number among all the packets waiting in the queues.

If it is desirable to assign different importance to different flows, e.g., to ensure that voice packets receive priority treatment, then *packet-by-packet weighted fair queuing* (WFQ) is used. WFQ plays a central role in QoS architectures and it is implemented in today's router products [Cisco, 1999; Cisco, 2006]. Organizations that manage their own intranets can employ WFQ-capable routers to provide QoS to their internal flows.

[Keshav, 1997] provides a comprehensive review of scheduling disciplines in data networks.

[Bhatti & Crowcroft, 2000] has a brief review of various packet scheduling algorithms

[Elhanany *et al*., 2001] reviews hardware techniques of packet forwarding through a router or switch

Packet scheduling disciplines are also discussed in [Cisco, 1995]

One of the earliest publications mentioning the leaky bucket algorithm is [Turner, 1986].

Section 5.3: Active Queue Management

Random Early Detection (RED) keeps the overall throughput high while maintaining a small average queue length, and tolerates transient congestion. When the average queue has exceeded a certain threshold, RED routers drop packets at random so that TCP connections back off at different times. This avoids the global synchronization effect of all connections. RED was proposed by Floyd and Jacobson [1993]. Sally Floyd maintains a list of papers on RED here: http://www.icir.org/floyd/red.html. Christiansen, *et al*., [2001] also provides an overview of various versions of RED and additional references. Srikant [2004] presents an in-depth account on RED techniques and their analysis.

Clark and Fang [1998] proposed an extension of RED to provide different levels of drop precedence for two classes of traffic. Their algorithm is called RED with IN/OUT or RIO for short. A device, located on the sourcing traffic side of a network boundary, serves a "policy meter." Packets are classified as being inside (IN) or outside (OUT), depending on whether they conform to the service allocation profile of a given sender/user. RIO uses twin RED algorithms for dropping packets, one for INs and one for OUTs. RIO chooses different parameters of RED algorithms for IN and OUT packets, which may be lined up in the same or different queues. When congestion sets in, RIO is able to preferentially drop OUT packets.

Explicit Congestion Notification (ECN) is described in RFC-3168. As expected, ECN reduces the number of packets dropped by a TCP connection, which, in turn, reduces latency and especially jitter, because packet retransmissions are avoided [RFC-2884]. This outcome is most dramatic when the TCP connection sends occasional isolated segments, which is common for interactive connections (such as remote logins) and transactional protocols (such as HTTP requests, the conversational phase of SMTP, or SQL requests. Such a sender will receive ECN notification, which it ignores because it sends only occasional isolated segments, but it benefits from the fact that its segment was not dropped. The reason for this effect is that the sender can detect a loss of an isolated segment only by an RTO timeout (which is relatively long), because there are no subsequent segments to generate duplicate ACKs. Effects of ECN on bulk transports are less clear because subsequent segments will soon generate duplicate ACKs and recent TCP versions use fast recovery to resend dropped segments in a timely manner (Section 2.2).

Section 5.4: Multiprotocol Label Switching (MPLS)

MPLS provides the ability to forward packets over arbitrary non-shortest paths, and emulate high-speed "tunnels" between IP-only (non-label-switched) domains. It offers a capability not available to conventionally routed solutions: the forwarding packets over arbitrary, non-shortest paths, which is particularly useful for managing network resources, known as "traffic engineering."

Label Distribution Protocol (LDP) is defined in RFC-3036 and is used to provide mechanisms for MPLS routers to process and route labeled traffic across an MPLS network.

Davie and Rekhter [2000] offer a very readable account of MPLS fundamentals, which, although dated, is still relevant to study because it explains well the basic concepts. A relatively recent and comprehensive review of MPLS is available in [De Ghein, 2007].

[Ziegelmann, 2007] Constrained Shortest Path First (CSPF)

Problems

Problem 5.1

Problem 5.2

Eight hosts, labeled *A*, *B*, *C*, *D*, *E*, *F*, *G*, and *H*, share a transmission link the capacity of which is 85. Their respective bandwidth demands are 14, 7, 3, 3, 25, 8, 10, and 18, and their weights are 3, 4, 1, 0.4, 5, 0.6, 2, and 1. Calculate the max-min weighted fair share allocation for these hosts. Show your work neatly, step by step.

Problem 5.3

Problem 5.4

Consider a packet-by-packet FQ scheduler that discerns three different classes of packets (forms three queues). Suppose a 1-Kbyte packet of class 2 arrives upon the following situation. The current round number equals 85000. There is a packet of class 3 currently in service and its finish number is 106496. There are also two packets of class 1 waiting in queue 1 and their finish numbers are $F_{1,1}$ = 98304 and $F_{1,2}$ = 114688.

Determine the finish number of the packet that just arrived. For all the packets under consideration, write down the order of transmissions under packet-by-packet FQ. Show the process.

Problem 5.5

Consider the following scenario for a packet-by-packet FQ scheduler and transmission rate equal 1 bit per unit of time. At time t=0 a packet of $L_{1,1}$=100 bits arrives on flow 1 and a packet of $L_{3,1}$=60 bits arrives on flow 3. The subsequent arrivals are as follows: $L_{1,2}$=120 and $L_{3,2}$=190 at t=100; $L_{2,1}$=50 at t=200; $L_{4,1}$=30 at t=250; $L_{1,3}$=160 and $L_{4,2}$=30 at t=300, $L_{4,3}$=50 at 350, $L_{2,2}$=150 and $L_{3,3}$=100 at t=400; $L_{1,4}$=140 at t=460; $L_{3,4}$=60 and $L_{4,4}$=50 at t=500; $L_{3,5}$=200 at t=560; $L_{2,3}$=120 at t=600; $L_{1,5}$=700 at t=700; $L_{2,4}$=50 at t=800; and $L_{2,5}$=60 at t=850. For every time new packets arrive, write down the sorted finish numbers. What is the actual order of transmissions under packet-by-packet FQ?

Problem 5.6

A transmitter works at a rate of 1 Mbps and distinguishes three types of packets: voice, data, and video. Voice packets are assigned weight 3, data packets 1, and video packets 1.5. Assume that initially arrive a voice packet of 200 bytes a data packet of 50 bytes and a video packet of 1000 bytes. Thereafter, voice packets of 200 bytes arrive every 20 ms and video packets every 40 ms. A data packet of 500 bytes arrives at 20 ms, another one of 1000 bytes at 40 ms and a one of 50 bytes at 70 ms. Write down the sequence in which a packet-by-packet WFQ scheduler would transmit the packets that arrive during the first 100 ms. Show the procedure.

Problem 5.7

Suppose a router has four input flows and one output link with the transmission rate of 1 byte/second. The router receives packets as listed in the table below. Assume the time starts at 0 and the "arrival time" is the time the packet arrives at the router. Write down the order and times at which the packets will be transmitted under:

(a) Packet-by-packet fair queuing (FQ)
(b) Packet-by-packet weighted fair queuing (WFQ), where flows 2 and 4 are entitled to twice the link capacity of flow 3, and flow 1 is entitled to twice the capacity of flow 2

Packet #	Arrival time [sec]	Packet size [bytes]	Flow ID	Departure order/ time under FQ	Departure order/ time under WFQ
1	0	100	1		
2	0	60	3		
3	100	120	1		
4	100	190	3		
5	200	50	2		
6	250	30	4		
7	300	30	4		
8	300	60	1		
9	650	50	3		
10	650	30	4		
11	710	60	1		
12	710	30	4		

Problem 5.8

[*Priority + Fair Queuing*] Consider a scheduler with three queues: one high priority queue and two non-priority queues that should share the resource that remains after the priority queue is served in a fair manner. The priority packets are scheduled to go first (lined up in their order of arrival), regardless of whether there are packets in non-priority queues. The priority packets are scheduled in a *non-preemptive* manner, which means that any packet currently being serviced from a non-priority queue is allowed to finish.

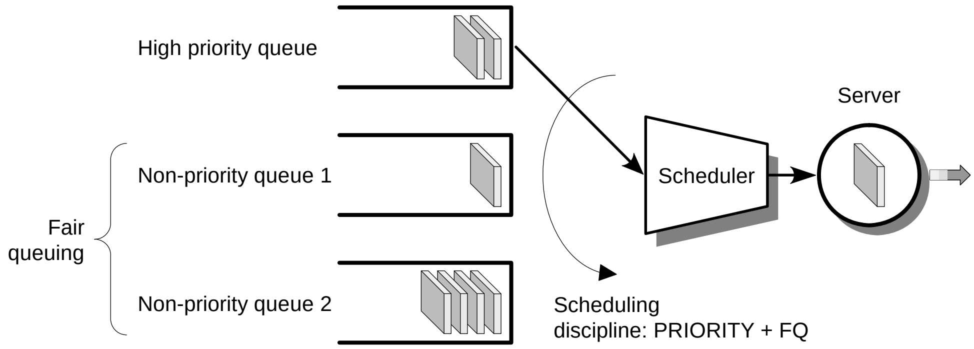

Modify the formula for calculating the packet finish number.

Assume that the first several packets arrive at following times:

> Priority queue: (arrival time $A_{1,1} = 5$, packet length $L_{1,1} = 2$); ($A_{1,2} = 8$, $L_{1,2} = 2$);
>
> First non-priority queue: ($A_{2,1} = 0$, $L_{2,1} = 6$); ($A_{2,2} = 7$, $L_{2,2} = 1$);
>
> Second non-priority queue: ($A_{3,1} = 1$, $L_{3,1} = 2$); ($A_{3,2} = 7$, $L_{3,2} = 1$);

Show the order in which these packets will leave the server and the departure times.

Problem 5.9

Consider the network in Figure 5-15 with the hosts attached as shown in the figure below. (As in Figure 5-15, routers *C*, *E*, and *F* are MPLS-capable.) Assume that the network starts in the initial state, where all IP routing tables are already built, but LFIBs (label forwarding information bases) are empty. Now assume that three hosts start sending data, first host `96.1.1.7` sends a packet to host `17.1.1.35`, then host `17.3.1.24` sends a packet to `10.2.5.35`, and finally `96.1.3.13` sends a packet to `17.3.1.24`. Assume that all LSPs (label switched paths) will be built based on the shortest paths found by the IP routing protocols and that the FECs (forwarding equivalence classes) will be determined only based on the destination IP addresses.

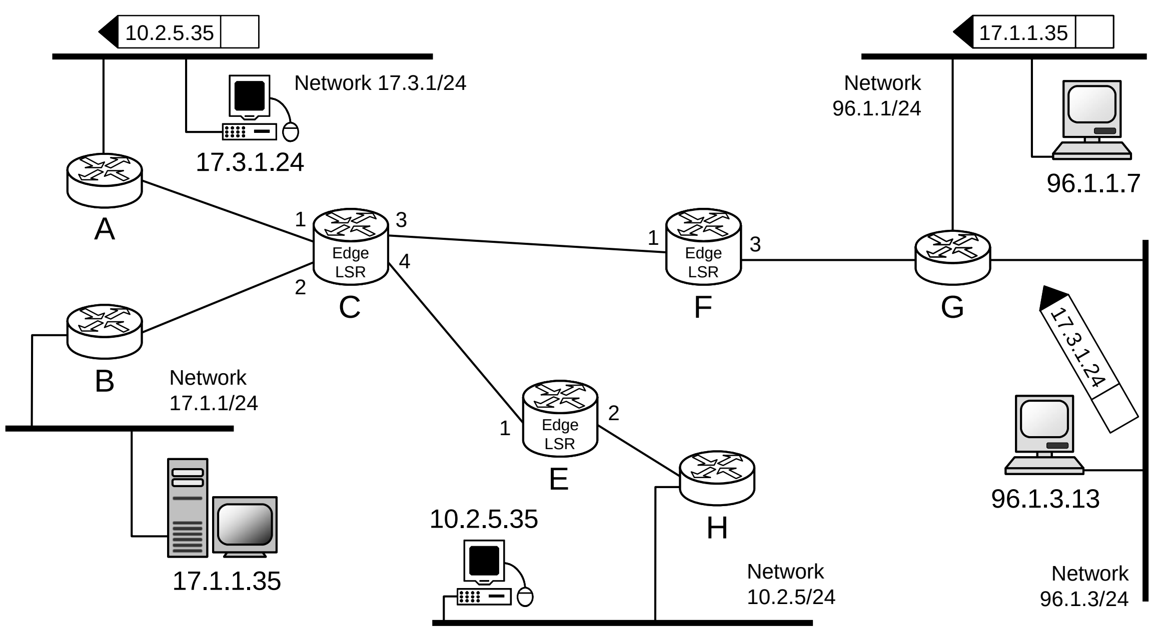

(a) Show step-by-step how LSPs will be built and what will be the entries of the LFIBs for MPLS-capable routers.
(b) For every instance of packet forwarding, indicate whether an LFIB or an ordinary IP forwarding table will be used to forward the packet. In case of LFIB-based forwarding, show the packet's MPLS label value.
(c) What is the minimum number of FECs and what is the minimum number of LSPs that needs to be set up?

Problem 5.10

Chapter 6
Wireless Networks

Contents

This chapter reviews wireless networks. The focus is on the network and link layers, and very little is mentioned about the physical layer of wireless networks. In addition, there is a little mention of infrastructure-based wireless networks and the focus is on infrastructure-less wireless networks.

6.1 Mesh Networks

In a multihop wireless ad hoc network, mobile nodes cooperate to form a network without the help of any infrastructure such as access points or base stations. The mobile nodes, instead, forward packets for each other, allowing nodes beyond direct wireless transmission range of each other to communicate over possibly multihop routes through a number of forwarding peer mobile nodes. The mobility of the nodes and the fundamentally limited capacity of the wireless channel, together with wireless transmission effects such as attenuation, multipath propagation, and interference, combine to create significant challenges for network protocols operating in an ad hoc network.

Figure 6-1

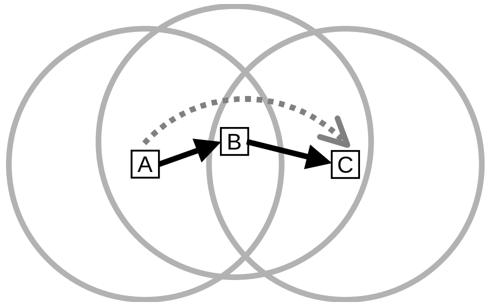

Figure 6-1: Example wireless mesh network: *A* communicates with *C* via *B*.

Figure 6-2: ------------- caption here -----------------.

Figure 6-2

6.2 Routing Protocols for Mesh Networks

In wired networks with fixed infrastructure, a communication endpoint device, known as "host," does not normally participate in routing protocols. This role is reserved for intermediary computing "nodes" that relay packets from a source host to a destination host. On the other hand, in wireless mesh networks it is common that computing nodes assume both "host" and "node"

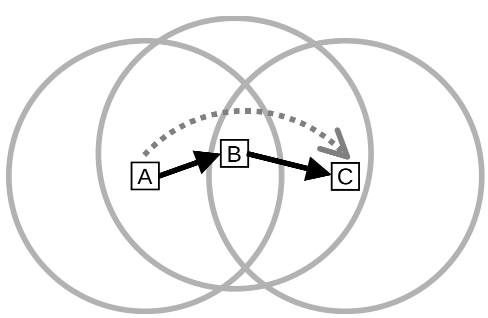

Figure 6-3: Example mobile ad-hoc network: *A* communicates with *C* via *B*.

roles—all nodes may be communication endpoints and all nodes may relay packets for other nodes. Therefore, in this chapter I use the terms "host" and "node" interchangeably.

Although there have been dozens of new routing protocols proposed for MANETs, the majority of these protocols actually rely on fundamental techniques that have been studied rigorously in the wired environment. However, each protocol typically employs a new heuristic to improve or optimize a legacy protocol for the purposes of routing in the mobile wireless environment. In fact, there are a few mechanisms that have received recent interest primarily because of their possible application to MANETs. There are two main classes of routing protocols:

- Proactive
    - Continuously update reachability information in the network
    - When a route is needed, it is immediately available
    - DSDV by Perkins and Bhagwat (SIGCOMM 94)
    - Destination Sequenced Distance vector
- Reactive
    - Routing discovery is initiated only when needed
    - Route maintenance is needed to provide information about invalid routes
    - DSR by Johnson and Maltz
    - AODV by Perkins and Royer
- Hybrid
    - Zone routing protocol (ZRP)

Centralized vs. localized solution:

Nodes in *centralized* solution need to know full network information to make decision; mobility or changes in activity status (power control) cause huge communication overhead to maintain the network information.

Nodes in *localized* algorithm require only local knowledge (direct neighbors, 2-hop neighbors) to make decisions. Majority of published solutions are centralized, compared with other centralized solutions.

Next, a brief survey of various mechanisms is given.

6.2.1 Dynamic Source Routing (DSR) Protocol

Source routing means that the sender must know in advance the complete sequence of hops to be used as the route to the destination. DSR is an *on-demand* (or *reactive*) ad hoc network routing protocol, i.e., it is activated only when the need arises rather than operating continuously in background by sending periodic route updates. DSR divides the routing problem in two parts: *Route Discovery* and *Route Maintenance*, both of which operate entirely on-demand. In Route Discovery, a node actively searches through the network to find a route to an intended destination

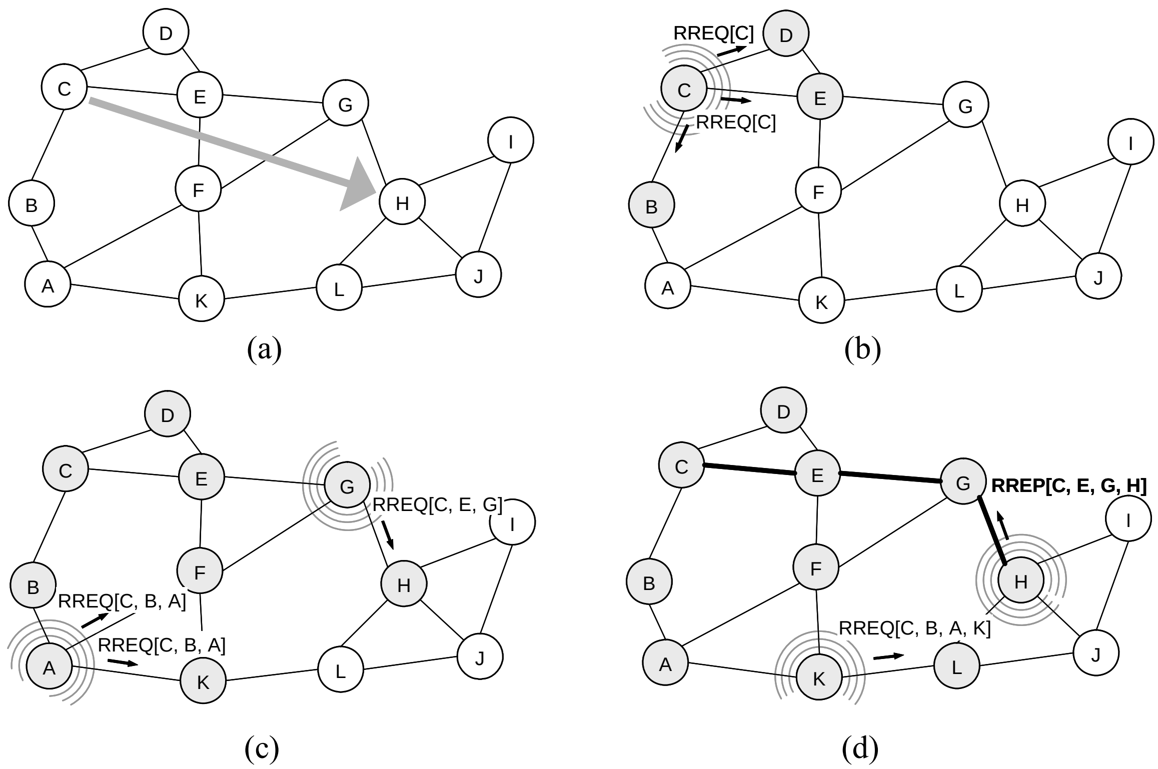

Figure 6-4: Route discovery in DSR: node C seeks to communicate to node H. Gray shaded nodes already received RREQ. The path in bold in (c) indicates the route selected by H for RREP. See text for details. (Note: the step where B and E broadcast RREQ is not shown.)

node. While using a route to send packets to the destination, the sending node runs the Route Maintenance process by which it determines if the route has broken, for example because two nodes along the route have moved out of wireless transmission range of each other.

An example is illustrated in Figure 6-1, where host *C* needs to establish a communication session with host *H*. A node that has a packet to send to a destination (*C* in our example) searches its Route Cache for a route to that destination. If no cached route is found, node *C* initiates Route Discovery by broadcasting a ROUTE REQUEST (RREQ) packet containing the destination node address (known as the *target* of the Route Discovery), a list (initially empty) of nodes traversed by this RREQ, and a *request identifier* from this source node. The request identifier, the address of this source node (known as the *initiator* of the Route Discovery), and the destination address together uniquely identify this Route Discovery attempt.

A node receiving a ROUTE REQUEST checks to see if it has previously forwarded a RREQ from this Discovery by examining the IP Source Address, destination address, and request identifier. For example, in Figure 6-4(b), nodes *B*, *E*, and *D* are the first to receive RREQ and they re-broadcast it to their neighbors. If the recipient of RREQ has recently seen this identifier, or if its own address is already present in the list in RREQ of nodes traversed by this RREQ, the node silently drops the packet. Otherwise, it appends its address to the node list and re-broadcasts the REQUEST. When a RREQ reaches the destination node, *H* in our example, this node returns a ROUTE REPLY (RREP) to the initiator of the ROUTE REQUEST. If an intermediary node receives a RREQ for a destination for which it caches the route in its Route Cache, it can send RREP back to the source without further propagating RREQ. The RREP contains a copy of the node list from the RREQ, and can be delivered to the initiator node by reversing the node list, by using a route back to the initiator from its own Route Cache, or "piggybacking" the RREP on a new ROUTE REQUEST targeting the original initiator. This path is indicated with bold lines in Figure 6-4(d). When the initiator of the request (node *C*) receives the ROUTE REPLY, it adds the newly acquired route to its Route Cache for future use.

In Route Maintenance mode, an intermediary node forwarding a packet for a source attempts to verify that the packet successfully reached the next hop in the route. A node can make this confirmation using a hop-to-hop acknowledgement at the link layer (such as is provided in IEEE 802.11 protocol), a passive acknowledgement (i.e., listen for that node sending packet to its next hop), or by explicitly requesting network- or higher-layer acknowledgement. Transmitting node can also solicit ACK from next-hop node. A packet is possibly retransmitted if it is sent over an unreliable MAC, although it should not be retransmitted if retransmission has already been attempted at the MAC layer. If a packet is not acknowledged, the forwarding node assumes that the next-hop destination is unreachable over this link, and sends a ROUTE ERROR to the source of the packet, indicating the broken link. A node receiving a ROUTE ERROR removes that link from its Route Cache.

In the basic version of DSR, every packet carries the entire route in the header of the packet, but some recent enhancements to DSR use implicit source routing to avoid this overhead. Instead, after the first packet containing a full source route has been sent along the route to the destination, subsequent packets need only contain a flow identifier to represent the route, and nodes along the route maintain flow state to remember the next hop to be used along this route based on the address of the sender and the flow identifier; one flow identifier can designate the default flow for this source and destination, in which case even the flow identifier is not represented in a packet.

A number of optimizations to the basic DSR protocol have been proposed [Perkins 2001, Chapter 5]. One example of such an optimization is *packet salvaging*. When a node forwarding a packet fails to receive acknowledgement from the next-hop destination, as described above, in addition to sending a ROUTE ERROR back to the source of the packet, the node may attempt to use an alternate route to the destination, if it knows of one. Specifically, the node searches its Route Cache for a route to the destination; if it finds one, then it salvages the packet by replacing the existing source route for the packet with the new route from its Route Cache. To prevent the possibility of infinite looping of a packet, each source route includes a *salvage count*, indicating how many times the packet has been salvaged in this way. Packets with salvage count larger than some predetermined value cannot be salvaged again.

In summary, DSR is able to adapt quickly to dynamic network topology but it has large overhead in data packets. The protocol does not assume bidirectional links.

6.2.2 Ad Hoc On-Demand Distance-Vector (AODV) Protocol

DSR includes source routes in packet headers and large headers can degrade performance, particularly when data contents of a packet are small. AODV attempts to improve on DSR by maintaining *routing tables* at the nodes, so that data packets do *not* have to contain routes. AODV retains the desirable feature of DSR that routes are maintained only between nodes which need to communicate.

ROUTE REQUEST packets are forwarded in a manner similar to DSR. When a node re-broadcasts a ROUTE REQUEST, it sets up a reverse path pointing towards the source. AODV assumes symmetric (bidirectional) links. When the intended destination receives a RREQ, it replies by sending a ROUTE REPLY. RREP travels along the reverse path set-up when RREQ is forwarded.

An intermediate node (not the destination) may also send a RREP, provided that it knows a more recent path than the one previously known to sender S. To determine whether the path known to an intermediate node is more recent, destination sequence numbers are used. The likelihood that an intermediate node will send a RREP when using AODV is not as high as in DSR. A new RREQ by node S for a destination is assigned a higher destination sequence number. An intermediate node, which knows a route but with a smaller sequence number, *cannot send* RREP.

A routing table entry maintaining a reverse path is purged after a timeout interval. Timeout should be long enough to allow RREP to come back. A routing table entry maintaining a forward path is purged if not used for an active_route_timeout interval. If no is data being sent using a particular routing table entry, that entry will be deleted from the routing table (even if the route may actually still be valid).

In summary, routes in AODV need not be included in the headers of data packets (unlike DSR, where every data packet carries the source route). Nodes maintain routing tables containing entries only for routes that are in active use. At most one next-hop per destination is maintained at each node, whereas DSR may maintain several routes for a single destination. Lastly, unused routes expire even if topology does not change.

6.3 More Wireless Link-Layer Protocols

Section 1.5.3 described Wi-Fi (IEEE 802.11). This section describes more wireless link-layer protocols and technologies.

6.3.1 IEEE 802.11n (MIMO Wi-Fi)

IEEE 802.11n builds on previous 802.11 standards (Section 1.5.3) by adding mechanisms to improve network throughput. 802.11n operates in the 2.4- and 5-GHz frequency bands. A key improvement is in the radio communication technology, but 802.11n is much more than just a new radio for 802.11. In addition to providing higher bit rates (as was done in 802.11a, b, and g), 802.11n significantly changed the frame format of 802.11. Specifically, 802.11n added *multiple-input multiple-output* (MIMO, pronounced *my-moh*) and 40-MHz channels to the physical layer (PHY), and *frame aggregation* to the MAC layer. It achieves a significant increase in the maximum raw data rate over the two previous standards (802.11a and 802.11g), from 54 Mbps to 600 Mbps, improves reliability, and increases transmission distance. At 300 feet, 802.11g

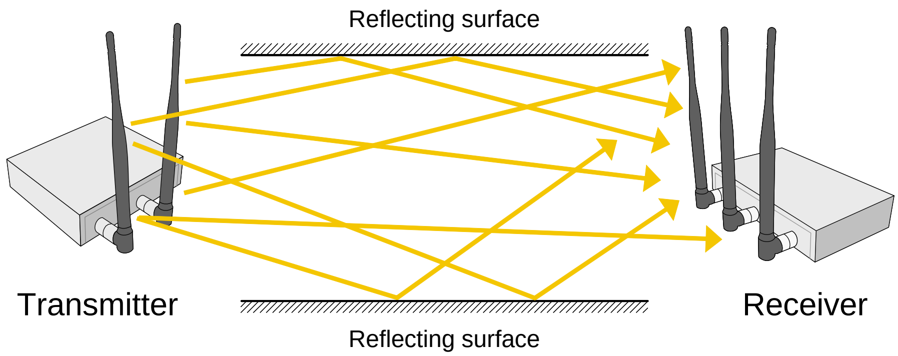

Figure 6-5: MIMO wireless devices with two transmitter and three receiver antennas. Notice how multiple radio beams are reflected from objects in the environment to arrive at the receiver antennas.

performance drops to 1 Mbps; on the other hand, at the same distance 802.11n networks operate at up to 70 Mbps, which is 70 times faster than 802.11g.

IEEE 802.11n-capable devices are also referred as *High Throughput (HT) devices*. An HT device declares that it is an HT device by transmitting the *HT Capabilities element*. The device also uses the HT Capabilities element to advertise which optional capabilities of the 802.11n standard it implements. The HT Capabilities element is carried as part of some control frames that wireless devices exchange during the connection setup or in periodical announcements. It is present in these frames: Beacon, Association Request, Association Response, Reassociation Request, Reassociation Response, Probe Request and Probe Response frames.

IEEE 802.11n standard modifies the frame formats used by 802.11n devices from those of "legacy" 802.11 devices. When 802.11n devices are operating in pure high-throughput mode, this is known as "greenfield mode," because it lacks any constraints imposed by prior technologies. This mode achieves the highest effective throughput offered by the 802.11n standard. To avoid rendering the network useless due to massive interference and collisions, the standard describes some mechanisms for backward compatibility with existing 802.11a/b/g deployments. These mechanisms are reviewed at the end of this section.

Physical (PHY) Layer Enhancements

A key to the 802.11n speed increase is the use of multiple antennas to send and receive more than one communication signals simultaneously, thus multiplying total performance of the Wi-Fi signal. This is similar to having two FM radios tuned to the same channel at the same time—the signal becomes louder and clearer. As for a receiver side analogy, people hear better with both ears than if one is shut. **Multiple-input multiple-output (MIMO)** is a technology that uses multiple antennas to resolve coherently more information than possible using a single antenna. Each 802.11n device has two radios: for transmitter and receiver. Although previous 802.11 technologies commonly use one transmit and two receive antennas, MIMO uses multiple independent transmit and receive antennas. This is reflected in the two, three, or even more antennas found on some 802.11n access points or routers (Figure 6-5). The network client cards on 802.11n mobile devices also have multiple antennas, although these are not that prominently visible. Each antenna can establish a separate (but simultaneous) connection with the corresponding antenna on the other device.

MIMO technology takes advantage of what is normally the enemy of wireless networks: *multipath propagation*. Multipath is the way radio frequency (RF) signals bounce off walls, ceilings, and other surfaces and then arrive with different amounts of delay at the receiver. MIMO is able to process and recombine these scattered and otherwise useless signals using sophisticated signal-processing algorithms.

A MIMO transmitter divides a higher-rate data stream into multiple lower-rate streams. (802.11n MIMO uses up to four streams.) Each of the unique lower-rate streams is then transmitted on the same spectral channel, but through a different transmit antenna via a separate spatial path to a corresponding receiver. The multiple transmitters and antennas use a process called *transmit beamforming* (TxBF) to focus the output power in the direction of the receivers. TxBF steers an outgoing signal stream toward the intended receiver by concentrating the transmitted radio energy in the appropriate direction. This increases signal strength and data rates. On the receiving end, multiple receivers and antennas reverse the process using *receive combining*.

The receiving end is where the most of the computation takes place. Each receiver receives a separate data stream and, using sophisticated signal processing, recombines the data into the original data stream. This technique is called Spatial Division Multiplexing (SDM). MIMO SDM can significantly increase data throughput as the number of resolved spatial data streams is increased. Spatial multiplexing combines multiple beams of data at the receiving end, theoretically multiplying throughput—but also multiplying the chances of interference. This is why the transmitter and the receiver must cooperate to mitigate interference by sending radio energy only in the intended direction. The transmitter needs feedback information from the receiver about the received signal so that the transmitter can tune each signal it sends. This feedback is available only from 802.11n devices, not from 802.11a, b, or g devices. This feedback is not immediate and is only valid for a short time. Any physical movement by the transmitter, receiver, or elements in the environment will quickly invalidate the parameters used for beamforming. The wavelength for a 2.4-GHz radio is only 120mm, and only 55mm for 5-GHz radio. Therefore, a normal walking pace of 1 meter per second will rapidly move the receiver out of the spot where the transmitter's beamforming efforts are most effective. In addition, transmit beamforming is useful only when transmitting to a single receiver. It is not possible to optimize the phase of the transmitted signals when sending broadcast or multicast transmissions.

The advantage of this approach is that it can achieve great throughput on a single, standard, 20-MHz channel while maintaining backward compatibility with legacy 801.11b/g devices. The net impact is that the overall signal strength (that is, *link budget*) is improved by as much as 5 dBi (dB isotropic). Although that may not sound significant, the improved link budget allows signals to travel farther, or, alternatively, maintains a higher data rate at a given distance as compared with a traditional 802.11g single transmitter/receiver product.

Each spatial stream requires a separate antenna at both the transmitter and the receiver. 802.11n defines many "$M \times N$" antenna configurations, ranging from "1×1" to "4×4." This refers to the number of transmit (M) and receive (N) antennas—for example, an access point with two transmit and three receive antennas is a "2×3" MIMO device. In addition, MIMO technology requires a separate radio frequency chain and analog-to-digital converter for each MIMO antenna. This translates to higher implementation costs compared to non-MIMO systems.

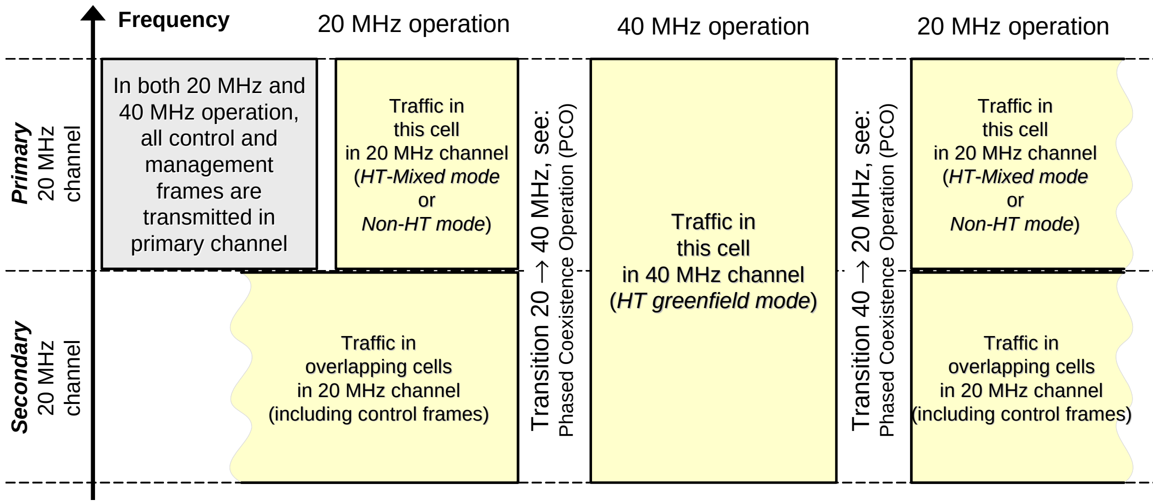

Figure 6-6: 802.11n channel bonding and 20/40 MHz operation. (Phased Coexistence Operation (PCO) is described later, in Figure 6-21.)

Channel bonding. In addition to MIMO, the physical layer of 802.11n can use double-wide channels that occupy 40 MHz of bandwidth. Legacy 802.11a, b, and g devices use 20-MHz-wide channels to transmit data. 802.11n can bond two 20-MHz channels that are adjacent in the frequency domain into one that is 40 MHz wide. That doubling of bandwidth results in a theoretical doubling of information-carrying capacity (data transmission rate). Up to four data-streams can be sent simultaneously using 20MHz or 40MHz channels. A theoretical maximum data rate of 600 Mbps can be achieved using four double-width channels (40 MHz). Although the initial intention of the 40MHz channel was for the 5 GHz band, because of the additional new spectrum, 40MHz channels are permitted in the 2.4 GHz band. Due to the limited spectrum and overlapping channels, 40MHz operation in 2.4 GHz requires special attention.

A 40-MHz channel is created by bonding two contiguous 20-MHz channels: a "primary" or "control" 20-MHz channel and a "secondary" or "extension" 20-MHz channel (Figure 6-6). **Primary channel** is the common channel of operation for all stations (including HT and non-HT) that are members of the BSS (Basic Service Set, defined in Section 1.5.3). To preserve interoperability with legacy clients, 802.11n access point transmits all control and management frames in the primary channel. All 20-MHz clients (whether HT or legacy non-HT) only associate to the primary channel, because the beacon frame is only transmitted on the primary channel. All transmissions to and from clients must be on the primary 20 MHz channel. Hence, all 40-MHz operation in 802.11n is termed "20/40 MHz." **Secondary channel** is a 20-MHz channel associated with a primary channel; it may be located in the frequency spectrum below or above the primary channel. It is used only by HT stations for creating a 40-MHz channel. A station is not required to react to control frames received on its secondary channel, even if it is capable of decoding those frames. The secondary channel of one BSS may be used by an overlapping BSS as its primary channel. If an access point detects an overlapping BSS whose primary channel is the access point's secondary channel, it switches to 20-MHz operation and may subsequently move to a different channel or pair of channels.

Phased Coexistence Operation (PCO) is an option in which an 802.11n access point alternates between using 20-MHz and 40-MHz channels. Before operating in the 40-MHz mode, the access point explicitly reserves both adjacent 20-MHz channels. This mechanism is described later in this section.

Based on the bandwidth used by devices in an 802.11n network, the operational modes can be classified as follows:

- *Legacy (non-HT) mode*. The operation is similar to IEEE 802.11a/b/g. This mode uses the primary 20-MHz channel for transmission.
- *Duplicate legacy mode*. In this mode, the devices use a 40-MHz channel bandwidth, but the same data are transmitted in the primary and secondary halves of the 40-MHz channel. This feature allows the station to send a control frame simultaneously on both 20-MHz channels, which improves efficiency. Examples are given later in this section.
- *High-throughput (HT) mode*. HT mode is available for both 20- and 40MHz channels. In this mode, supporting one and two spatial streams is mandatory. A maximum of four spatial streams is supported.
- *HT duplicate mode*. This mode uses the modulation and coding scheme (MCS) #32 that provides the lowest transmission rate in a 40-MHz channel (6 Mbps), as well as longer transmission range.

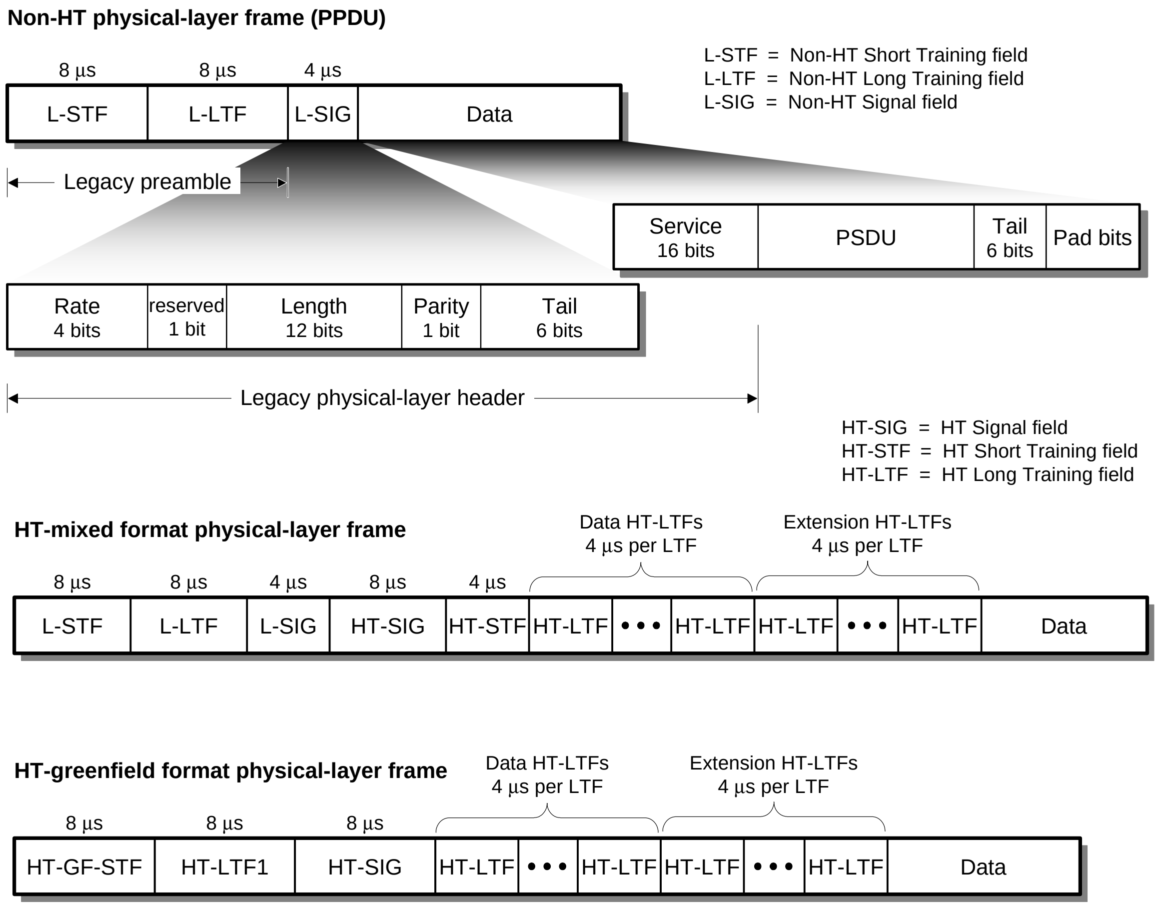

Figure 6-7: 802.11n physical-layer frame formats. Compare to Figure 1-71(b).

Figure 6-7 shows the physical-layer frame formats supported by 802.11n: the legacy format and new formats. Two new formats, called *HT formats*, are defined for the PLCP (PHY Layer Convergence Protocol): *HT-mixed format* and *HT-greenfield format*. There is also an MCS-32 frame format used for the *HT duplicate mode*. In addition to the HT formats, there is a non-HT duplicate format, used in the *duplicate legacy mode*, which duplicates a 20-MHz non-HT (legacy) frame in two 20-MHz halves of a 40-MHz channel.

The legacy **Non-HT frame format** (top row in Figure 6-7) is the 802.11a/g frame format and can be decoded by legacy stations. (Notice that this "legacy" 802.11a/g frame format is different from 802.11b legacy format, shown in Figure 1-71(b). Both are "legacy" in the sense that they predate 802.11n.) The preamble uses short and long training symbols. This allows legacy receivers to detect the transmission, acquire the carrier frequency, and synchronize timing. The physical-layer header contains the legacy Signal field (L-SIG) which indicates the transmission data rate (Rate subfield, in Mbps) and the payload length of the physical-layer frame (Length subfield, in bytes in the range 1–4095), which is a MAC-layer frame.

The **HT-mixed format** (middle row in Figure 6-7) starts with a preamble compatible with the legacy 802.11a/g. The legacy Short Training Field (L-STF), the legacy Long Training Field (L-LTF) and the legacy Signal field (L-SIG) can be decoded by legacy 802.11a/g devices. The rest of the HT-Mixed frame has a new format, and cannot be decoded by legacy 802.11a/g devices.

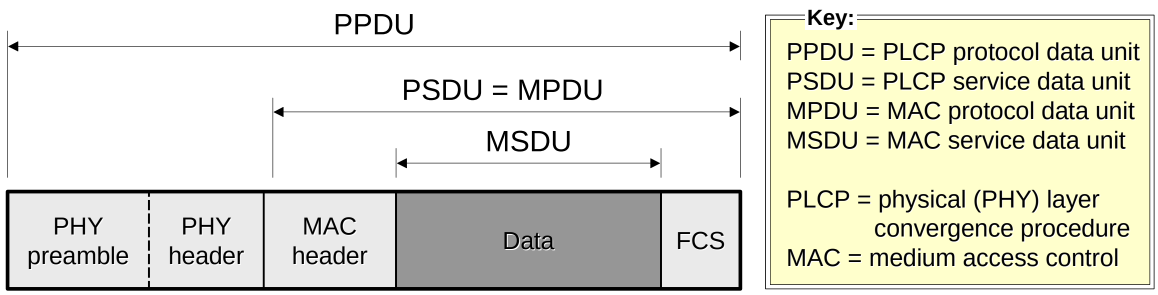

Figure 6-8: Terminology review for the 802.11 frame structure. Compare to Figure 1-71.

The HT preambles are defined in HT-mixed format and in HT-greenfield format to carry the information required to operate in a system with multiple transmit and multiple receive antennas. The HT-SIG field contains the information about the modulation scheme used, channel, bandwidth, length of payload, coding details, number of HT training sequences (HT-LTFs), and tail bits for the encoder. The number of HT-LTFs is decided by the antenna configuration and use of space-time block codes. HT training sequences are used by the receiver for estimating various parameters of the wireless MIMO channel.

The **HT-greenfield format** (bottom row in Figure 6-7) is completely new, without any legacy-compatible part. The preamble transmission time is reduced as compared to the mixed format. Support for the HT Greenfield format is optional and the HT devices can transmit using both 20-MHz and 40-MHz channels.

When an 802.11n access point is configured to operate in Mixed Mode (for example, 802.11b/g/n mode), the access point sends and receives frames based on the type of a client device. By default, the access point always selects the optimum rate for communicating with the client based on wireless channel conditions.

MAC Layer Enhancement: Frame Aggregation

First, the reader may find it useful to review Figure 6-8 for the terminology that will be used in the rest of this section. Figure 6-9 shows the 802.11n MAC frame format. Compared to the legacy 802.11 (Figure 1-71(a)), the change comprises the insertion of the High Throughput (HT) Control field and the change in the length of the frame body. The maximum length of the frame body is 7955 bytes (or, octets) and the overall 802.11n frame length is 8 Kbytes.

Every frame transmitted by an 802.11 device has a significant amount of fixed overhead, including physical layer header, MAC header, interframe spaces, and acknowledgment of transmitted frames (Figure 6-11(a)). (The reader should also check Figure 2-20 and the discussion in Section 2.5) At the highest of data rates, this overhead alone can be longer than the entire data frame. In addition, contention for the channel and collisions also reduce the maximum effective throughput of 802.11. 802.11n addresses these issues by making changes in the MAC layer to improve on the inefficiencies imposed by this fixed overhead and by contention losses.

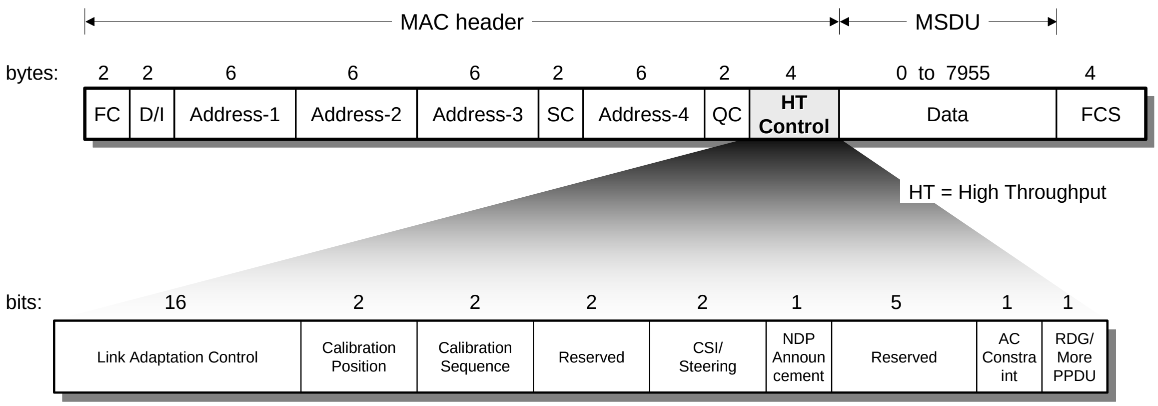

Figure 6-9: 802.11n link-layer frame format. Compare to Figure 1-71(a).

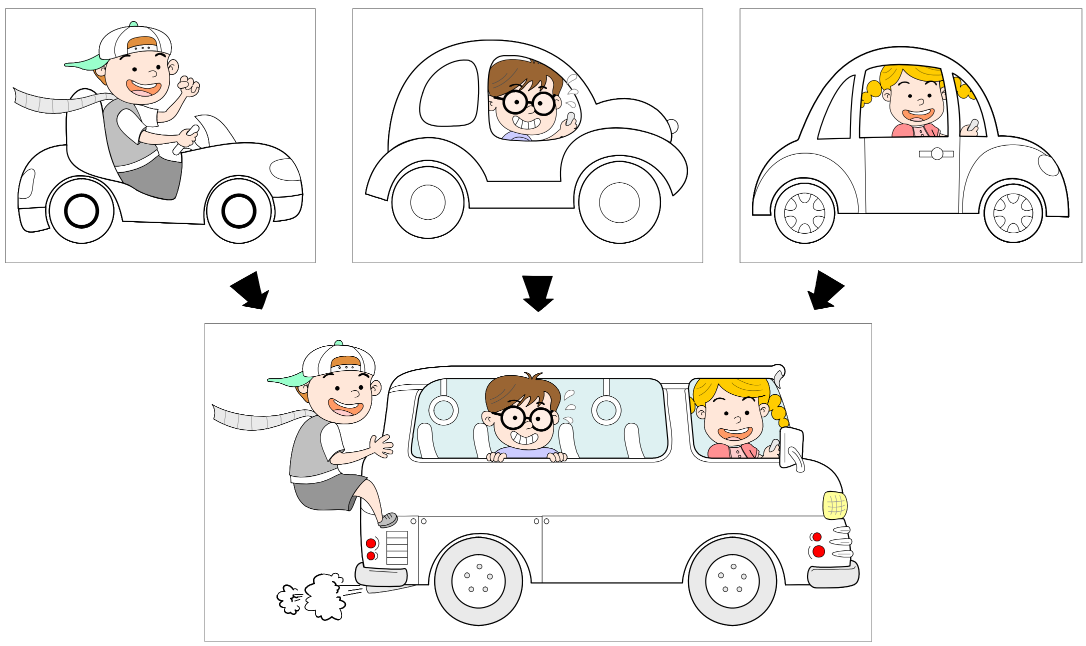

Figure 6-10: Packet aggregation analogy.

To reduce the link-layer overhead, 802.11n employs the mechanism known as *packet aggregation*, which is the process of joining multiple packets together into a single transmission unit, in order to reduce the overhead associated with each transmission. It is equivalent to a group of people riding a bus, rather than each individually riding a personal automobile (Figure 6-10). Generally, packet aggregation is useful in situations where each transmission unit may have significant overhead (preambles, headers, CRC, etc.) or where the expected packet size is small compared to the maximum amount of information that can be transmitted. Because at link layer packets are called frames, the mechanism is correspondingly called "frame aggregation."

Frame aggregation is essentially putting the payloads of two or more frames together into a single transmission. Frame aggregation is a feature of the IEEE 802.11e and 802.11n standards that increases throughput by sending two or more data frames in a single transmission (Figure 6-11(b)). Because control information needs to be specified only once per frame, the ratio of payload data to the total volume of data is higher, allowing higher throughput. In addition, the

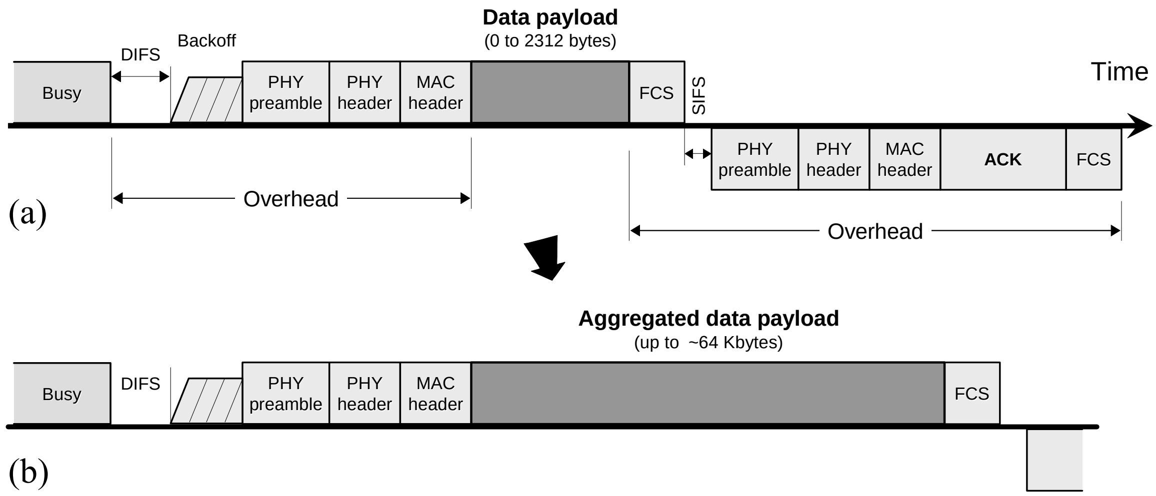

Figure 6-11: (a) Link-layer overhead in legacy IEEE 802.11. (b) 802.11n frame aggregation.

reduced number of frame transmissions significantly reduces the waiting time during the CSMA/CA backoff procedure as well as the number of potential collisions. The maximum frame size is also increased in 802.11n, to accommodate these large, aggregated frames. The maximum frame size is increased from 4 KB to 64 KB. (64 KB frame size is achieved by sending multiple 8 KB frames in a burst, as explained later.)

There are several limitations of frame aggregation. First, all the frames that are aggregated into a transmission must be sent to the same destination; that is, all the frames in the aggregated frame must be addressed to the same mobile client or access point. Second, all the frames to be aggregated have to be ready for transmission at the same time, potentially delaying some frames while waiting for additional frames, in order to attempt to send a larger aggregate frame. Third, the maximum frame size that can be successfully sent is affected by a factor called *channel coherence time*. The time for frame transmission must be shorter than the channel coherence time. Channel coherence time depends on how quickly the transmitter, receiver, and other objects in the environment are moving. When the things are moving faster, the channel data rate is reduced, and therefore the allowed maximum frame size becomes smaller.

Although frame aggregation can increase the throughput at the MAC layer under ideal channel conditions, a larger aggregated frame will cause each station to wait longer before its next chance for channel access. Thus, there is a tradeoff between throughput and delay (or, latency) for frame aggregation at the MAC layer (as throughput increases, latency increase as well). Furthermore, under error-prone channels, corrupting a large aggregated frame may waste a long period of channel time and lead to a lower MAC efficiency.

The ability to send multiple frames without entering the backoff procedure and re-contending for the channel first appeared in the IEEE 802.11e MAC. This mechanism reduces the contention and backoff overhead and thus enhances the efficiency of channel utilization. The notion of **transmit opportunity (TXOP)** is used to specify duration of channel occupation. During TXOP period of time, the station that won channel access can transmit multiple consecutive data frames without re-contending for the channel. If the station determines that it allocated too long TXOP and currently does not have more data to transmit, it may explicitly signal an early completion of its TXOP. This action, known as **truncation of TXOP**, prevents waste by allowing other stations to use the channel. Until the NAV has expired, even if the transmitting station has no data to send and the channel is sensed as idle, other stations do not access the medium for the remaining TXOP. The TXOP holder performs truncation of TXOP by transmitting a **CF-End** (Contention-Free-End) frame, if the remaining TXOP duration is long enough to transmit this frame. CF-End frame indicates that the medium is available. Stations that receive the CF-End frame reset their NAV and can start contending for the medium without further delay.

The frame aggregation can be performed within different sub-layers of the link layer. The 802.11n standard defines two types of frame aggregation: *MAC Service Data Unit (MSDU)* aggregation and *MAC Protocol Data Unit (MPDU)* aggregation. Both aggregation methods group several data frames into one large frame and reduce the overhead to only a single radio preamble for each frame transmission (Figure 6-11(b)). However, there are slight differences in the two aggregation methods that result in differences in the efficiency gained (MSDU aggregation is more efficient). These two methods are described here.

- **MAC Service Data Units (MSDUs) Aggregation**

MSDU aggregation exploits the fact that most mobile access points and most mobile client protocol stacks use Ethernet as their “native” frame format (Figure 1-59). It collects Ethernet frames to be transmitted to a single destination and wraps them in a single 802.11n frame. This is efficient because Ethernet headers are much shorter than 802.11 headers (compare Figure 1-59 and Figure 1-71). For this reason, MSDU aggregation is more efficient than MPDU aggregation.

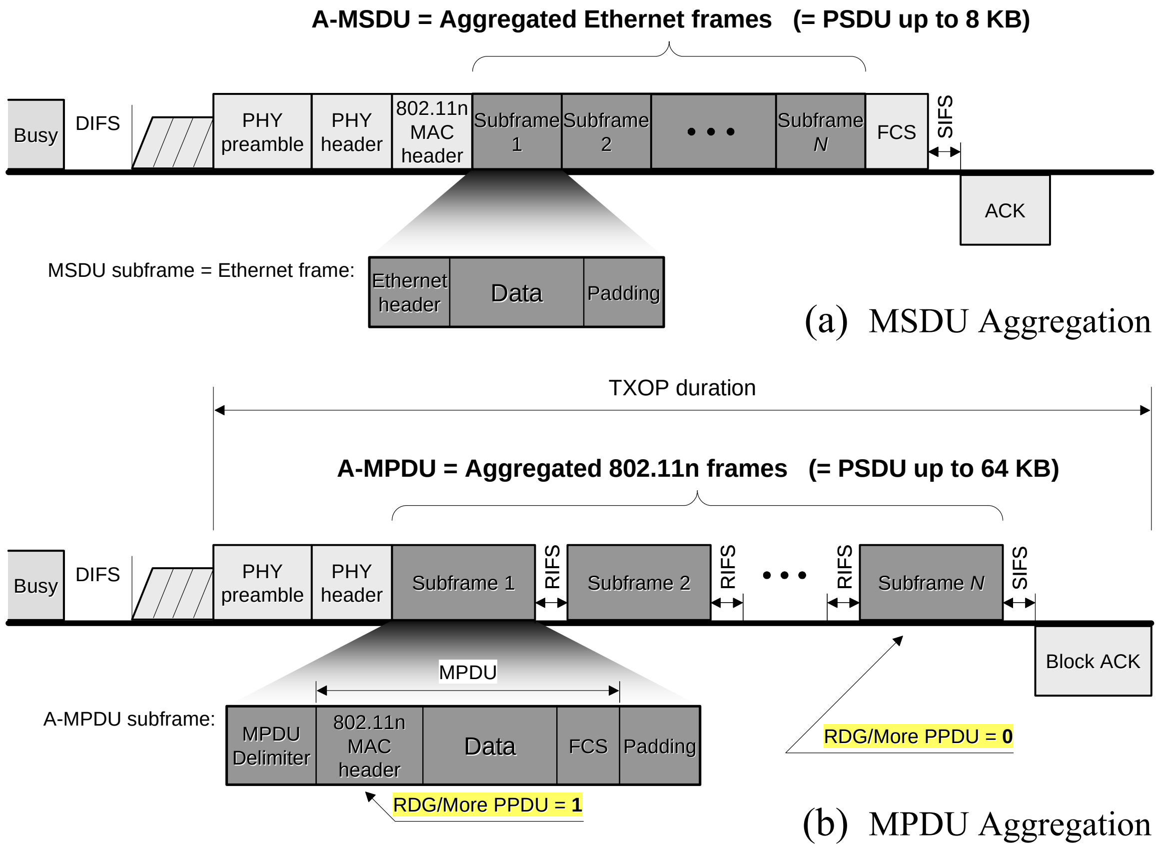

Figure 6-12: 802.11n Frame aggregation methods: (a) MAC Service Data Unit aggregation (A-MSDU); (b) MAC Protocol Data Unit aggregation (A-MPDU).

MSDU aggregation allows several MAC-level service data units (MSDUs) to be concatenated into a single Aggregated MSDU (A-MSDU). Figure 6-12(a) shows the frame format for A-MSDU. In MSDU aggregation, the aggregated payload frames share not just the same physical (PHY) layer header, but also the same 802.11n MAC header. The resulting 802.11n frames can be up to 8 Kbytes in size.

When the source is a mobile device, the aggregated frame is sent to the access point, where the constituent Ethernet frames are forwarded to their ultimate destinations. When the source is an access point, all of the constituent frames in the aggregated frame must be destined to a single mobile client, because there is only a single destination in each mobile client.

With MSDU aggregation, the entire, aggregated frame is encrypted once using the security association of the destination of the outer 802.11n frame wrapper. A restriction of MSDU aggregation is that all of the constituent frames must be of the same quality-of-service (QoS) level. For example, it is not permitted to mix voice frames with best-effort frames.

If no acknowledgement is received, the whole 802.11n frame must be retransmitted. That is, an A-MSDU aggregate fails as a whole even if just one of the enclosed MSDUs contains bit errors.

- **MAC Protocol Data Units (MPDUs) Aggregation**

MPDU aggregation also collects Ethernet frames to be transmitted to a single receiver, but it converts them into 802.11n frames. Normally this is less efficient than MSDU aggregation, but it may be more efficient in environments with high error rates, because of a mechanism called block acknowledgement (described later). This mechanism allows each of the aggregated data frames to be individually acknowledged or retransmitted if affected by an error.

MPDU aggregation scheme enables aggregation of several MAC-level protocol data units (MPDUs) into a single PHY-layer protocol data unit (PPDU). Figure 6-12(b) shows the frame format for an Aggregated MPDU (A-MPDU). A-MPDU consists of a number of MPDU delimiters each followed by an MPDU. Except when it is the last A-MPDU subframe in an A-MPDU, padding bytes are appended to make each A-MPDU subframe a multiple of 4 bytes in length, which can facilitate subframe delineation at the receiver. A-MPDU allows bursting 802.11n frames up to 64 KB.

The purpose of the *MPDU delimiter* (4 bytes long) is to locate the MPDU subframes within the A-MPDU such that the structure of the A-MPDU can usually be recovered when one or more MPDU delimiters are received with errors. Subframes are sent as a burst (not a single unbroken transmission). The subframes are separated on the air from one other by the *Reduced Inter-Frame Space* (RIFS) interval of 2 μs duration (compared to SIFS interval which is 16 μs).[19] Figure 6-12(b) also indicates that the sender uses the "RDG/More PPDU" bit of the HT Control field in the MAC frame (Figure 6-9) to inform the receiver whether there are more subframes in the current burst. If the "RDG/More PPDU" field is set to "1," there will be one or more subframes to follow the current subframe; otherwise, the bit value "0" indicates that this is the last subframe of the burst.

Subframes of an A-MPDUs burst can be acknowledged individually with a single Block-Acknowledgement (described in the next subsection). The MPDU structure can be recovered even if one or more MPDU delimiters are received with errors. Unlike A-MSDU where the whole aggregate needs to be retransmitted, only unacknowledged MPDU subframes need to be retransmitted.

Summary of the characteristics for the two frame aggregation methods:

- MSDU aggregation is more efficient than MPDU aggregation, because the Ethernet header is much shorter than the 802.11 header.

- MPDU structure can be recovered even if one or more MPDU subframes are received with errors; conversely, an MSDU aggregate fails as a whole—even if just one of the enclosed MSDUs contains bit errors the whole A-MSDU must be retransmitted.

- A-MPDU is performed in the software whereas A-MSDU is performed in the hardware.

[19] RIFS is a means of reducing overhead and thereby increasing network efficiency. A transmitter can use RIFS after a transmission when it does not expect to receive immediately any frames, which is the case here. Note that RIFS intervals can only be used within a Greenfield HT network, with HT devices only and no legacy devices.

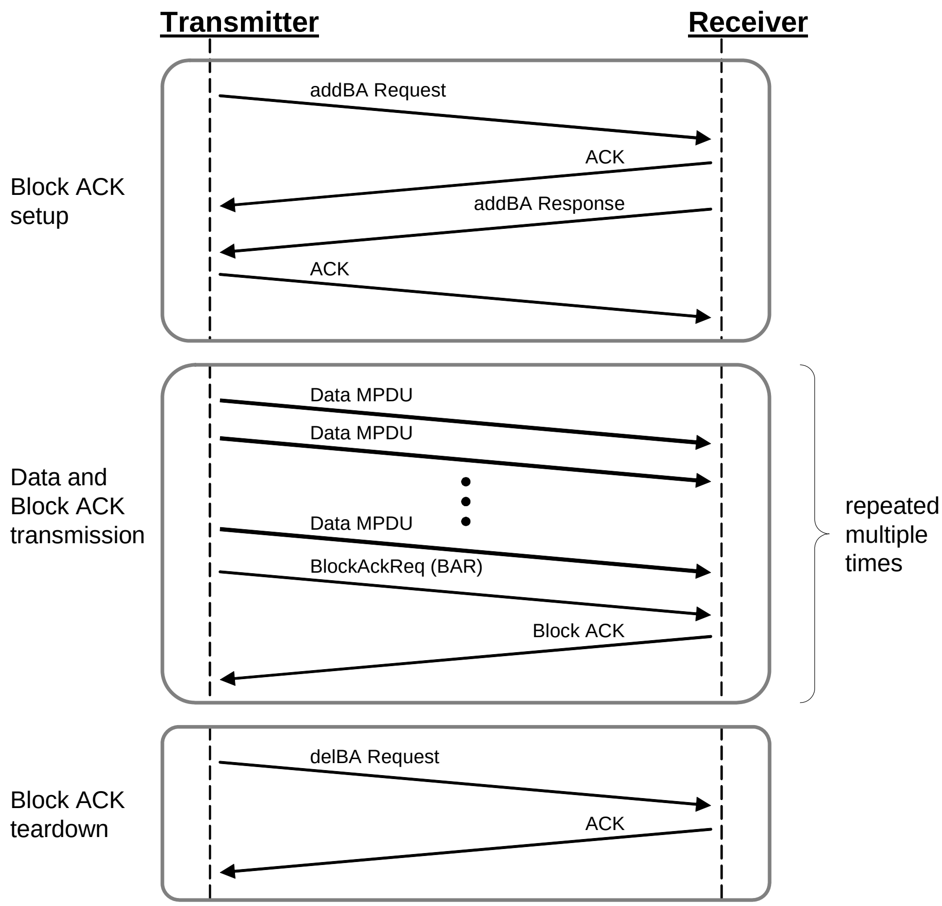

Figure 6-13: Initiation, use, and termination of 802.11n block acknowledgements.

MAC Layer Enhancement: Block Acknowledgement

Rather than sending an individual acknowledgement following each data frame, 802.11n introduces the technique of confirming a burst of up to 64 frames with a single **block acknowledgement** (Block ACK or BACK) frame. The Block ACK mechanism significantly reduces overhead due to bursts of small frames. Block acknowledgment was initially defined in IEEE 802.11e as an optional scheme to improve the MAC efficiency. The 802.11n standard made the Block ACK mechanism mandatory to support by all the HT devices. The Block ACK contains a bitmap to acknowledge selectively individual frames of a burst. This feature is comparable to selective acknowledgements of TCP, known as TCP SACK (Chapter 2).

Figure 6-13 shows how the Block ACK capability is activated, used, and deactivated by sending action frames. *Action frames* are used to request a station to take action on behalf of another. To initiate a Block ACK session, the transmitter sends an Add-Block-Acknowledgment request (addBA, also written as ADDBA). The addBA request indicates a starting frame sequence number and a window size of frame sequence numbers that the receiver should expect as part of the transmission. The receiver can choose to accept or reject the request and informs the transmitter by an addBA response frame. If the receiver rejects the addBA request, the session will continue with the legacy sequential transmit/acknowledgment exchanges. If the receiver

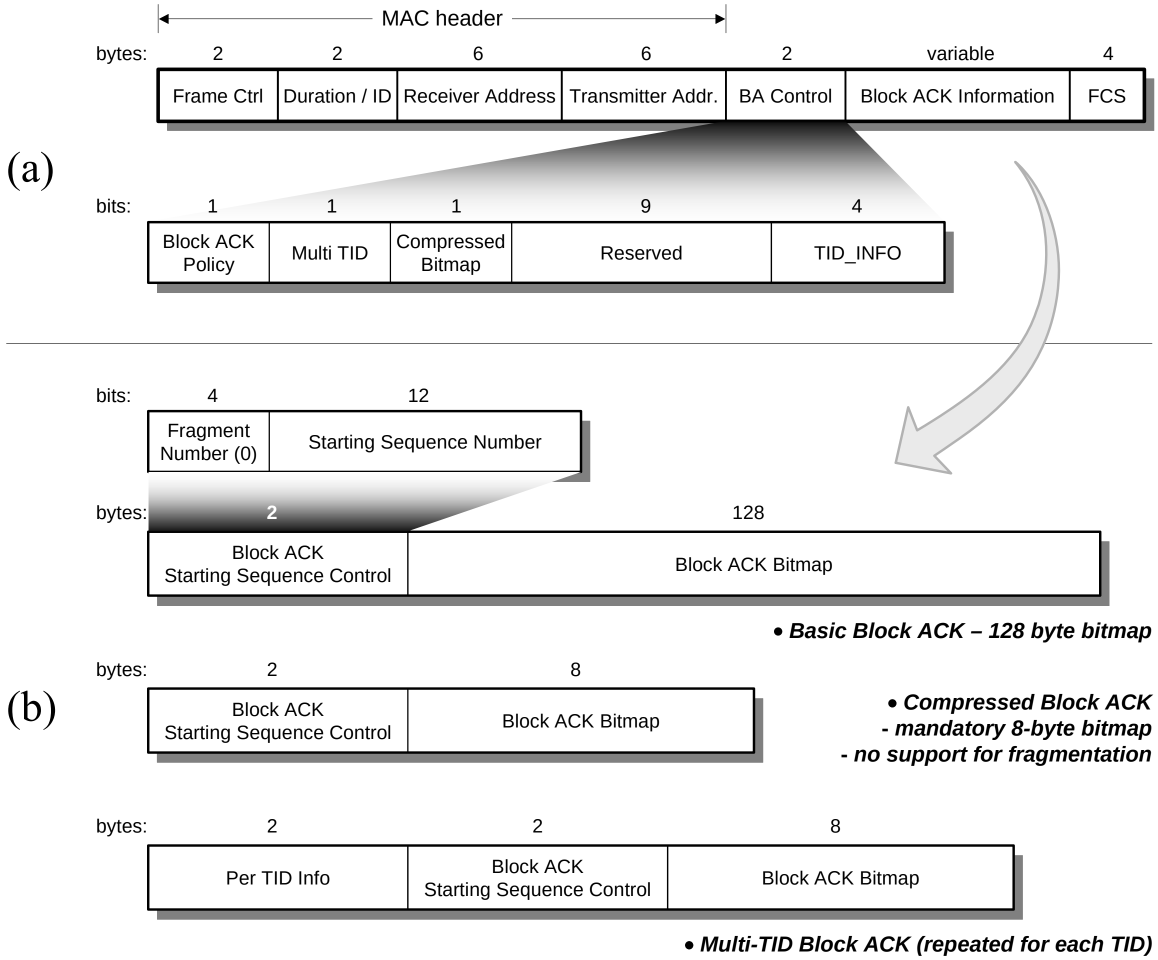

Figure 6-14: (a) 802.11n block acknowledgement frame format. (b) Format of the Block ACK Information field for the three variants of the Block ACK frame.

accepts the addBA request, the transmitter can send multiple frames without waiting for ACK frames. The receiver silently accepts frames that have sequence numbers within the current window. Only after the transmitter solicits a Block ACK by sending a *Block ACK Request* (BlockAckReq or BAR), the receiver responds by a Block ACK response frame indicating the sequence numbers successfully received. Frames that are received outside of the current window are dropped. This cycle may be repeated many times. Finally, when the transmitter does not need Block ACKs any longer, it terminates the BA session by sending a Delete-Block-Acknowledgment request (delBA or DELBA).

The Block ACK carries ACKs for individual frames as bitmaps. The exact format depends on the encoding. Figure 6-14(a) shows the format of Block ACK frames. The subfields of the Block ACK Control field are as follows:

- The Block ACK Policy bit specifies whether the sender requires acknowledgement immediately following BlockAckReq (bit value “0”), or acknowledgement can be delayed (bit value “1”).

- The values of the Multi-TID and Compressed Bitmap fields determine which of three possible Block ACK frame variants is represented (Figure 6-15(a)). The Block ACK frame variants are shown in Figure 6-14(b).

Multi-TID	Compressed Bitmap	Block ACK frame variant
0	0	Basic Block ACK
0	1	Compressed Block ACK
1	0	reserved
1	1	Multi-TID Block ACK

(a)

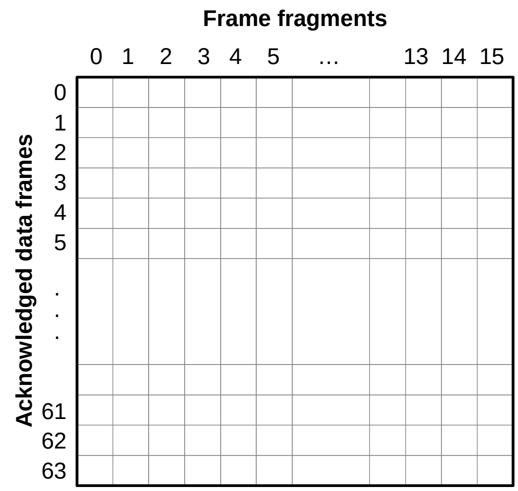

(b)

Figure 6-15: (a) 802.11n Block ACK frame variant encoding. (b) Block ACK Bitmap subfield (128 bytes long = 64×16 bits) of a Basic Block ACK frame variant. Each bit represents the received status (success/failure) of a frame fragment.

- The meaning of the TID_INFO subfield of the BA Control field depends on the Block ACK frame variant type. For the first two variants (Basic Block ACK and Compressed Block ACK), the TID_INFO subfield of the BA Control field contains the TID for which a Block ACK frame is requested. The *traffic identifier (TID)* is assigned by upper-layer protocols to inform the MAC protocol about the type of data that it is asked to transmit. This is important for MAC protocols that support quality of service (QoS), such as 802.11e and 802.11n. Therefore, the first two BA variants are capable of acknowledging only traffic of a single identifier type. A Block ACK frame could be extended to include the bitmaps for multiple TIDs. This extended Block ACK frame variant is called Multi-TID Block ACK (MTBA). More details are provided later.

Figure 6-14(b) shows the structure of the three variants of the Block ACK frame:

- **Basic Block ACK variant.** The Basic Block ACK variant is inherited from the IEEE 802.11e standard. The BA Information field within the Basic Block ACK frame contains the Block ACK Starting Sequence Control subfield and the Block ACK Bitmap, as shown in the top row of Figure 6-14(b). The Starting Sequence Number subfield (12-bit unsigned integer) of the Block ACK Starting Sequence Control field contains the sequence number of the first data frame (MPDU) that this Block ACK is acknowledging. This is the same number as in the previously received BlockAckReq frame to which this Block ACK is responding. When the transmitter receives a Block ACK, based on this number it knows to which BlockAckReq it corresponds. The Fragment Number subfield is always set to 0.

Before describing the BA Bitmap structure, it is necessary to mention the fragmentation mechanism in 802.11. The process of partitioning a MAC-level frame (MPDU) prior to transmission into smaller MAC-level frames is called **fragmentation**. Fragmentation creates smaller frames to increase reliability, by increasing the probability of successful transmission in cases where channel characteristics limit reception reliability for longer frames. The reader may remember IP packet fragmentation (Section 1.4.1), which is done for different reasons, and where the fragments are reassembled only at the final destination. Conversely, defragmentation in

802.11 is accomplished at each immediate receiver. In the 802.11e and 802.11n standards, each MAC frame can be partitioned into up to 16 fragments.

The 128-byte long Block ACK Bitmap subfield represents the received status of up to 64 frames. In other words, the bitmap size is 64×16 bits (Figure 6-15(b)). That is, because each MAC-level frame can be partitioned into up to 16 fragments, 16 bits (2 bytes) are allocated to acknowledge each frame. Each bit of this bitmap represents the received status (success/failure) of a frame fragment. Two bytes are equally allocated even if the frame is not actually fragmented or is partitioned into less than 16 fragments. Suppose a frame has 11 fragments; then 11 bits are used, and remaining 5 bits are not used. Even so, this frame will consume 16 bits in the bitmap. If the frame is not fragmented, only one bit is used. Obviously, in cases with no fragmentation it is not efficient to acknowledge each frame using 2 bytes when all is needed is one bit. The overhead problem occurs also when the number of frames acknowledged by a Block ACK is small, because the bitmap size is fixed to 128 bytes. Thus, using two bytes per acknowledged frame in the bitmap results in an excessive overhead for Block ACK frames.

To overcome the potential overhead problem, 802.11n defines a modified Block ACK frame, called Compressed Block ACK.

- **Compressed Block ACK variant.** This Block ACK frame variant uses a reduced bitmap of 8 bytes, as shown in the middle row of Figure 6-14(b). Fragmentation is not allowed when the compressed Block ACK is used. Accordingly, a compressed Block ACK can acknowledge up to 64 non-fragmented frames. The bitmap size is reduced from 1024 (64×16) bits to 64 (64×1) bits.

The BA Information field within the Compressed Block ACK frame comprises the Block ACK Starting Sequence Control field and the Block ACK bitmap. The Starting Sequence Number subfield of the Block ACK Starting Sequence Control field is the sequence number of the first MSDU or A-MSDU for which this Block ACK is sent. The Fragment Number subfield of the Block ACK Starting Sequence Control field is set to 0.

The 8-byte Block ACK Bitmap within the Compressed Block ACK frame indicates the received status of up to 64 MSDUs and A-MSDUs. Each bit that is set to 1 acknowledges the successful reception of a single MSDU or A-MSDU in the order of sequence number, with the first bit of the bitmap corresponding to the MSDU or A-MSDU with the sequence number that matches the Starting Sequence Number field value.

Figure 6-16 shows an example using Compressed Block ACK frames. Here we assume that the transmitter sends aggregate A-MPDUs with 32 subframes. Bitmap bit position n is set to 1 to acknowledge the receipt of a frame with the sequence number equal to (Starting Sequence Control + n). Bitmap bit position n is set to 0 if a frame with the sequence number (Starting Sequence Control + n) has not been received. For unused fragment numbers of an aggregate frame, the corresponding bits in the bitmap are set to 0. For example, the Block ACK bitmap of the first Block ACK in Figure 6-16 contains [`7F FF FF FF 00 00 00 00`]. The first byte corresponds to the first 8 frames, but read right to left (that is why `7F` instead of `F7`). This means that, relative to the Starting Sequence Number 146, the first four frames and sixth to eight frames are successfully received. The fifth frame is lost (sequence number 150). The second byte corresponds to the second 8 frames, also read right to left, and so on. The last 32 bits are all zero because the A-MPDU contained 32 subframes. In the second transmission, the transmitter resends frame #150 and additional 32 frames (starting with the sequence number 179 up to #211).

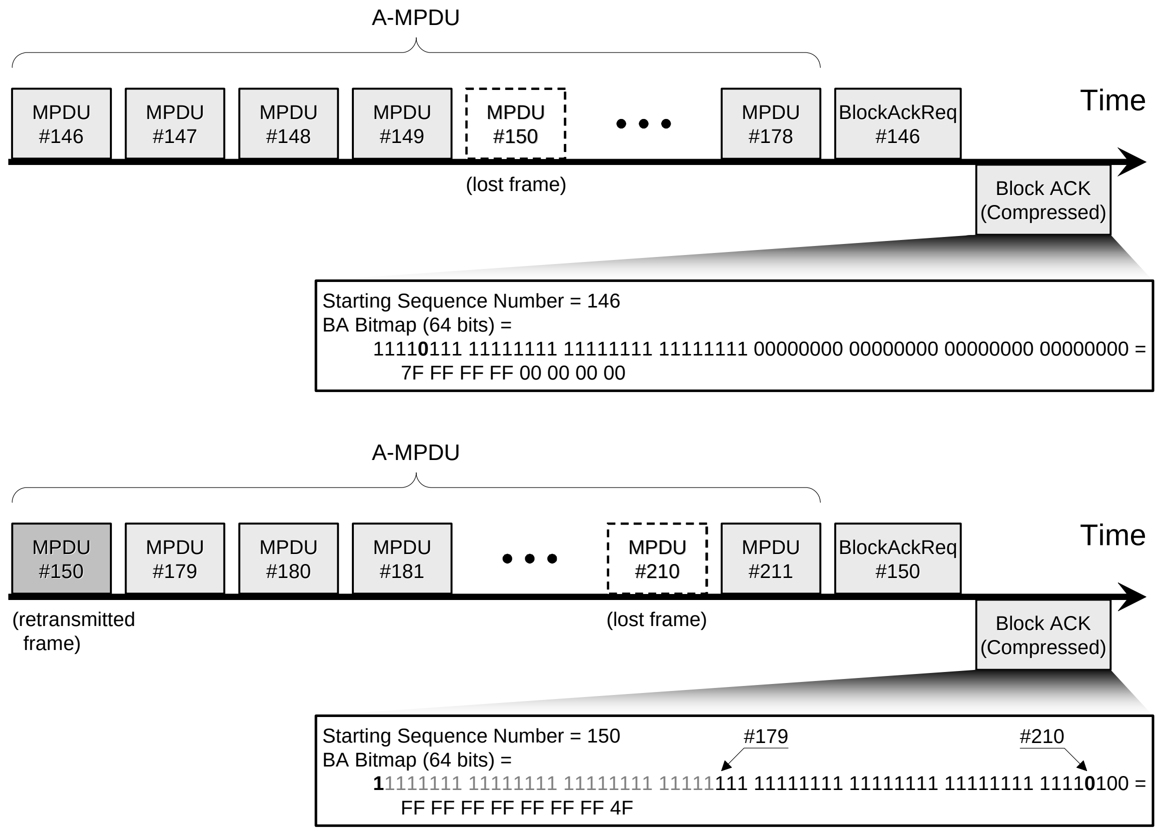

Figure 6-16: 802.11n Block ACK example using the Compressed Block ACK frame variant.

As seen, if a frame is not acknowledged, the sequence numbers can keep moving forward while the sending station keeps retrying that frame. However, when the span between the sequence number of the next frame to be sent and the retry frame becomes 64, the sending unit has to decide what to do. It can stop aggregating while it keeps retrying the old frame, or it can simply drop that frame.

- **Multi-TID Block ACK variant.** The TID_INFO subfield of the BA Control field of the Multi-TID Block ACK frame contains the number of traffic identifiers (TIDs), less one, for which information is reported in the BA Information field. For example, a value of 2 in the TID_INFO field means that information for 3 TIDs is present.

The BA Information field within the Multi-TID Block ACK frame contains one or more instances of the Per TID Info, Block ACK Starting Sequence Control field and the Block ACK Bitmap, as shown in the bottom row of Figure 6-14(b).

The Starting Sequence Number subfield of the Block ACK Starting Sequence Control field is the sequence number of the first MSDU or A-MSDU for which this Block ACK is sent. The first instance of the Per TID Info, Block ACK Starting Sequence Control and Block ACK Bitmap fields that is transmitted corresponds to the lowest TID value, with subsequent instances following ordered by increasing values of the Per TID Info field. The 8-byte Block ACK bitmap within the Multi-TID Block ACK frame functions the same way as for the Compressed Block ACK frame variant.

MAC Layer Enhancement: Reverse Direction (RD) Protocol

The 802.11n also specifies a *bidirectional data transfer* method, known as Reverse Direction (RD) protocol. Conventional transmit opportunity (TXOP) operation described above and already present in IEEE 802.11e allows efficient *unidirectional* transfer of data: the station holding the TXOP can transmit multiple consecutive data frames without reentering backoff procedure. The 802.11n RD protocol provides more efficient *bidirectional* transfer of data between two 802.11 devices during a TXOP by eliminating the need for either device to contend for the channel. This is achieved by piggybacking of data from the receiver on acknowledgements (ACK frame).

Reverse direction mechanism is useful in network services with bidirectional traffic, such as VoIP and online gaming. It allows the transmission of feedback information from the receiver and may enhance the performance of TCP, which requires bidirectional transmission (TCP data segments in one direction and TCP ACK segments in the other). (See Section 2.5 for more discussion.) The conventional TXOP operation only helps the forward direction transmission but not the reverse direction transmission. For application with bidirectional traffic, their performance degrades due to contention for the TXOP transmit opportunities. Reverse direction mechanism allows the holder of TXOP to allocate the unused TXOP time to its receivers to enhance the channel utilization and performance of reverse direction traffic flows.

Before the RD protocol, each unidirectional data transfer required the initiating station to contend for the channel. If RTS/CTS is used, the legacy transmission sequence of RTS (Request To Send) - CTS (Clear To Send) - DATA (Data frame) - ACK (Acknowledgement) allows the sender to transmit only a single data frame in forward direction (Figure 6-17(a)). In the bidirectional data transfer method (i.e., with the RD protocol), once the transmitting station has obtained a TXOP, it may essentially grant permission to the other station to send information back during its TXOP.

Reverse data transmission requires that two roles be defined: RD initiator and RD responder. *RD initiator* is the station that holds the TXOP and has the right to send Reverse Direction Grant (RDG) to the *RD responder*. The RD initiator sends its permission to the RD responder using a Reverse Direction Grant (RDG) in the "RDG/More PPDU" bit of the HT Control field in the MAC frame (Figure 6-9). The RD initiator grants permission to the RD responder by setting this bit to "1." When the RD responder receives the data frame with "RDG/More PPDU" bit set to "1," it decides whether it will send to the RD initiator more frames immediately following the one just received. It first sends an acknowledgement fro the received frame in which the "RDG/More PPDU" bit is set to "1" if one or more data frames will follow the acknowledgement, or with the bit set to "0" otherwise. For the bidirectional data transfer, the reverse DATAr frame can contain a Block ACK to acknowledge the previous DATAf frame. The transmission sequence will then become RTS-CTS-DATAf-DATAr-ACK (Figure 6-17(b)).

If the "RDG/More PPDU" bit in the acknowledgement frame is set to "1," the RD initiator will wait for the transmission from the RD responder, which will start with SIFS or *Reduced Inter-Frame Space* (RIFS) interframe time after the RDG acknowledgement is sent. A transmitter can use RIFS after a transmission when it does not expect to receive immediately any frames, which is the case here. If there is still data to be sent from the RD responder, it can set "RDG/More PPDU" bit in the data frame header to "1" to notify the initiator. The RD initiator still has the right to accept or reject the request. To allocate the extended TXOP needed for additional reverse

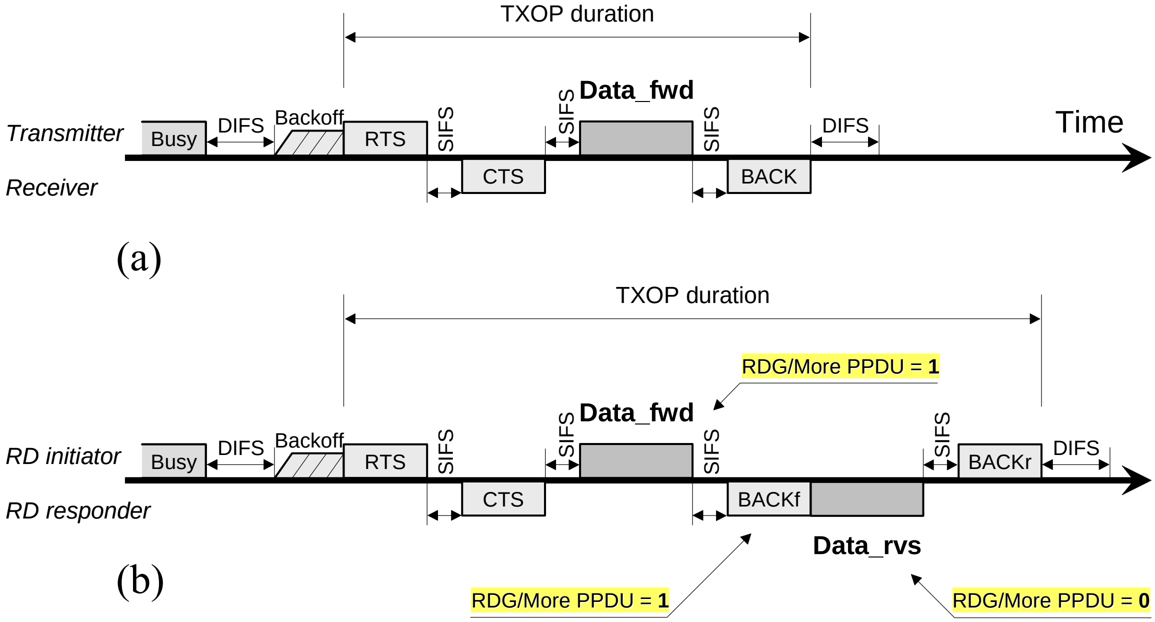

Figure 6-17: 802.11n Reverse Direction (RD) protocol. (a) Unidirectional RTS/CTS access scheme. (b) Bidirectional RTS/CTS access scheme. RD initiator invites RD responder to send reverse traffic by setting the RPG/MorePPDU flag to "1." RD responder sends zero or more frames and sets RPG/MorePPDU to "0" in the last frame of the "RD response burst."

frames, the initiator will set to "1" the "RDG/More PPDU" bit in the acknowledgement frame or the next data frame. To reject the new RDG request, the initiator sets "RDG/More PPDU" to "0."

Backward Compatibility

802.11n devices transmit a signal that cannot be decoded by devices built to an earlier standard. To avoid rendering the network useless due to massive interference and collisions, 802.11n devices must provide backwards compatibility. Compatibility with legacy 802.11 devices has been a critical issue continuously faced throughout the evolution of 802.11. For example, 802.11g provides a mechanism for operation with 802.11b devices. Similarly, 802.11n has a number of mechanisms to provide backward compatibility with 802.11 a, b, and g devices. These mechanisms ensure that the legacy stations are aware of 802.11n transmissions in the same area and do not interfere with them. The cost for achieving this protection is the throughput degradation for 802.11n.

Because the 802.11 MAC layer operation is based on carrier sense multiple access (CSMA/CA) protocol, it is essential for the station that won access to the channel to inform other stations how long it will transmit, to avoid being interrupted. The mechanism of announcing the duration of the transmission is called **protection mechanism**, and different options have emerged in the evolution of 802.11 wireless LANs. Before transmitting at a high data rate, the station must attempt to update the *network allocation vector* (NAV) of non-HT stations that do not support the high data rate, so that they can defer their transmission. (See Section 1.5.3 for the description of NAV.) The duration information has to be transmitted at physical data-rates that are decodable by the legacy stations (the pure 802.11n transmission is not).

Three different *operating modes* are defined for 802.11n devices (actually, four, but one is a kind of sub-mode and omitted here for simplicity). The legacy **Non-HT operating mode** sends data frames in the old 802.11a/g format (shown in the top row of Figure 6-7) so that legacy stations can understand them. However, only 802.11a and g stations understand Non-HT mode format—802.11b stations predate 802.11a/g and do not understand it. Non-HT mode is used by 802.11n devices only to communicate with legacy 802.11 devices, rather than with other 8021.11n devices. It cannot be used with 40-MHz channels (Figure 6-6). At the transmitter, only one transmitting antenna is used in Non-HT mode. Receive diversity is exploited in this mode. An 802.11n device using Non-HT delivers no better performance than 802.11a/g. This mode gives essentially no performance advantage over legacy networks, but offers full compatibility.

The legacy operating mode is a Non-HT (High Throughput) mode, whereas the Mixed and Greenfield modes are HT modes. In **Mixed operating mode**, frames are transmitted with a preamble compatible with the legacy 802.11a/g (middle row in Figure 6-7). The legacy Short Training Field (L-STF), the legacy Long Training Field (L-LTF) and the legacy Signal field (L-SIG) can be decoded by legacy 802.11a/g devices. The rest of the HT-Mixed frame has a new format, and cannot be decoded by legacy 802.11a/g devices.

In **Greenfield operating mode**, high throughput frames are transmitted without any legacy-compatible part (bottom row in Figure 6-7). In this mode, there is no provision to allow a legacy device to understand the frame transmission. Receivers enabled in this mode should be able to decode frames from the legacy mode, mixed mode, and the Greenfield mode transmitters. The preamble is not compatible with legacy 802.11a/g devices and only 802.11n devices can communicate when using the Greenfield format. Support for the Greenfield format is optional and the HT devices can transmit using both 20-MHz and 40-MHz channels.

When a Greenfield device is transmitting, the legacy systems may detect the transmission, and therefore avoid collision, by sensing the presence of a radio signal, using the carrier-sensing mechanism in the physical layer. However, legacy devices cannot decode any part of an HT Greenfield frame. Therefore, they cannot set their NAV and defer the transmission properly. They must rely on continuous physical-layer carrier sensing to detect the busy/idle states of the medium. In the worst case, HT Greenfield transmissions will appear as noise bursts to the legacy devices (and vice versa).

The HT Mixed mode is mandatory to support and transmissions can occur in both 20-MHz and 40-MHz channels. Support for the HT Greenfield mode is optional; again, transmissions can occur in both 20-MHz and 40-MHz channels (Figure 6-6). Support for Non-HT Legacy mode is mandatory for 802.11n devices, and transmissions can occur only in 20-MHz channels.

An 802.11n access point (AP) starts in the Greenfield mode, assuming that all stations in the BSS (Basic Service Set) will be 802.11n capable. If the access point detects a legacy (non-HT) 802.11a/b/g device (at the time when it associates to the access point or from transmissions in an overlapping network), the access point switches to the mixed mode. 802.11n stations are communicating mutually using the mixed mode, and with legacy stations using the non-HT mode. When non-HT stations leave the BSS, the access point, after a preset time, will switch back from the Mixed mode to the Greenfield mode. The same is true of when the access point ceases to hear nearby non-HT stations; it will switch back to the Greenfield mode.

The following protection mechanisms (described later) are defined for 802.11n to work with legacy stations:

- Transmit control frames such as RTS/CTS or CTS-to-self using a legacy data rate, before the HT transmissions. For control frame transmissions, use 20-MHz non-HT frames or 40-MHz non-HT duplicate frames (Figure 6-6).
- L-SIG TXOP protection
- Transmit the first frame of a transmit opportunity (TXOP) period using the non-HT frame that requires a response frame (acknowledgement), which is also sent as a non-HT frame or non-HT duplicate frame. After this initial exchange, the remaining TXOP frames can be transmitted using HT-Greenfield format and can be separated by RIFS (Reduced Inter Frame Spacing).
- Using the HT-Mixed frame format, transmit a frame that requires a response. The remaining TXOP may contain HT-Greenfield frames and/or RIFS sequences.

The first two protection schemes are extension of the protection mechanisms that have been introduced in the migration from 802.11b to 802.11g. Use of control frames such as RTS/CTS or CTS-to-self is a *legacy compatibility mode*. L-SIG TXOP protection is a *mixed compatibility mode* (uses HT-mixed frame format) and is optional to implement. The last two schemes are applicable only in the presence of TXOP, which is a feature that might be enabled only for certain services, such as voice and video transmission.

In an 802.11n HT coverage cell that operates in 20/40-MHz channels, there may be legacy 802.11 devices (operating in the primary 20-MHz channel) along with the 40-MHz HT devices. Furthermore, there may be an overlapping cell with legacy 802.11 devices operating in the secondary channel of this cell. A protection mechanism must take into account both cases and provide protection for the 40-MHz HT devices against interference from either source (i.e., legacy devices inside or outside this cell). Next, we review those protection mechanisms.

- **Control Frames Protection: RTS/CTS Exchange, CTS-to-self, and Dual-CTS**

We have already seen in Section 1.5.3 how RTS/CTS (Request-to-Send/Clear-to-Send) exchange is used for protection of the current transmission. The RTS/CTS frames let nearby 802.11 devices—including those in different but physically overlapping networks—set their network allocation vector (NAV) and defer their transmission. This mechanism is called “virtual carrier sensing” because it operates at the MAC layer, unlike physical-layer carrier sensing. Transmission of RTS/CTS frames helps avoid hidden station problem irrespective of transmission rate and, hence, reduces the collision probability.

802.11g introduced another NAV-setting protection mechanism (also adopted in 802.11n), called **CTS-to-self mechanism**. CTS-to-self allows a device to transmit a short CTS frame, addressed to itself, that includes the NAV duration information for the neighboring legacy devices, which will protect the high-rate transmission that will follow. The advantage of the CTS-to-self NAV distribution mechanism is lower network overhead cost than with the RTS/CTS NAV distribution mechanism—instead of transmitting two frames separated by a SIFS interval, only one frame is transmitted. However, CTS-to-self is less robust against hidden nodes and collisions than RTS/CTS. Stations employing NAV distribution should choose a mechanism that is appropriate

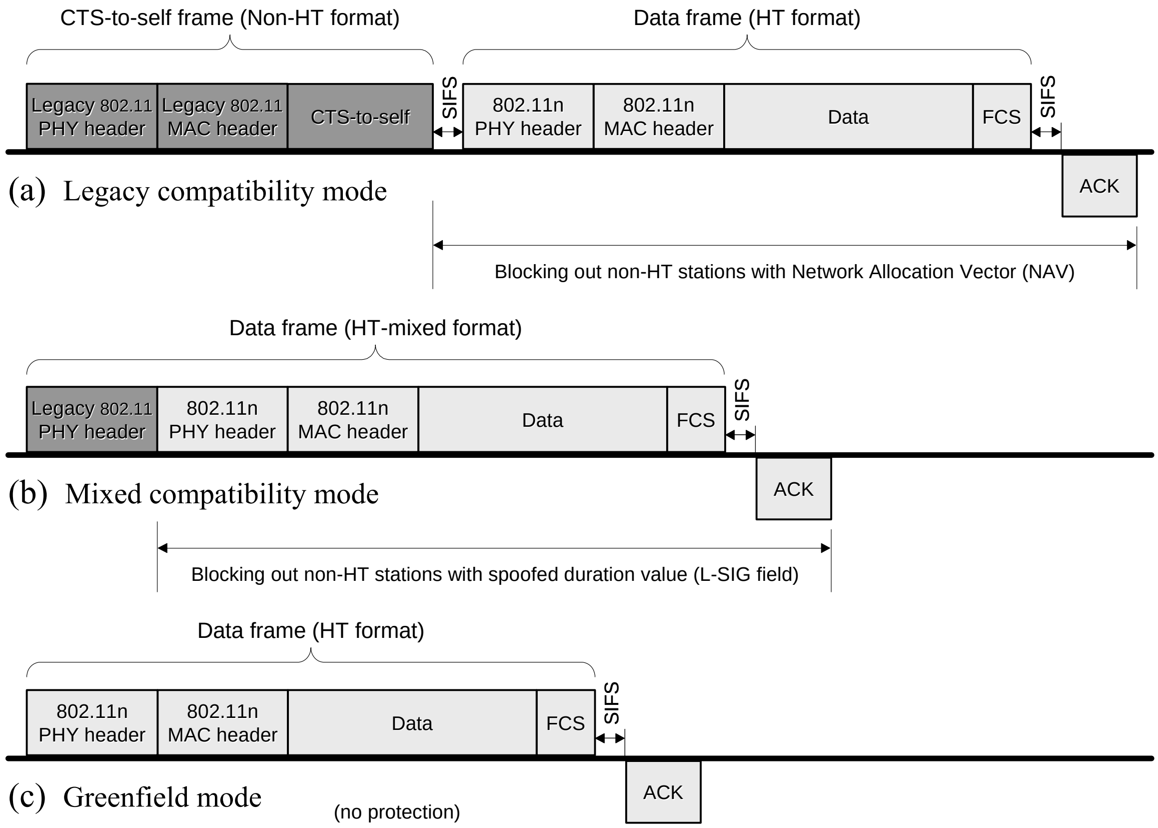

Figure 6-18: 802.11n backwards compatibility modes. (a) Using control frames for NAV distribution. (b) L-SIG TXOP protection. (c) Greenfield mode offers no protection.

for the given network conditions. If errors occur when employing the CTS-to-self mechanism, stations should switch to the more robust RTS/CTS mechanism.

HT protection requires 802.11n devices to announce their intent to transmit by sending legacy-format control frames prior to HT data transmission (Figure 6-18(a)). The CTS-to-self frame must be transmitted using one of the legacy data rates that a legacy device will be able to receive and decode. Transmission rate of the control frames depends on the type of legacy device that is associated in the BSS. If both 802.11b and 802.11g devices are associated, then 802.11b rates (known as Clause 15 rates) are used to transmit protection frames because 802.11g stations can decode such frames.

In **Dual-CTS protection mode**, the RTS receiver transmits two CTS frames, one in Non-HT mode and another in HT mode, so the subsequent data frame is protected by a legacy CTS and an HT CTS. The dual-CTS feature can be enabled or disabled by setting the Dual CTS Protection subfield in beacon frames. Dual-CTS protection has two benefits. First, using the legacy RTS/CTS or legacy CTS-to-self frames to reset NAV timers prevents interference with any nearby 802.11a/b/g cells. Second, it resolves the hidden node problem within the 802.11n cell.

Figure 6-19 shows an example network with an 802.11n access point (AP) and two mobile stations, one 802.11n (station *A*) and the other legacy 802.11g (station *B*). When traffic is generated by station *A*, it first sends an RTS to the AP. The AP responds with two CTS frames, one in HT and the other in legacy format. Station *A* is then free to transmit the data frame, while

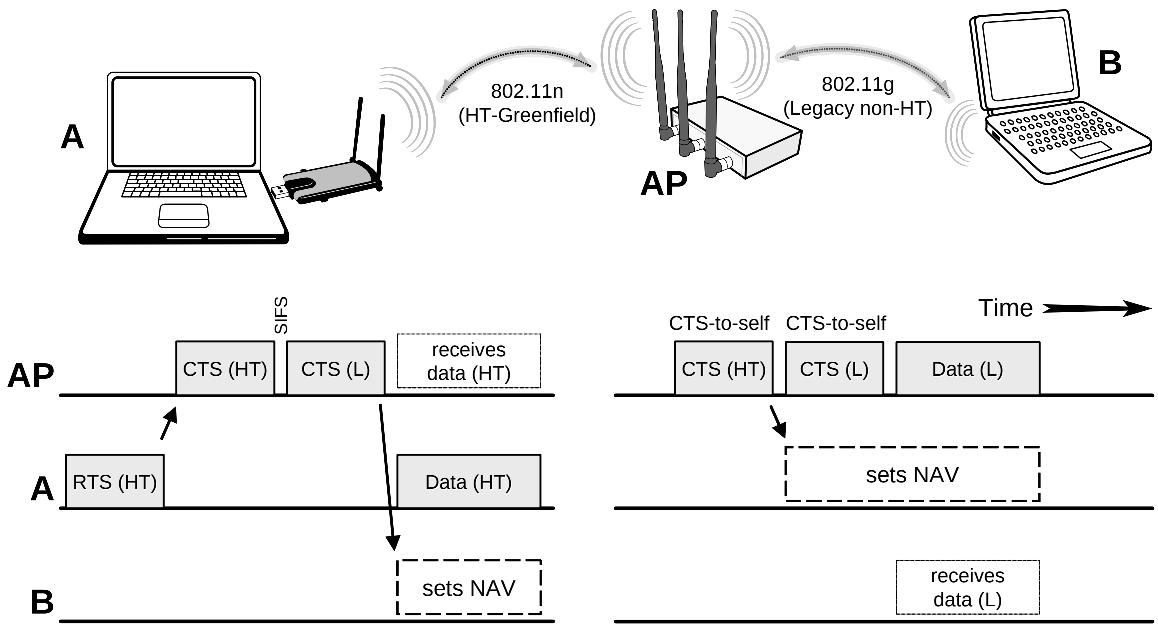

Figure 6-19: Example of 802.11n Dual-CTS protection (CTS-to-self).

other stations in the same and neighboring networks (e.g., station *B*) set their NAV correctly so they do not transmit over the authorized frame, interfering with it. Later, when the AP has traffic to send to station *B*, it uses dual CTS-to-self frames to perform the same function (Figure 6-19).

Dual-CTS makes the 802.11n network a good neighbor to overlapping or adjacent legacy 802.11 networks. It also solves the hidden-station problem where different clients in a cell may not be able to hear each other's transmissions, although, by definition they all can hear the AP and its CTS frames. However, the use of control frames further reduces the data throughput of the network. Although RTS/CTS frames are short (20 and 14 bytes, respectively), it takes more time to transmit them at the legacy rate of 6 Mbps than it takes to transmit 500 bytes of data at 600 Mbps. Therefore, HT protection significantly reduces an 802.11n W-LAN's overall throughput.

- **L-SIG TXOP Protection**

In *Legacy Signal field Transmit Opportunity* (L-SIG TXOP) protection mechanism, protection is achieved by transmitting the frame-duration information in a legacy-formatted physical header, and then transmitting the data at an 802.11n high rate (Figure 6-18(b)). Each frame is sent in an HT-mixed frame format. A legacy device that receives and successfully decodes an HT-mixed frame defers its own transmission based on the duration information present in the legacy Signal (L-SIG) field (see Figure 6-7). Such legacy clients remain silent for the duration of the forthcoming transmission. Following the legacy physical header, the 802.11n device sends the remaining part of the frame using 802.11n HT rates and its multiple spatial streams. L-SIG TXOP protection is also known as *PHY layer spoofing*.

The Rate and Length subfields of the L-SIG field (Figure 6-7) determine the duration of how long non-HT stations should defer their transmission:

$$\text{L-SIG Duration} = (\text{legacy Length} / \text{legacy Rate})$$

This value should be equal to the duration of the remaining HT part of the HT-mixed format frame. The Rate parameter should be set to the value 6 Mbps. Non-HT stations are not able to

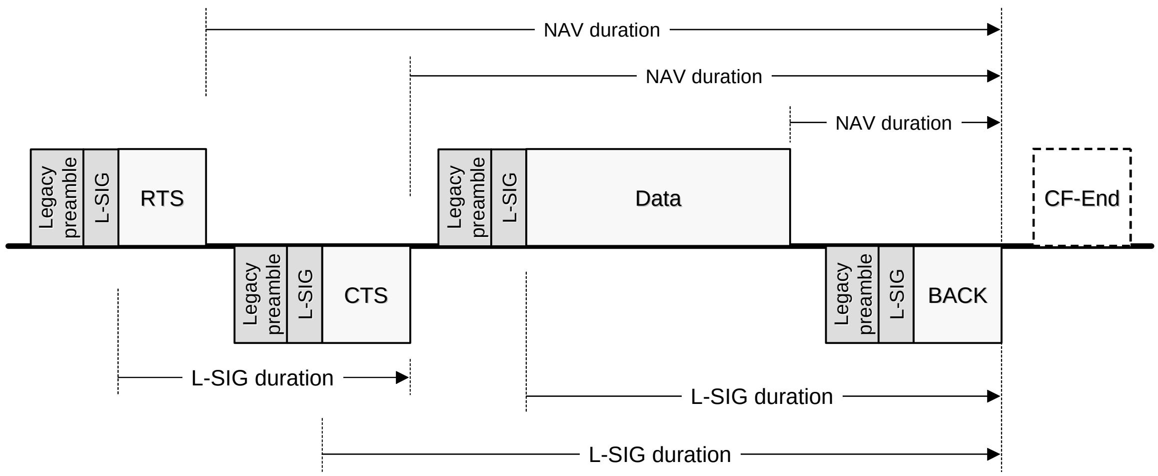

Figure 6-20: 802.11n L-SIG TXOP protection: Example of L-SIG duration setting.

receive any frame that starts throughout the L-SIG duration. Therefore, no frame may be transmitted to a non-HT station during an L-SIG protected TXOP.

Figure 6-20 illustrates an example of how L-SIG Durations are set when using L-SIG TXOP Protection. In this example, an L-SIG TXOP protected sequence starts with an RTS/CTS initial handshake, which provides additional protection from hidden stations. Any initial frame exchange may be used that is valid for the start of a TXOP. The term "L-SIG TXOP protected sequence" includes these initial frames and any subsequent frames transmitted within the protected L-SIG duration.

The TXOP holder should transmit a CF-End frame starting a SIFS period after the L-SIG TXOP protected period (Figure 6-20). Because legacy stations are unable to distinguish a Mixed-mode acknowledgement frame from other Mixed-mode frames, they may mistakenly infer that ACK frame is lost. As a result, they would wait unnecessarily until the EIFS time elapses (see Figure 1-75(b)), which leads to potential unfairness or a "capture effect." CF-End enables other stations to avoid such undesirable effects. Note that this is not an instance of TXOP truncation (described earlier), because here the CF-End frame is not transmitted to reset the NAV.

All HT-mixed mode frames contain the L-SIG field, so is not necessary to send special control frames to announce the medium reservation duration explicitly. An 802.11n station must indicate whether it supports L-SIG TXOP Protection in its L-SIG TXOP Protection Support capability field in Association-Request and Probe-Response frames. The mixed mode can be used in a 40-MHz channel, but to make it compatible with legacy clients, all broadcast and non-aggregated control frames are sent on a legacy 20-MHz channel as defined in 802.11a/b/g, to be interoperable with those clients (Figure 6-6). And, of course, all transmissions to and from legacy clients must be within a legacy 20-MHz channel. L-SIG TXOP protection mechanism is not applicable when 802.11b stations are present, because the Signal field (L-SIG) is encoded in 802.11g frame format that 802.11b devices do not understand. The double physical-layer header (legacy plus 802.11n headers) adds overhead, reducing the throughput. However, it makes possible for 802.11n stations to take advantage of HT features for the remaining part of the frame transmission.

- **Phased Coexistence Operation (PCO)**

Another mechanism for coexistence between 802.11n HT cells and nearby legacy 802.11a/b/g cells is known as *Phased Coexistence Operation (PCO)*. This is an optional mode of operation that divides time into slices and alternates between 20-MHz and 40-MHz transmissions. The HT access point designates time slices for 20-MHz transmissions in both primary and secondary 20-MHz channels, and designates time slices for 40-MHz transmissions. This operation is depicted in Figure 6-6 and now we describe the mechanism for transitioning between the phases. The algorithm for deciding when to switch the phase is beyond the scope of the 802.11n standard.

The phased coexistence operation (PCO) of 802.11n is illustrated in Figure 6-21, where an 802.11n coverage cell (BSS-1) is overlapping a legacy 802.11 cell (BSS-2). Stations *A* and *B* are associated with BSS-1 and station *C* is associated with BSS-2, but it can hear stations in BSS-1 and interfere with their transmissions. Only station *A* is capable of transmitting and receiving frames in the 40-MHz channel. As explained earlier (Figure 6-6), a 40-MHz channel is formed by bonding two contiguous 20-MHz channels, one designated as *primary channel* and the other as *secondary channel*. In this example, BSS-2 happens to operate in what BSS-1 considers its own secondary channel, i.e., the secondary channel of BSS-1 is the primary channel for BSS-2. In 20-MHz phase, all stations contend for medium access in their respective primary channels. When the 802.11n access point wishes to use a 40-MHz channel, it needs to reserve explicitly both adjacent 20-MHz channels. The access point is coordinating the phased operation of the associated stations with 20-MHz and 40-MHz bandwidth usage.

The bottom part of Figure 6-21 shows the MAC-protocol timing diagram for reservation and usage of the 40-MHz channel. Transitions back and forth between 20-MHz and 40-MHz channels start with the Beacon frame or Set-PCO-Phase frame. The 802.11n access point (AP) accomplishes the reservation by setting the NAV timers of all stations with appropriate control frames transmitted on the respective channels. The access point uses CTS-to-self frames to set the NAV timers. As Figure 6-21 depicts, the AP transmits both CTS-to-self frames simultaneously using the *duplicate legacy mode* (described earlier in this section). Although control frames are transmitted only on the primary channel, the secondary channel of BSS-1 is the primary channel of BSS-2, so station *C* will receive the second CTS-to-self and react to it. This feature improves efficiency (but notice that it could not be exploited in Figure 6-19). When the NAV timer is set, the station is blocked from transmission until the NAV timer expires. However, as seen in Figure 6-21 station *A* will also set its own NAV, which means that station *A* too will be blocked. This is why the AP will transmit a CF-End frame in the HT-Greenfield format on the 40-MHz channel, so that only station *A* will decode it and start contending for access to the 40-MHz channel. Recall that CF-End truncates TXOP and clears the NAV timers of the clients that receive this frame.

To end the 40-MHz phase, the HT access point first sends a Set-PCO-Phase frame so station *A* knows the 40-MHz phase is over. Next, to release the 40-MHz channel, the AP uses two CF-End frames sent simultaneously on both 20-MHz channels using the duplicate legacy mode. This will truncate the remaining TXOP for the legacy clients (stations *B* and *C*). Thereafter, all stations may again contend for medium access on their respective 20-MHz primary channels.

Reservation of the 40-MHz channel may not happen smoothly, because traffic in BSS-2 may continue for a long time after the access point transmitted the Beacon frame or Set-PCO-Phase frame (see the initial part of the timing diagram in Figure 6-21). If the secondary channel continues to be busy after the phase transition started, the stations in BSS-1 are not allowed to transmit on the primary 20-MHz channel because their NAV timers are set. If waiting for

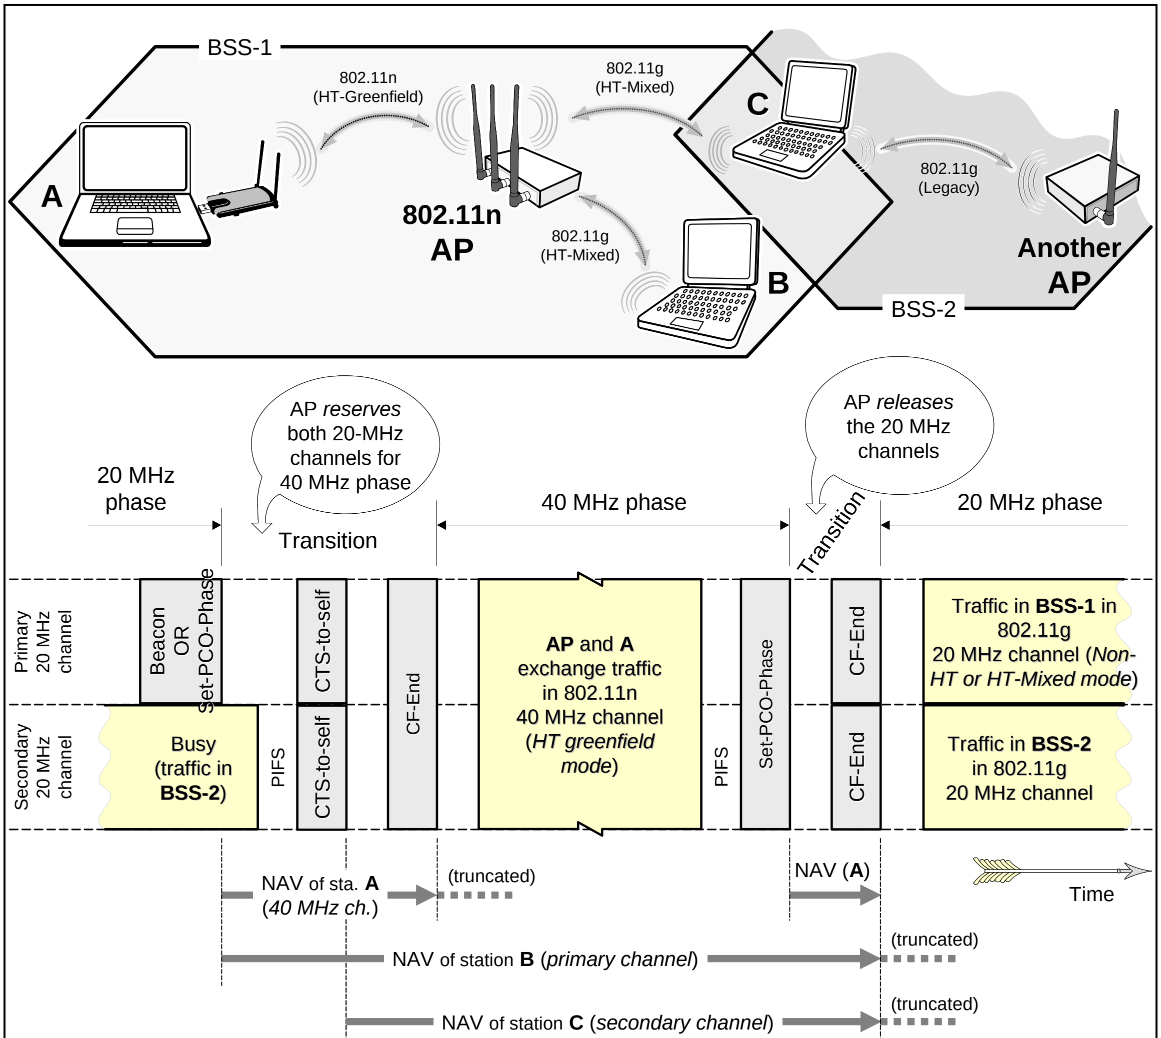

Figure 6-21: Example timing diagram for the (optional) 802.11n phased coexistence operation (PCO).

reservation of the secondary 20-MHz channel exceeds a given threshold, the access point may decide to terminate the transition process and go back to the 20-MHz phase.

Phased coexistence operation (PCO) makes an 802.11n access point more tolerant of nearby legacy APs operating on overlapping channels and might improve throughput in some situations. However, once again, this option reduces throughput due to transmission of many control frames (CTS-to-self and CF-End). In addition, switching back and forth between channels could potentially increase delay jitter for data frames, and therefore PCO mode would not be recommended for real-time multimedia traffic.

The cost of backwards compatibility features is additional overhead on every 802.11n transmission. This reduces the benefits of all the 802.11n improvements, resulting in significantly lower effective throughput by 802.11n devices in mixed environments. The HT-Mixed format will most likely be the most commonly used frame format because it supports both HT and legacy 802.11a/g devices. The HT-Mixed format is also considered mandatory and transmissions can occur in both 20-MHz and 40-MHz channels. It can be expected that protection mechanisms will be in use in the 2.4-GHz band (802.11b and 802.11g) until nearly every legacy device has disappeared. This is because there are too few channels available in that band to effectively overlay pure 802.11n wireless LANs in the same areas as legacy 2.4-GHz W-LANs. Given the larger number of channels available in the 5-GHz band in many countries, it is possible that two completely separate W-LANs could be operating in the same area in the 5-GHz band, with 802.11a operating on one set of channels and 802.11n operating on a different, nonintersecting set of channels. This would allow 802.11n to operate in pure high-throughput mode (HT-Greenfield mode), achieving the highest effective throughput offered by the 802.11n standard.

6.3.2 WiMAX (IEEE 802.16)

The IEEE 802.16 standard is also known as WiMAX, which stands for Worldwide Interoperability for Microwave Access. WiMAX, as approved in 2001, addressed frequencies from 10 to 66 GHz, where extensive spectrum is available worldwide but at which the short wavelengths introduce significant deployment challenges. A new effort will extend the air interface support to lower frequencies in the 2–11 GHz band, including both licensed and license-exempt spectra. Compared to the higher frequencies, such spectra offer the opportunity to reach many more customers less expensively, although at generally lower data rates. This suggests that such services will be oriented toward individual homes or small to medium-sized enterprises.

Medium Access Control (MAC) Protocol

The IEEE 802.16 MAC protocol was designed for point-to-multipoint broadband wireless access applications. It addresses the need for very high bit rates, both uplink (to the Base Station) and downlink (from the BS). Access and bandwidth allocation algorithms must accommodate hundreds of terminals per channel, with terminals that may be shared by multiple end users.

6.3.3 ZigBee (IEEE 802.15.4)

ZigBee is a specification for a suite of high level communication protocols using small, low-power digital radios based on the IEEE 802.15.4-2003 standard for wireless personal area networks (WPANs), such as wireless headphones connecting with cell phones via short-range radio. The technology defined by the ZigBee specification is intended to be simpler and less expensive than other WPANs, such as Bluetooth (Section 6.3.4). ZigBee is targeted at radio frequency (RF) applications that require a low data rate, long battery life, and secure networking.

6.3.4 Bluetooth

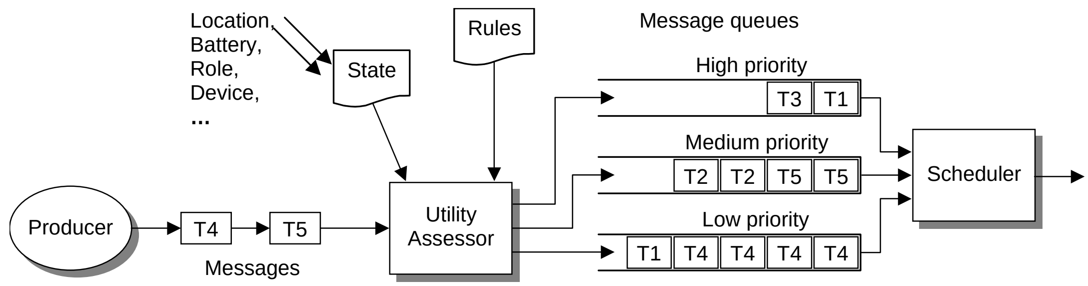

Figure 6-22: Priority-based scheduling of the messages generated by a producer. Messages are labeled by data types of the data they carry (T_1, …, T_5).

6.4 Wi-Fi Quality of Service

There is anecdotal evidence of W-LAN spectrum congestion; Unlicensed systems need to scale to manage user "QoS." Density of wireless devices will continue to increase; ~10x with home gadgets; ~100x with sensors/pervasive computing

Decentralized scheduling

We assume that each message carries the data of a single data type. The messages at the producer are ordered in priority-based queues. The priority of a data type is equal to its current utility, $U(T_i \mid S_j)$. Figure 6-22 shows the architecture at the producer node. Scheduler works in a round-robin manner, but may have different strategies for sending the queued messages, called queuing discipline. It may send all high priority messages first, or it may assign higher probability of sending to the high-priority messages, but the low-priority messages still get non-zero probability of being sent.

It is not clear whose rules for assigning the utilities should be used at producers: producer's or consumer's. If only the consumer's preferences are taken into the account, this resembles to the filtering approach for controlling the incoming information, e.g., blocking unsolicited email messages. One of the drawbacks of filtering is that does not balance the interests of senders and recipients: filtering is recipient-centric and ignores the legitimate interests of the sender [**Error! Reference source not found.**]. This needs to be investigated.

6.5 Summary and Bibliographical Notes

A collection of articles on mobile ad hoc networks, particularly the routing protocol aspect, is available in [Perkins, 2001]. [Murthy & Manoj, 2004] provide a comprehensive overview of ad hoc networks.

The major enhancement in IEEE 802.11e MAC protocol is providing Quality-of-Service (QoS), which is lacking in the legacy IEEE 802.11 MAC protocol. In IEEE 802.11e, enhanced distributed channel access (EDCA) is introduced to enhance legacy IEEE 802.11 DCF operation. EDCA is a contention-based channel access mechanism. QoS support is provided with different access categories (ACs). Four ACs are used in EDCA, each with an independent backoff mechanism and contention parameters. The parameters of ACs are set differently to provide differentiated QoS priorities for ACs.

The Block ACK has a potential to be more efficient than the regular ACK policy. However, the Basic Block ACK frame (defined in IEEE 802.11e and adopted in 802.11n) includes a Block ACK Bitmap of 128 bytes, and the efficiency of the Block ACK might be seriously compromised when the number of frames acknowledged by a Block ACK is small or the frames are not fragmented.

802.11n

The objective of IEEE 802.11n standard is to increase the throughput beyond 100 Mbps as well as extending the effective range from previous 802.11a/b/g standards. Use of Multiple Input Multiple Output (MIMO) technology along with OFDM (MIMO-OFDM) and doubling the channel bandwidth from 20-MHz to 40-MHz helps increase the physical (PHY) rate up to 600 Mbps. The data rates supported in an 802.11n network range from 6.5 Mbps to 600 Mbps. Support for 2 spatial streams is mandatory at the access point and up to 4 spatial streams can be used. The PHY layer enhancements are not sufficient to achieve the desired MAC throughput of more than 100 Mbps due to rate-independent overheads. To overcome this limitation, frame aggregation at the MAC layer is used in 802.11n to improve the efficiency. In the aggregate MSDU (A-MSDU) scheme, multiple MSDUs form a MPDU i.e., an 802.11n MAC-level frame (A-MSDU) consists of multiple subframes (MSDUs), either from different sources or destinations. The aggregate MPDU (A-MPDU) scheme can be used to aggregate multiple MPDUs into a single PSDU. Both aggregation schemes have their pros and cons along with associated implementation aspects. However, most product manufacturers support both features.

802.11b and 802.11g devices operate in the 2.4 GHz band and the 5 GHz band is used by 802.11a devices. The 802.11n-based devices can operate in both bands and hence backward compatibility with the respective legacy devices in the bands are an important feature of the standard. Most of the benefits of 802.11n will only be realized when 802.11n-capable clients are used with similar infrastructure, and even a few legacy (802.11a/b/g) clients in the cell will drastically reduce overall performance compared to a uniform 802.11n network. For quite a long time, 802.11n will need to operate in the presence of legacy 802.11a, b, and g devices. This mixed-mode operation will continue until all the devices in an area have been upgraded or replaced with 802.11n devices. However, sometimes protection mechanism is needed even in an 802.11n-only network as the devices can have different capabilities. Hence, protection schemes are not only used for coexistence with legacy devices but also for interoperability with various different operating modes of 802.11n devices. Each protection mechanism has a different impact on the performance and 802.11n devices will operate more slowly when working with legacy Wi-Fi devices.

802.11n and HT technology is so complex that an entire book dedicated to the topic would probably not be able to cover fully every aspect of HT. Section 6.3.1 highlight only some of the key features of 802.11n. MIMO technology is only briefly reviewed. 802.11n link adaptation and receiver feedback information is not reviewed at all. Other issues that were not covered include security, power management, and quality-of-service (QoS). The reader should consult other sources for these topics.

Xiao and Rosdahl [2002; 2003] have shown that control overhead is a major reason for inefficient MAC. The overhead is large either when the data rate is high or when the frame size is small. Throughput in 802.11 has an upper bound even the data rate goes to infinity. [Xiao, 2005] …

Thornycroft [2009] offers a readable introduction to 802.11n. Perahia and Stacey [2008] provide an in-depth review of 802.11n. The reader should consult the IEEE 802.11n standard for a more technical walk-through of the newly introduced enhancements.

Wang and Wei [2009] investigated the performance of the IEEE 802.11n MAC layer enhancements: frame aggregation, block acknowledgement, and reverse direction (RD) protocol. They conclude that VoIP performance is effectively improved with 802.11n MAC enhancements.

Problems

Chapter 7
Network Monitoring

Contents

7.1 Introduction

See: http://www.antd.nist.gov/

Wireless link of a mobile user does not provide guarantees. Unlike wired case, where the link parameters are relatively stable, stability cannot be guaranteed for a wireless link. Thus, even if lower-level protocol layers are programmed to perform as best possible, the application needs to know the link quality. The "knowledge" in the wired case is provided through quality guarantees, whereas here link quality knowledge is necessary to adapt the behavior.

Adaptation to the dynamics of the wireless link bandwidth is a frequently used approach to enhance the performance of applications and protocols in wireless communication environments [Katz 1994]. Also, for resource reservation in such environments, it is crucial to have the knowledge of the dynamics of the wireless link bandwidth to perform the admission control.

7.2 Available Bandwidth Estimation

In general, an accurate available bandwidth estimation is essential in monitoring if the different flows are living up to the required Quality-of-Service (QoS). For instance, streaming applications could adapt their sending rate to improve the QoS depending on a real-time knowledge of the end-to-end available bandwidth. Within a mobile-communications core network, the available bandwidth could also be used as an input to take decisions concerning issues such as load control, admission control, handover and routing. However, the scale of the different systems, the different traffic characteristics and the diversity of network technologies make this characterization of the end-to-end available bandwidth a challenging task.

One possible way to implement available bandwidth estimation would be to deploy special software or hardware on each router of the network. However, the cost in time and money of new equipment, maintenance of new nodes and software development makes this impractical. Moreover, this wide-scale deployment of specialized routers, which are continuously reporting bandwidth properties, might overwhelm the network. Another limitation is the lack of control over hosts and routers across autonomous domains.

An alternative is to run software on end hosts, which is usually called **active probing**. In this approach, the available bandwidth is inferred rather than directly measured. Ideally, a probing scheme should provide an accurate estimate as quickly as possible, while keeping the increased load on the network to the necessary minimum. There are several obstacles for measuring the available bandwidth by active probing. First, the available bandwidth is a time-varying metric. Second, the available bandwidth exhibits variability depending on the observing time-scale. Third, in the current networks increasingly intelligent devices are being used for traffic prioritization.

A **narrow link** (bottleneck) is a communication link with a small upper limit on the bandwidth. Conversely, a **tight link** (overused) is a communication link with a small available bandwidth but the overall bandwidth may be relatively large.

7.2.1 Packet-Pair Technique

A *packet pair* consists of two packets, usually of the same size, that are sent back-to-back via a network path. Unlike the techniques mentioned in the previous section, packet-pair probing directly gives a value for the capacity of the *narrow link*, with no additional information about the capacities of other links on the path. They assume FIFO queuing model in network routers and probe packets could be ICMP (Internet Control Message Protocol) packets. Packet-pair techniques for bandwidth measurement are based on measuring the transmission delays that packets suffer on their way from the source to the destination. The idea is to use inter-packet time to estimate the characteristics of the bottleneck link. If two packets travel together so that they are queued as a pair at the bottleneck link with no packet intervening between them, then their inter-packet spacing is proportional to the time needed to transmit the second packet of the pair on the bottleneck link (**Error! Reference source not found.**).

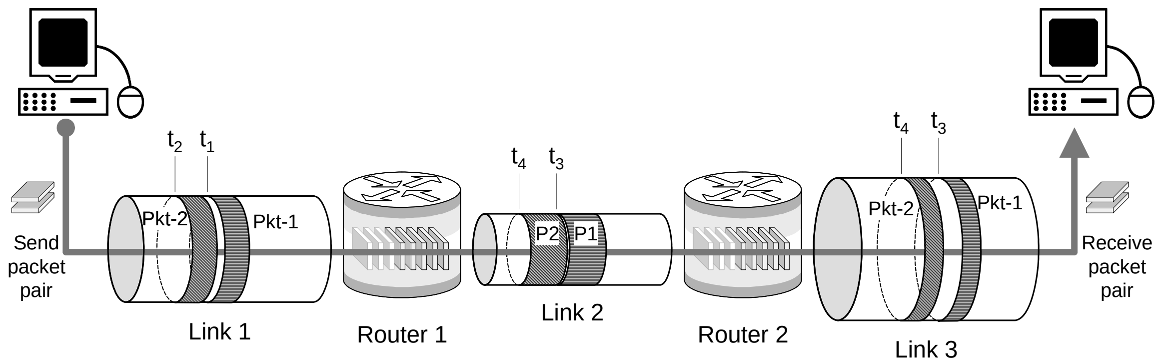

Figure 7-1: Packet-pair technique for bandwidth measurement. The packet spacing on the bottleneck link (Link 2) will be preserved on downstream links (Link 3).

Recall that packet transmission delay is computed using Eq. (1.2) in Section 1.3 as $t_x = L/R$, where L is the packet length and R is the link transmission rate. Le us assume that the receiver measures the difference of arrival times for a packet pair as $\Delta = t_4 - t_3$. Then, the transmission rate of the bottleneck link (i.e., link bandwidth), b, can be computed as:

$$b = L/\Delta \tag{6.1}$$

Note that the inter-packet spacing at the source must be sufficiently small so that:

$$t_2 - t_1 \leq L/b = t_4 - t_3 \tag{6.2}$$

The *packet-pairs* technique is useful for measuring the bottleneck capacity, i.e., the minimum capacity along the path. The main advantage of this method is that it performs the measurement of inter-arrival times between packets only at the end host. This fact avoids the problem of asymmetric routing, ICMP dependency and link-layer effects of RTT-based Capacity Estimation methods. On the other hand, this technique is very sensitive, not only to the probing packet size and user time resolution, but also to the cross-traffic.

PathMon [Kiwior, et al., 2004]

7.3 Dynamic Adaptation

Holistic QoS, system level

Computation and communication delays

Combinatorial optimization

Adaptive Service Quality

It may be that only two options are offered to the customers by the server: to be or not to be processed. In other words, quality of servicing is offered either in the fullest or no servicing at all. But, it may be that the server offers different options for customers "in hurry." In this case, we can speak of different *qualities of service*—from no service whatsoever, through partial service, to full service. The spectrum of offers may be discrete or continuous. Also, servicing options may be *explicitly* known and advertised as such, so the customer simply chooses the option it can afford. The other option is that servicing options are *implicit*, in which case they could be specified by servicing time or cost, or in terms of complex circumstantial parameters. Generally, we can say that the customer specifies the *rules* of selecting the quality of service in a given rule specification language.

Associated with processing may be a cost of processing. Server is linked with a certain resource and this resource is limited. Server capacity C expresses the number of customers the server can serve per unit of time and it is limited by the resource availability.

Important aspects to consider:

- Rules for selecting QoS
- Pricing the cost of service
- Dealing with uneven/irregular customer arrivals
- Fairness of service
- Enforcing the policies/agreements/contracts
    - Admission control
    - Traffic shaping

7.3.1 Data Fidelity Reduction

Compression

Simplification

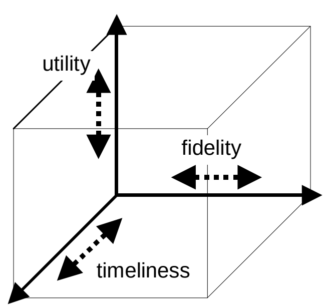

Figure 7-2: Dimensions of data adaptation.

Abstraction

Conversion (different domains/modalities)

We consider the model shown in Figure 6-?? where there are multiple clients producing and/or consuming dynamic content. Some shared content may originate or be cached locally while other may originate from remote and change with time. The data that originates locally may need to be distributed to other clients. The clients have local computing resources and share some global resources, such as server(s) and network bandwidth, which support information exchange. Although producers and consumers are interrelated, it is useful to start with a simpler model where we consider them independently to better understand the issues before considering them jointly. We first consider individual clients as *data consumers* that need to visualize the content with the best possible quality and provide highest interactivity. We then consider clients as *data producers* that need to update the consumers by effectively and efficiently employing global resources.

We will develop a formal method for maximizing the utility of the shared content given the limited, diverse, and variable resources. Figure 8-1 illustrates example dimensions of data adaptation; other possible dimensions include modality (speech, text, image, etc.), security, reliability, etc. The user specifies the rules R for computing the utilities of different data types that may depend on contextual parameters. We define the *state* of the environment as a touple containing the status of different environmental variables. For example, it could be defined as: state = (time, location, battery energy, user's role, task, computer type). The location may include both the sender and the receiver location. Given a state S_j the utility of a data type T_i is determined by applying the user-specified rules: $U(T_i \mid S_j) = R(T_i \mid S_j)$. We also normalize the utilities because it is easier for users to specify relative utilities, so in a given state S_j the utilities of all data types are: $\sum_i U(T_i \mid S_j) = 1$.

Our approach is to vary the fidelity and timeliness of data to *maximize the sum of the utilities* of the data the user receives. Timeliness is controlled, for example, by the parameters such as update frequency, latency and jitter. Fidelity is controlled by parameters such as the detail and accuracy of data items and their structural relationships. Lower fidelity and/or timeliness correspond to a lower demand for resources. Our method uses nonlinear programming to select those values for fidelity and timeliness that maximize the total data utility, subject to the given resources. Note that the user can also require fixed values for fidelity and/or timeliness, and seek an optimal solution under such constraints.

7.3.2 Application Functionality Adaptation

7.3.3 Computing Fidelity Adaptation

Review CMU-Aura work

7.4 Summary and Bibliographical Notes

When deploying a real-time or multimedia application, you need to establish a thorough baseline of current network activity on all segments that will host the application. You need to understand the degree to which latency, jitter and packet loss affect your network before deploying a real-time or multimedia application. You must understand current network load and behavior, including any areas where latency is elevated or highly variable. In many networks, traffic loads may vary substantially over time. As loads increase, inconsistent packet delivery rates are probable. Thus, increasing loads form the foundation for excessive latency and jitter—which are two of the most prevalent inhibitors for consistent application performance. When collecting baseline metrics, remember that network behavior varies widely as various business activities occur. Be sure to create a baseline that reflects all major phases and facets of your network's activities.

V. Firou, J. Le Boudec, D. Towsley, and Z. Zhang, "Theories and Models for Internet Quality of Service," Proceedings of the IEEE, Special Issue on Internet Technology, August 2002.

Various forms of the packet-pair technique were studied by [Bolot, 1993], [Carter & Crovella, 1996], [Paxson, 1997a], and [Lai & Baker, 1999].

[Jain & Dovrolis, 2002] investigated how to deal with cross-traffic, using statistical methods

[Hu & Steenkiste, 2003], available bandwidth discovery: Initial Gap Increasing (IGI)
estimate both upper limit and background traffic and subtract both. Problem: presumes that the bottleneck link is also the tight link

PathMon [Kiwior, et al., 2004]

Problems

Chapter 8
Internet Protocols

Contents

This chapter describes several important protocols that are used in the current Internet. I feel that these protocols are not critical for the rest of this text. However, the reader may feel otherwise, so I included them for completeness. Also, a student new to the field may wish to know about practical implementations of the concepts and algorithms described in the rest of this text.

Although this chapter is about the Internet protocols, the key Internet protocols are not reviewed here. Because of their great importance, IP and TCP are described early one, in Chapters 1 and 2, respectively.

8.1 Internet Protocol Version 6 (IPv6)

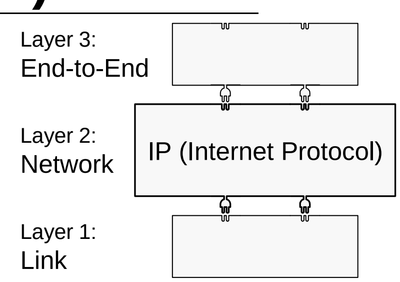

Internet Protocol version 4 (Sections 1.4.1 and 1.4.4) was first developed in the 1970s and the main document that defines IPv4 functionality (RFC-791) was published in 1981. Many things could not be envisioned at such early stage of Internet development, especially shortage of address space availability. Internet Protocol version 6 (IPv6) was developed primarily to

Visit http://en.wikipedia.org/wiki/Internet_reference_model for more details on the Internet reference model

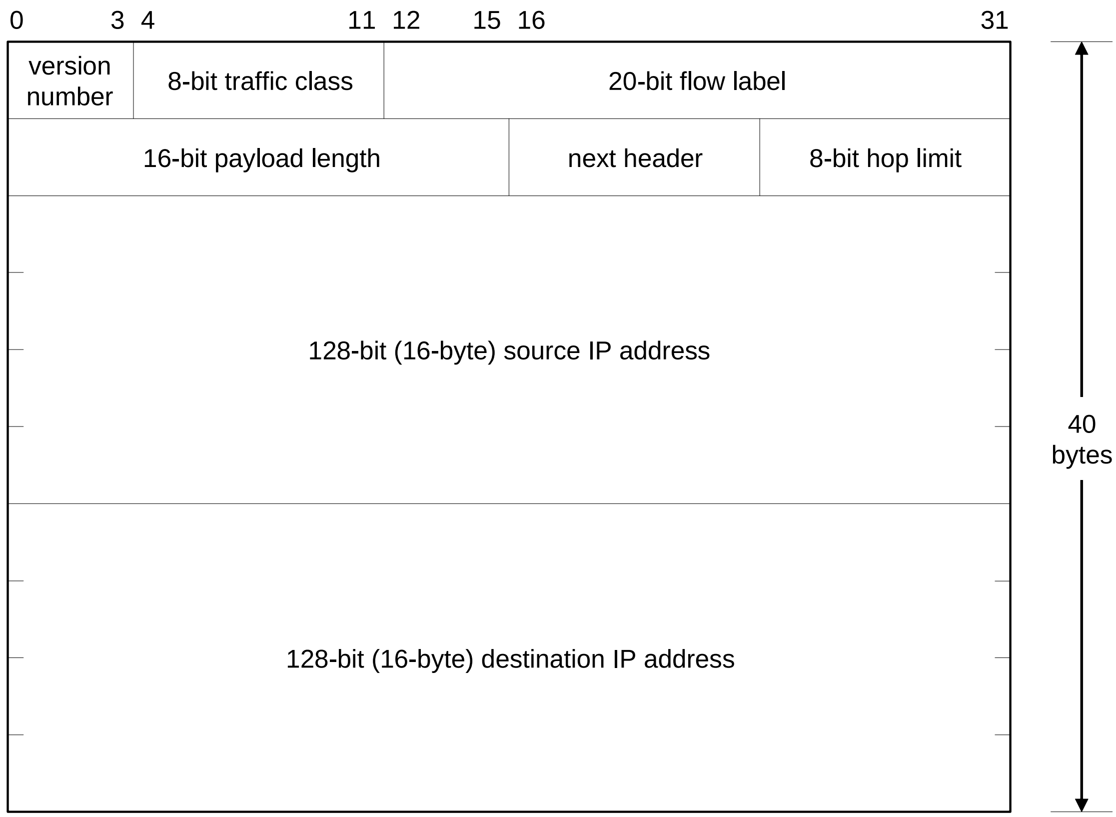

Figure 8-1: The header format of IPv6 datagrams.

address the rapidly shrinking supply of IPv4 addresses, but also to implement some novel features based on the experience with IPv4.

Figure 8-1 shows the format of IP version 6 datagram headers. It is much simpler than the IPv4 header (Figure 1-36), but also double the size of a default IPv4 header, mainly because of longer IP addresses in IPv6. The IPv6 header fields are as follows:

Version number: This field indicates version number, and the value of this field for IPv6 datagrams is 6.

Traffic class: The 8-bit Traffic Class field is available for use by originating nodes or forwarding routers to identify and distinguish between different classes or priorities of IPv6 packets. This filed is equivalent to IPv4 Type-of-Service (Figure 1-36), and it is intended to provide various forms of “differentiated service” or DiffServ for IP packets (Section 3.3.5). Experiments are currently underway to determine what sorts of traffic classifications are most useful for IP packets.

Flow label: The 20-bit Flow Label field may be used by a source to label packets for which it requests special handling by the IPv6 routers, such as non-default quality of service or “real-time” service. This feature is still experimental and evolves as more experience is gained with IntServ flow support in the Internet (Section 3.3.4). Hosts or routers that do not support the functions of the Flow Label field are required to set the field to zero when originating a packet; pass the field on unchanged when forwarding a packet; and ignore the field when receiving a packet.

Payload length: This field tells how many bytes follow the 40-byte header in the IP datagram. Unlike IPv4 Datagram Length, IPv6 Payload Length does not count the datagram header.

Next header: This 8-bit selector identifies the type of header immediately following the IPv6 header within the IPv6 datagram. Currently, there are six extension headers (optional), which may follow the main IPv6 header. If a header is the last IP header, this field identifies the type of the upper-layer protocol to pass the payload to at the destination, and uses the same values as the IPv4 "User Protocol" field (Figure 1-36). See more about extension headers in Section 8.1.2.

Hop limit: This 8-bit unsigned integer specifies how long a datagram is allowed to remain in the Internet, to catch packets that are stuck in routing loops. It is decremented by one by each node that forwards the packet. The packet is discarded if Hop Limit is decremented to zero. It is equivalent to the "Time-to-Live" (TTL) field in IPv4 datagrams.

Source IP address: This address identifies the end host that originated the datagram.

Destination IP address: This is the IP address of the intended recipient of the packet (possibly not the ultimate recipient, if a *Routing header* is present, as described in the next section).

Notice that there are none of the fields related to packet fragmentation from IPv4 (Figure 1-36). This is because IPv6 takes a different approach to fragmentation. To simplify the work of routers and speed up their performance, IPv6 assumes that router do not perform any fragmentation. IPv6 requires that every link in the Internet have a maximum transmission unit (MTU) of 1280 bytes or greater. On any link that cannot convey a 1280-byte packet in one piece, link-specific fragmentation and reassembly must be provided at a protocol layer below IPv6.

Links that have a configurable MTU (for example, PPP links Section 1.5.1) must be configured to have an MTU of at least 1280 bytes. It is strongly recommended that IPv6 nodes implement Path MTU Discovery (described in RFC-1981), in order to dynamically discover and take advantage of path MTUs greater than 1280 bytes. This rule makes fragmentation less likely to occur in the first place. In addition, when a host sends an IPv6 packet that is too large, instead of fragmenting it, the router that is unable to forward it drops the packet and sends back an error message. This message tells the originating host to break up all future packets to that destination.

8.1.1 IPv6 Addresses

The IPv6 address space is 128-bits (2^{128}) in size, which translates to the exact number of 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses. That seems like enough for all purposes that currently can be envisioned.

A new notation (*hexadecimal colon notation*) has been devised to writing 16-byte IPv6 addresses. A 128-bit address is divided into eight sections, each two bytes long. It is written as eight groups of four hexadecimal digits (total 32) with colons between the groups, like this:

2000:0000:0000:0000:0123:4567:89AB:CDEF

Some simplifications have been authorized for special cases. For example, if an address has a large number if consecutive zeros, the zero fields can be omitted and replaced with a double colon "::".The above example address can be written compactly as 2000::123:4567:89AB:CDEF. Notice that only leading or intermediary zeroes can be abbreviated, but not the trailing zeroes.

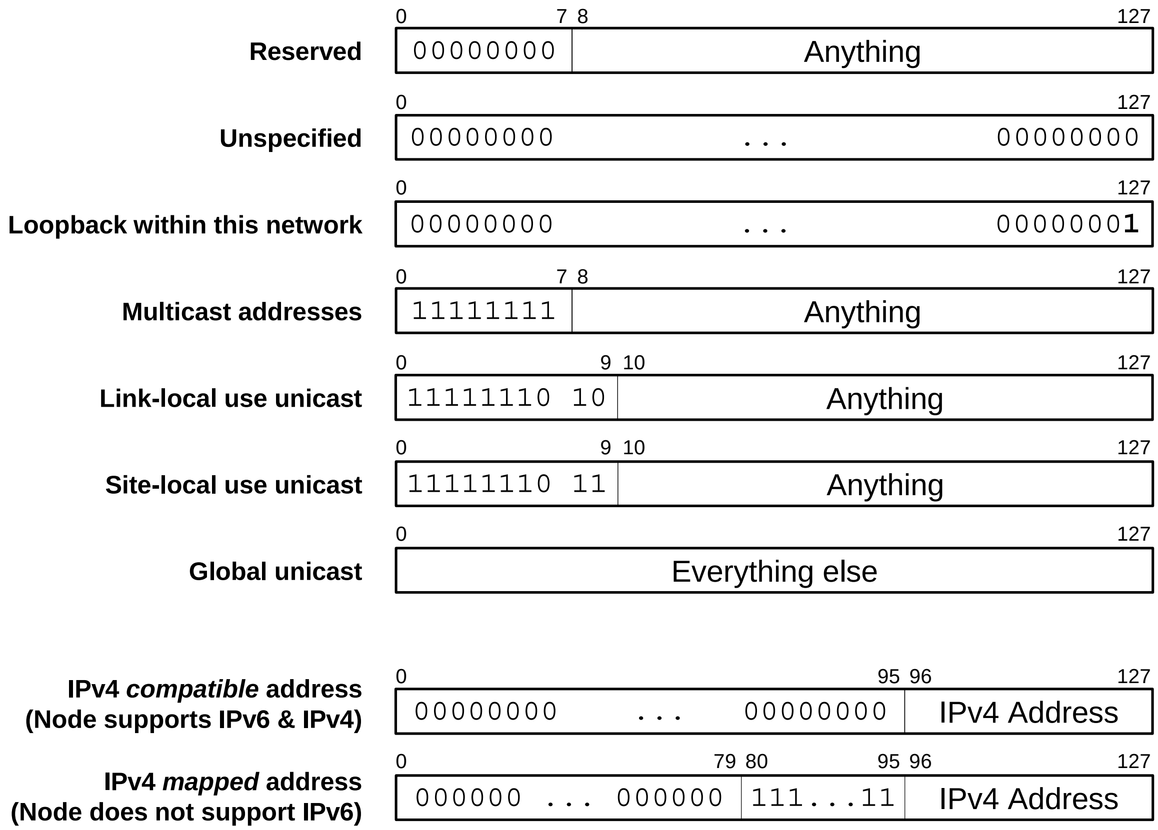

Figure 8-2: Selected address prefix assignments for IPv6, excerpted from RFC-2373.

Also, this type of abbreviation is allowed only once per address; if there are two runs of zeroes, only one of the can be abbreviated.

As with IPv4 CIDR scheme, the notation A/m designates a subset of IPv6 addresses (subnetwork) where A is the prefix and the mask m specifies the number of bits that designate the subset, beginning from left to right. For example, the notation: 2000:0BA0:01A0::/48 implies that the part of the IPv6 address used to represent the subnetwork has 48 bits. Because each hexadecimal digit has 4 bits, the prefix representing the subnetwork is formed by 48/4 = 12 digits, that is: "2000:0BA0:01A0." The remaining (128 − 48)/4 = 20 digits would be used to represent the network interfaces inside the subnetwork.

Similar to CIDR in IPv4 (Section 1.4.4), IPv6 addresses are classless. However, the address space is hierarchically subdivided depending on the leading bits of an address. A variable number of leading bits specify the **type prefix** that defines the purpose of the IPv6 address. To avoid ambiguity, the prefix codes are designed such that no code is identical to the first part of any other code. The current assignment of prefixes is shown in Figure 8-2. Notice that two special addresses ("unspecified address" and loopback address) are assigned out of the reserved 00000000-format prefix space. The address 0:0:0:0:0:0:0:0 is called the *unspecified address*. It indicates the absence of an address, and must never be assigned to any node as a destination address. However, it may be used by a host in the Source Address field of an IPv6 packet initially, when the host wants to learn its own address. The unicast address 0:0:0:0:0:0:0:1 is called the *loopback address*. It may be used by a node to send an IPv6 packet to itself. It may never be assigned to any physical interface. It may be thought of as being associated with a virtual interface (e.g., the loopback interface).

The IPv6 addresses with embedded IPv4 addresses are assigned out of the reserved 00000000-format prefix space. There are two types of IPv6 addresses that contain an embedded IPv4 address (see bottom of Figure 8-2). One type are the so-called *IPv4-compatible addresses*. These addresses are used by IPv6 routers and hosts that are directly connected to an IPv4 network and use the "tunneling" approach to send packets over the intermediary IPv4 nodes (Section 8.1.3). This address format consists of 96 bits of 0s followed by 32 bits of IPv4 address. Thus, and IPv4 address of 128.6.29.131 can be converted to an IPv4-compatible ::128.6.29.131. The other type is so called *IP-mapped addresses*. These addresses are used to indicate IPv4 nodes that do not support IPv6. The format of these addresses consist of 80 bits of 0s, followed by 16 bits of 1s, and then by 32 bits of IPv4 address. An example would be written as ::FFFF:128.6.29.131.

IPv6 allows three types of addresses:

- **Unicast:** An identifier for a single network interface. A packet sent to a unicast address should be delivered to the interface identified by that address.
- **Anycast:** A prefix identifier for a set of network interfaces. These typically belong to different nodes, with addresses having the same subnet prefix. A packet sent to an anycast address should be delivered to only one of the interfaces identified by that prefix. The interface selected by the routing protocol is the "nearest" one, according to the protocol's distance metrics. For example, all the routers of a backbone network provider could be assigned a single anycast address, which would then be used in the routing header. One expected use of anycasting is "fuzzy routing," which means sending a packet through "one router of network *X*." The anycast address will also be used to provide enhanced routing support to mobile hosts.
- **Multicast:** An identifier for a set of network interfaces, typically belonging to different nodes that may or may not share the same prefix. A packet sent to a multicast address should be delivered to all the interfaces identified by that address.

There are no broadcast addresses in IPv6; their function is taken over by multicast addresses.

It is anticipated that unicast addressing will be used for the vast majority of traffic under IPv6, just as is the case for older one, IPv4. It is for this reason that the largest of the assigned blocks of the IPv6 address space is dedicated to unicast addressing. RFC-2374 assigned the Format Prefix 2000::/3 (a "001" in the first three bits of the address) to unicast addresses. However, RFC-3587 invalidated this restriction. Although currently only 2000::/3 is being delegated by the IANA, implementations should not make any assumptions about 2000::/3 being special. In the future, the IANA might be directed to delegate currently unassigned portions of the IPv6 address space for the purpose of Global Unicast as well.

Figure 8-3(a) shows the general format for IPv6 global unicast addresses as defined in RFC-3587. The *Global Routing Prefix* is a (typically hierarchically-structured) value assigned to a site (a cluster of subnets or links), the *Subnet ID* is an identifier of a subnet within the site, and the *Interface ID* identifies the network interfaces on a link. The global routing prefix is designed to be structured hierarchically by the Regional Internet Registries (RIRs) and Internet Service

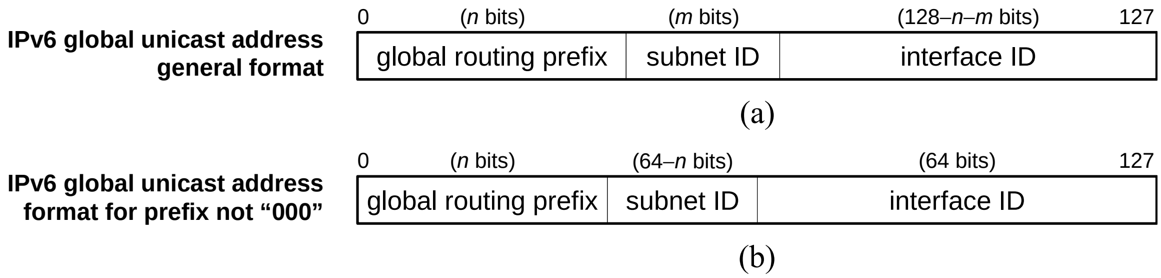

Figure 8-3: Format of IPv6 global unicast addresses.

Table 8-1: IPv6 extension headers.

Extension header	Description
Hop-by-hop options	Miscellaneous information for routers.
Destination options	Additional information for the destination node.
Routing	Loose list of routers to visit (similar to IPv4 source routing).
Fragmentation	Information about datagram fragmentation and reassembly.
Authentication	Verification of the sender’s identity.
Encrypted security payload	Information about the encrypted contents.

Providers (ISPs). The subnet-ID field is designed to be structured hierarchically by site administrators.

RFC-3587 also requires that all unicast addresses, except those that start with binary value “000,” have Interface IDs that are 64-bits long and to be constructed in Modified 64-bit Extended Unique Identifier (EUI-64) format. The format of global unicast address in this case is shown in Figure 8-3(b). This includes global unicast address under the 2000::/3 prefix (starting with binary value “001”) that is currently being delegated by the IANA.

An IPv6 Address [RFC-4291] may be administratively assigned using DHCPv6 [RFC-3315] in a manner similar to the way IPv4 addresses are, but may also be autoconfigured, facilitating network management. Autoconfiguration procedures are defined in [RFC-4862] and [RFC-4941]. IPv6 neighbors identify each other’s addresses using either Neighbor Discovery (ND) [RFC-4861] or SEcure Neighbor Discovery (SEND) [RFC-3971].

8.1.2 IPv6 Extension Headers

IPv6 header is relatively simple (compared to IPv4), because features that are rarely used or less desirable are removed. However, some of these features occasionally are still needed, so IPv6 has introduced the concept of an optional *extension header*. These headers can be supplied to provide extra information, and are placed between the IPv6 header and the upper-layer header in a packet. Extension headers allow the extension of the protocol if required by new technologies or applications. Six kinds of extension headers are defined at present (Table 8-1). Each one is optional, and if present, each is identified by the Next Header field of the preceding header. If more than one is present, they must appear directly after the main header, and preferably in the

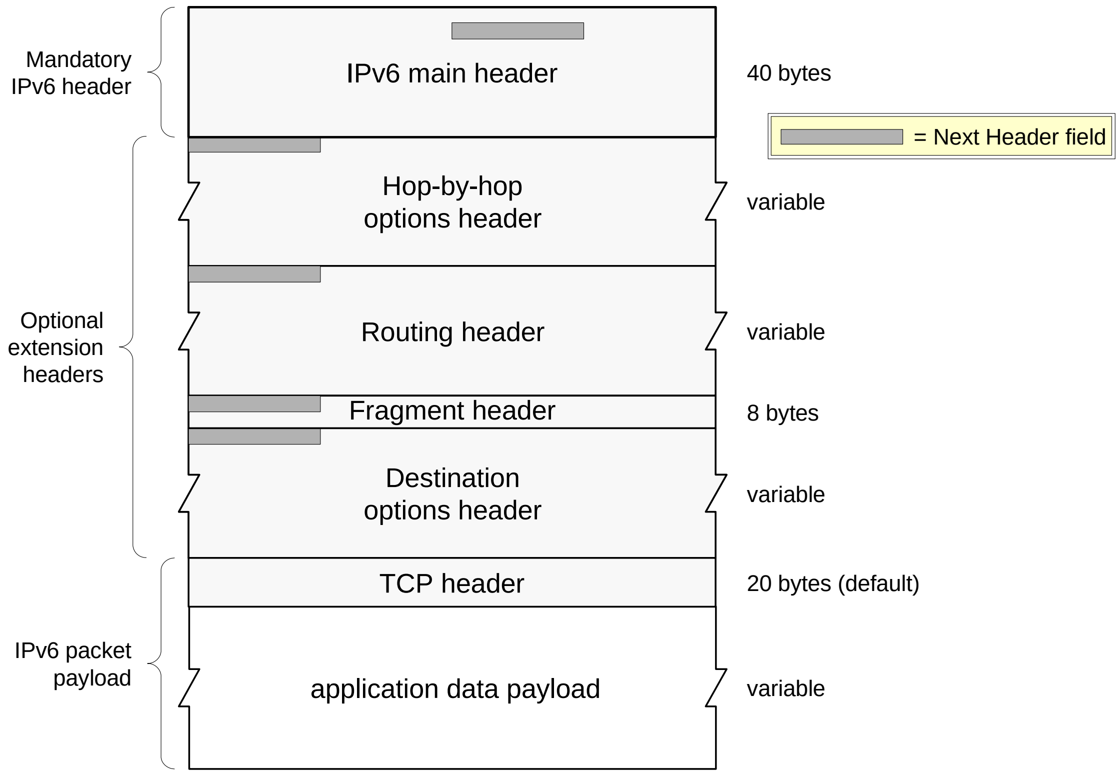

Figure 8-4: Example of IPv6 extension headers.

order shown in Table 8-1. After the extension headers follows the upper-layer header (e.g., TCP header), which is the header of the payload contained in this IPv6 datagram.

Figure 8-4 shows an example of an IPv6 datagram that includes an instance of each extension header, except those related to security. Note that the main IPv6 header and each extension header include a Next Header field. This field identifies the type of the immediately following header. If the next header is an extension header, then this field contains the type identifier of that header. Otherwise, this field contains the identifier of the upper-layer protocol to which the datagram will be delivered. In the latter case, the same values are used as for the IPv4 Protocol field. In the example in Figure 8-4, the upper-layer protocol is TCP and the payload carried by this IPv6 datagram is a TCP segment.

With one exception, extension headers are not examined or processed by any node along a packet's delivery path, until the packet reaches the node (or, in the case of multicast, each of the set of nodes) identified in the Destination Address field of the IPv6 header. There, regular demultiplexing on the Next Header field of the IPv6 header invokes the module to process the first extension header, or the upper-layer header if no extension header is present. The contents and semantics of each extension header determine whether or not to proceed to the next header. Therefore, extension headers must be processed strictly in the order they appear in the packet. A receiver must not, for example, scan through a packet looking for a particular kind of extension header and process that header prior to processing all preceding ones.

The exception referred to in the preceding paragraph is the Hop-by-Hop Options header, which carries information that must be examined and processed by every node along a packet's delivery path, including the source and destination nodes. The Hop-by-Hop Options header, when present,

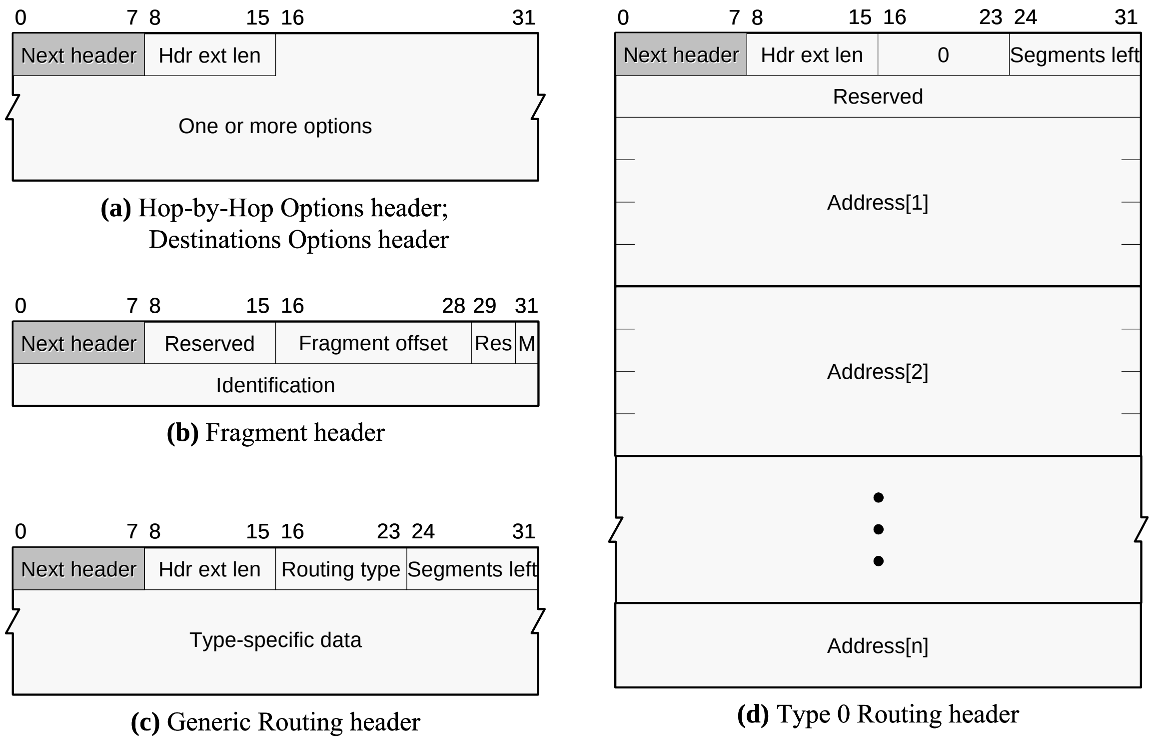

Figure 8-5: Format of IPv6 extension headers.

must immediately follow the main IPv6 header. Its presence is indicated by the value zero in the Next Header field of the main IPv6 header.

Hop-by-Hop Options Header

The Hop-by-Hop Options header carries optional information for the routers that will be visited by this IPv6 datagram. This header must be examined by every node along a packet's delivery path. The Hop-by-Hop Options header is identified by a Next Header value of 0 in the main IPv6 header, and has the following format (Figure 8-5(a)):

- **Next Header (8 bits):** Identifies the type of header immediately following the Hop-by-Hop Options header. Uses the same values as the IPv4 Protocol field.
- **Hdr Ext Len (8-bits):** Length of the Hop-by-Hop Options header in 64-bit units, not including the first 64 bits.
- **Options:** A variable-length field, of length containing one or more options, such that the complete Hop-by-Hop Options header is long an integer multiple of 64-bits. Each option is defined by three sub-fields: Option Type (8 bits), which identifies the option; Length (8 bits), which specifies the length of the Option Data field (in bytes); and Option Data, which is a variable-length specification of the option.

If, as a result of processing a header, a node is required to proceed to the next header but the Next Header value in the current header is unrecognized by the node, it should discard the packet and send an ICMP Parameter Problem message to the source of the packet, with an ICMP Code value of 1 ("unrecognized Next Header type encountered") and the ICMP Pointer field containing the offset of the unrecognized value within the original packet. The same action should be taken if a node encounters a Next Header value of zero in any header other than an IPv6 header.

8.1.3 Transitioning from IPv4 to IPv6

There are two approaches for gradually introducing IPv6 in the public Internet, which is based on IPv4:

- **Dual-stack approach:** IPv6 nodes also have a complete IPv4 implementation. Such a node, referred to as and IPv6/IPv4 node in RFC-4213, has the ability to send and receive both IPv4 and IPv6 datagrams.
- **Tunneling approach:** Any two IPv6 nodes that are connected via intermediary IPv4 routers (that are not IPv6-capable) create a "tunnel" between them. That is, the sending node takes the entire IPv6 datagram (header and payload included) and puts it in the data (payload) filed of an IPv4 datagram. This datagram is then addressed to the receiving IPv6 node and sent to the first intermediary node in the tunnel.

8.2 Routing Protocols

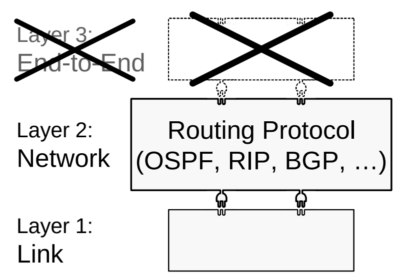

This section reviews several currently most popular Internet routing protocols.

8.2.1 Routing Information Protocol (RIP)

Routing Information Protocol (RIP) is a distance-vector routing protocol (Section 1.4.3), described in RFC-1058. Similar to OSPF (described next in Section 8.2.2) RIP is also used for routing within individual autonomous domains. Unlike OSPF, which scales to large intranets, RIP is useful for small subnets because of its simplicity of implementation and configuration, where its inadequacies are not prominent. RIP inadequacies include poor dealing with link failures and lack of support for multiple metrics. In addition, unlike OSPF, RIP for IP internetworks cannot be subdivided and no route summarization is done beyond the summarizing for all subnets of a network identifier. As a result, RIP networks are "flat."

The format of RIP version 2 route-advertisement packets is shown in Figure 8-6. The first four bytes of a RIP message contain the RIP header. The Command field is used to specify the purpose of this packet. For example, the possible values include: "1" which means a request for the receiver node to send all or part of its routing table; and "2" which symbolizes a response message containing all or part of the sender's routing table. This message may be sent in response to a request or poll, or it may be an update message generated by the sender.

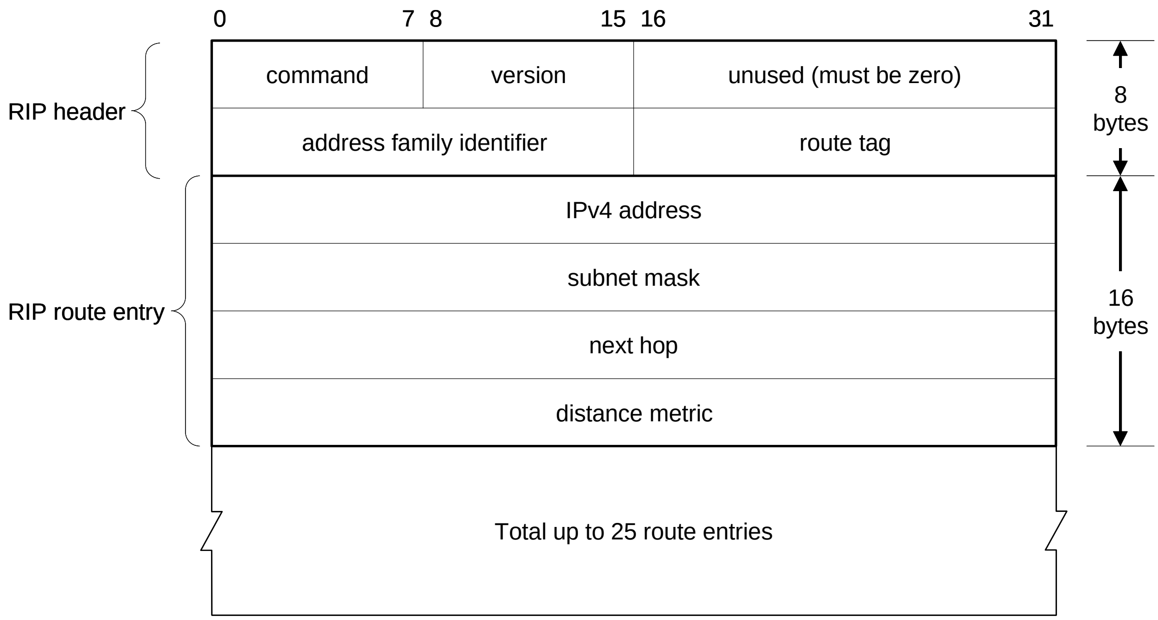

Figure 8-6: Routing Information Protocol (RIP) version 2 packet format for IPv4 addresses.

The Version field of the header specifies version number of the RIP protocol that the sender uses. The value should be "2" for RIP messages that use authentication or carry information in any of the newly defined fields (RIP version 2, defined in RFC-1723). The contents of the Unused field (two bytes) shall be ignored. The Address Family Identifier filed indicates what type of address is specified in the entries. The address family identifier for IPv4 is 4. The Route Tag field is an attribute assigned to a route that must be preserved and readvertised with a route. The intended use of the Route Tag is to provide a method of separating "internal" RIP routes (routes for networks within the RIP routing domain) from "external" RIP routes, which may have been imported from another routing protocol.

The remainder of a RIP message is composed of route entries. There may be between 1 and 25 route entries and each is 20 bytes long. The IP Address is the usual 4-byte IPv4 address. The Subnet Mask field contains the subnet mask that is applied to the IP address to yield the non-host portion of the address (recall Section 1.4.4). If this field is zero, then no subnet mask has been included for this entry.

The Next Hop field identifies the immediate next hop IP address to which packets to the destination (specified by the IP Address of this route entry) should be forwarded. Specifying a value of 0.0.0.0 in this field indicates that routing should be via the originator of this RIP advertisement packet. An address specified as a next hop must be directly reachable on the logical subnet over which the advertisement is made. The purpose of the Next Hop field is to eliminate packets being routed through extra hops in the system. This is particularly useful when RIP is not being run on all of the routers on a network, but some other routing protocols are used, as well. Note that Next Hop is an "advisory" field. That is, if the provided information is ignored, a possibly sub-optimal, but still valid, route may be taken. If the received Next Hop is not directly reachable, it should be treated as 0.0.0.0.

The Metric field of a route entry specifies the distance to the destination node identified by the IP Address of this entry. RIP takes the simplest approach to link cost metrics, a hop-count metric with all link costs being equal to 1. Valid distances are 1 through 15, with 16 defined as infinity. This limits RIP to running on relatively small networks, with no paths longer than 15 hops.

RIP, like most distance vector routing protocols, announces its routes in an unsynchronized and unacknowledged manner. Peer routers exchange distance vectors every 30 seconds, and a router is declared dead if a peer does not hear from it for 180 s, which is the hold-down timer period. RIP uses split horizon with poisoned reverse to tackle the counting-to-infinity problem.

Triggered updates allow a RIP router to announce changes in metric values almost immediately rather than waiting for the next periodic announcement. The trigger is a change to a metric in an entry in the routing table. For example, networks that become unavailable can be announced with a hop count of 16 through a triggered update. Note that the update is sent *almost immediately*, where a time interval to wait is typically specified on the router. If triggered updates were sent by all routers immediately, each triggered update could cause a cascade of broadcast traffic across the IP internetwork.

8.2.2 Open Shortest Path First (OSPF)

Open Shortest Path First (OSPF) is a link-state routing protocol (Section 1.4.2) that is currently the preferred protocol for interior routing—routing within individual *autonomous systems*, i.e., internetworks controlled by a single organization, also known as *intranets*. Autonomous systems (ASs) and inter-domain routing are described in Section 1.4.5.

The router broadcasts its link-state advertisements (LSAs) to *all* other routers in its autonomous system, not just to its neighboring routers. A router broadcasts link-state advertisements whenever there is a change in link status (e.g., outage or changed link cost). It also broadcasts a link's state periodically (at least once every 30 minutes), even if the link's state has not changed. An OSPF cost advertised in LSAs is a unitless metric that indicates the degree of preference for using a link. The network administrator can configure the cost of individual links to represent delay, data rate, monetary cost, or other factors.

Each router gathers the received LSAs into a database called the *link state database* (LSDB). By synchronizing LSDBs between all neighboring routers, each router has each other router's LSA in its database. Therefore, every router has the same LSDB. From the LSDB, entries for the router's routing table are calculated using the algorithm described in Section 1.4.2 to determine the least-cost path, the path with the lowest accumulated cost, to each network in the AS internetwork.

OSPF allows introducing additional level of hierarchy, in addition to autonomous systems. Each AS that runs OSPF can be configured into **areas**, where each area behaves like an independent network. Different areas exchange information via routers that belong to several areas, known as **area-border routers**. Each OSPF area runs its own OSPF protocol, and each router in an area broadcasts its LSAs only to routers within its area and each router's LSDB includes only the state of this area's links. Exactly one OSPF area is configured to act as the **backbone area** that routes traffic between the other areas within the same AS. There must be at least one area-border router in each area, connecting the area to the backbone. Each area-border router maintains several LSDBs, one for each area to which it belongs.

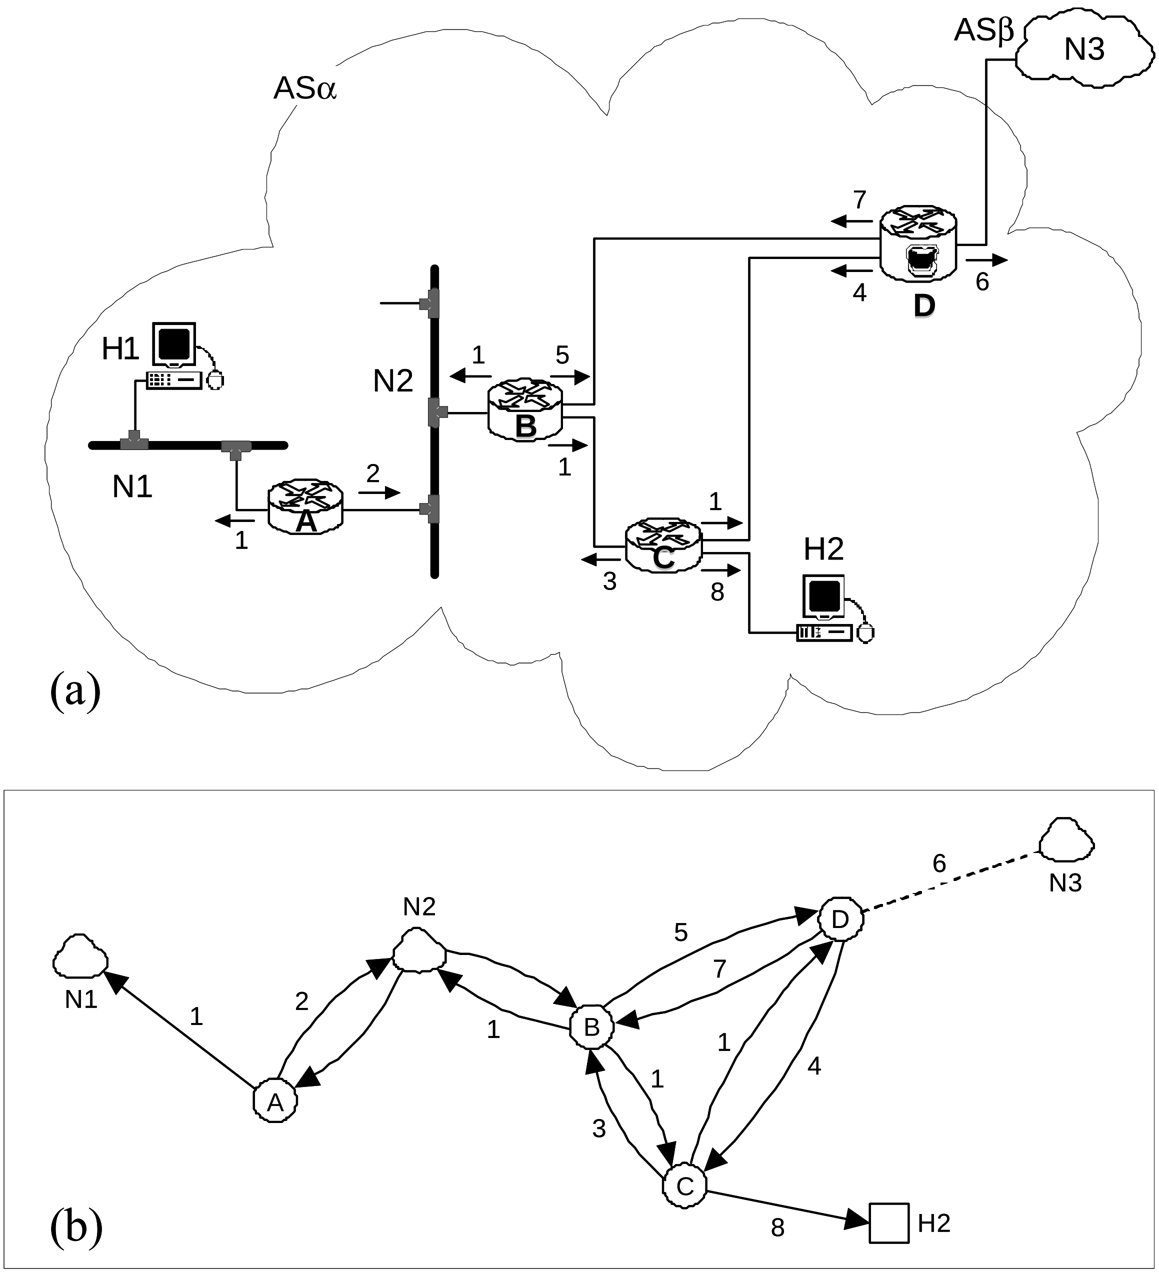

Figure 8-7: (a) Example of an autonomous system running OSPF. (b) Directed graph representation of the same AS. Notice that link costs may be asymmetric for both directions.

OSPF represents network topology as a *directed graph*. An example is shown in Figure 8-7, where autonomous system ASα is running OSPF as its interior gateway protocol. The vertices of the link-state database graph represent *routers* and *networks*. A graph edge connects two routers when they are attached via a physical point-to-point link. An edge connecting a router to a network indicates that the router has an interface on the network. Networks can be either *transit* or *stub* networks. **Transit network** is capable of carrying data traffic that is originated and destined externally to this network. A transit network is represented by a vertex having both incoming and outgoing edges. A stub network's vertex has only incoming edges. For example, in Figure 8-7(b), *N*2 is a transit network and *N*1 is a stub network. The mapping is as follows:

- Two routers connected by a point-to- point link are represented as two router vertices directly connected by a pair of edges, one in each direction. For example, in Figure 8-7 the cost of the edge from router *B* to *C* is “1” and from *C* to *B* is “3.”
- When several routers are attached to a broadcast network (transit network), the graph shows all routers bidirectionally connected to the network vertex. For example, in Figure 8-7 network *N*2 has routers *A* and *B* attached and therefore its edges in Figure 8-7(b) are bidirectional.
- If a network has only one attached router (i.e., a stub network), the network appears on the end of a stub connection in the graph. See, for example, network *N*1 in Figure 8-7.
- Hosts attached directly to routers appear on the graph as stub networks. See, for example, network *H*2 in Figure 8-7. (Host *H*1 is not shown in the graph because it is not attached directly to a router.)
- If a router is connected to other autonomous systems (so called “speaker node,” see Section 1.4.5), then the cost to each network in the other AS must be obtained from an exterior routing protocol, such as BGP (Section 8.2.3). Such a network is represented as a stub. For example, in Figure 8-7 router *D* is a speaker node and it is connected to network *N*3 in ASβ.

The cost of a link is associated with the output port of each router interface, so each end of the link may see the link cost differently. As already noted, the link cost is configurable by the network administrator. Arcs of the link-state database graph are labeled with the cost of the corresponding router output interface, as seen in Figure 8-7. Arcs without labels have a cost of zero. Arcs leading from transit networks to routers always have a cost of 0. See for example arcs from *N*2 to routers *A* and *B* in Figure 8-7(b).

All OSPF messages begin with the same header (Figure 8-8). Current (as of 2010) *Version* of the OSFP protocol is 2. There are five payload *Type*s of OSPF packets, as follows: 1 = Hello; 2 = Database Description; 3 = Link State Request; 4 = Link State Update; 5 = Link State Acknowledgment. The packet *Length* is given in bytes and includes the standard OSPF header. The *Address* identifies the source router of this packet. *Area ID* is a 32-bit number identifying the OSPF routing area to which the source router of this packet belongs. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labeled with the backbone Area ID of `0.0.0.0`. The *Checksum* field represents the standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit Authentication field. This checksum is calculated as the 16-bit one’s complement of the one’s complement sum of all the 16-bit words in the packet, excepting the Authentication field. If the packet’s length is not an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. *Authentication Type* identifies the authentication scheme to be used for the packet. Finally, *Authentication* is a 64-bit field for use by the authentication scheme.

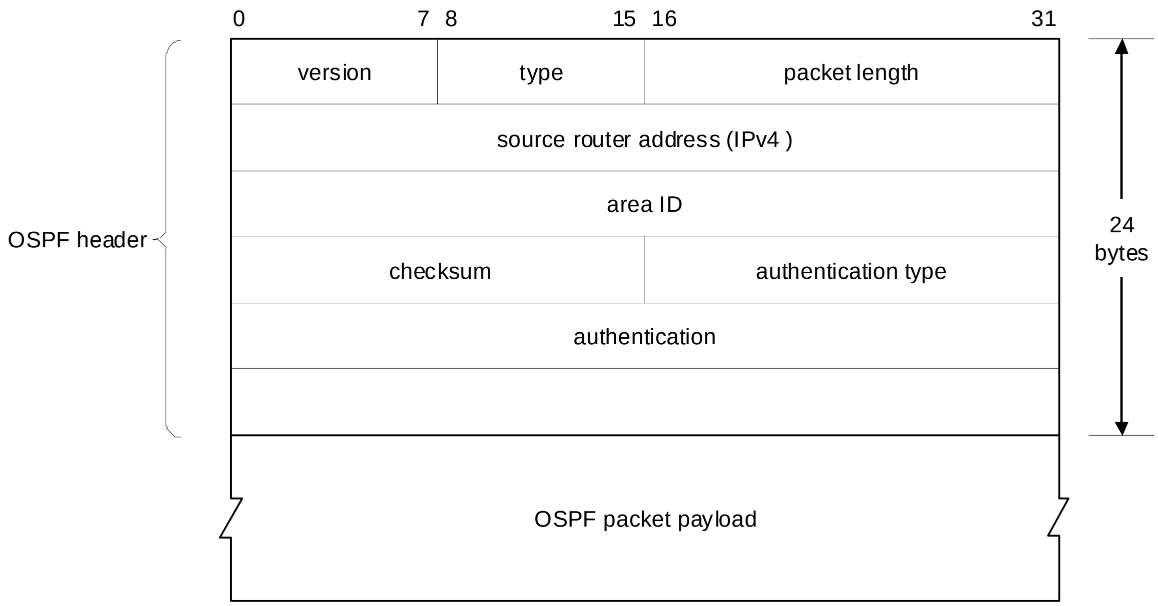

Figure 8-8: OSPF packet format for IPv4 addresses.

Link State Update packets are OSPF packet type 4. These packets implement the flooding of LSAs. Each Link State Update packet carries a collection of LSAs one hop further from their origin. Several LSAs may be included in a single packet. The payload format for type-4 packets is shown in Figure 8-9. There is one common LSA header for all LSA advertisement types, shown in the top part of Figure 8-9. The LSA advertisement type is specified in the *Type* field, see the top row in Figure 8-9.

As seen, these link-state advertisements (LSAa) are more complex than LSAs described in Section 1.4.2 for a basic version of link state routing. The complexity derives from the more complex link-state database graph representation for OSPF (Figure 8-7). For example, a router running OSPF may generate link-state advertisements that advertise one or more networks that are directly connected to this router. A router may also advertise a direct point-to- point link to another router.

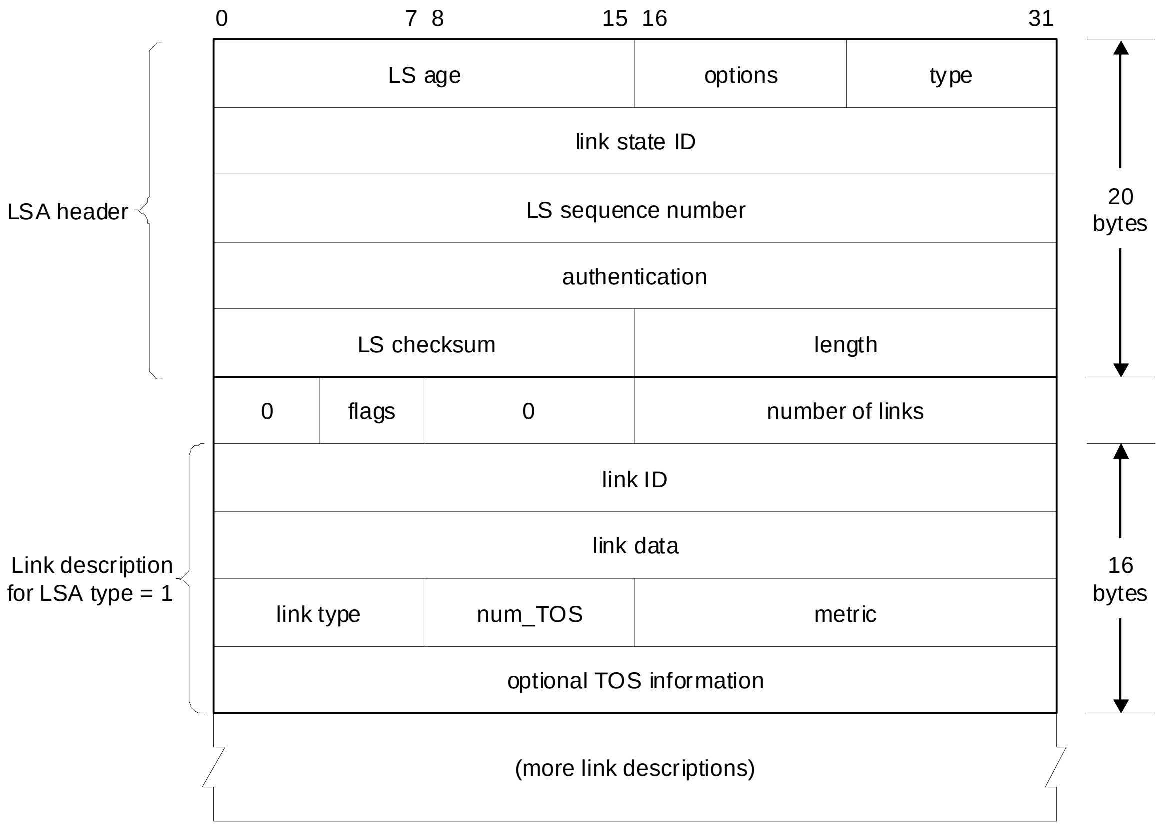

Figure 8-9: OSPF packet payload format for Type=4 packets (Link State Update packets). These packets carry one or more Link State Advertisements (LSAs). The bottom part shows the format of link description for Type 1 LSAs (specified in the LSA header in the top row).

OSPF has fast convergence rate. It can detect and propagate topology changes faster than a distance-vector routing protocol and it does not suffer from the counting-to-infinity problem (described in Section 1.4.3). OSPF-calculated routes are always loop-free. With OSPF, an autonomous system can be subdivided into contiguous groups of networks called areas. Routes within areas can be summarized to minimize route table entries. Areas can be configured with a default route summarizing all routes outside the AS or outside the area. As a result, OSPF can scale to large and very large internetworks.

8.2.3 Border Gateway Protocol (BGP)

Section 1.4.5 presented the key challenges that arise due to independent administrative entities that compete for profit to provide global Internet access. The key questions for an Autonomous System (AS) can be summarized as follows:

- What routing information to advertise to other ASs; how to process the routing information received from other ASs; and, what of the received information to readvertise?
- How can an AS achieve a consistent picture of the Internet viewed by all of its routers, so that for a given data packet each router would make the same forwarding decision (as if each had access to the routing tables of all the border routers within this AS)?

The inter-AS (or, external gateway routing protocol) routing protocol needs to decide whether to forward routing advertisement packets (import/export policies) and whether to disseminate reachability of neighboring ASs at the risk of having to carry their transit traffic unrecompensed.

Border Gateway Protocol (BGP) is an inter-Autonomous System routing protocol that addresses the above requirements. BGP is extremely complex and many issues about its operation are still not well understood. The main complexity of an external routing is not in the protocol for finding routes. Rather, the complexity lies in how border gateways (or, "BGP speakers") are configured to implement the business preferences, and in how external routes are learned from other ASs are disseminated internally within an AS. As will been seen later, there are two keys for this capability: (1) provider's *filtering policies* for processing and redistributing the received route advertisements, which are kept confidential from other ASes; and, (2) BGP *path attributes* that are included in route announcements and used when applying the local filtering policies.

BGP is a *path-vector routing protocol* (Section 1.4.5), where distance vectors are annotated not only with the entire path used to compute each distance, but also with path attributes that describe the advertised paths to destination prefixes. For example, the attributes include preference values assigned to an advertised path by the routers through which this advertisement passed. Unlike, distance-vector, path-vector routing converges quickly to correct paths and guarantees freedom from loops. However, there is an issue of large routing tables needed for path-vector routing. We will see later how BGP addresses this issue.

Routing Between and Within Autonomous Systems

In Section 1.4.5 we saw that the inter-AS routing protocol must exchange routing advertisements between different domains, as well as disseminate the received information within its own AS.

BGP routers use TCP (Chapter 2) on a well-known port (179) to communicate with each other, instead of layering the routing message directly over IP, as is done in other Internet routing protocols. Interior Gateway Protocols (IGPs), such as RIP (Section 8.2.1) and OSPF (Section 8.2.2) rely on periodic updates that carry the entire routing tables of the sending routers containing all active routes. Unlike this, BGP sends only **incremental updates** containing only the routing entries that have changed since the last update (or transmission of all active routes). TCP ensures reliable delivery and simplifies the error management in the routing protocol. However, routing updates are subject to TCP congestion control, which can lead to complicated network dynamics and performance problems. For example, routing updates might be delayed waiting for TCP sender to time out.

BGP neighbors, or peers, are established by manual configuration between routers to create a TCP session on port 179. Because TCP is a connection-oriented protocol with precisely identified endpoints, each BGP router maintains a separate TCP session with each other BGP router to which it is connected. There is typically one such BGP TCP connection for each link that directly connects two speaker routers in different ASs. There are also TCP connections between the speaker routers within the same AS (if there is more than one speaker within the AS), known as *internal peering*. Unlike BGP speakers in different ASs that are typically directly connected at the link layer, BGP speakers within the same AS are usually connected via non-speaker routers (i.e., at the network layer). For each TCP connection, the two routers at each end are called **BGP peers** and the TCP connection over which BGP messages are sent is called a **BGP session**. A BGP

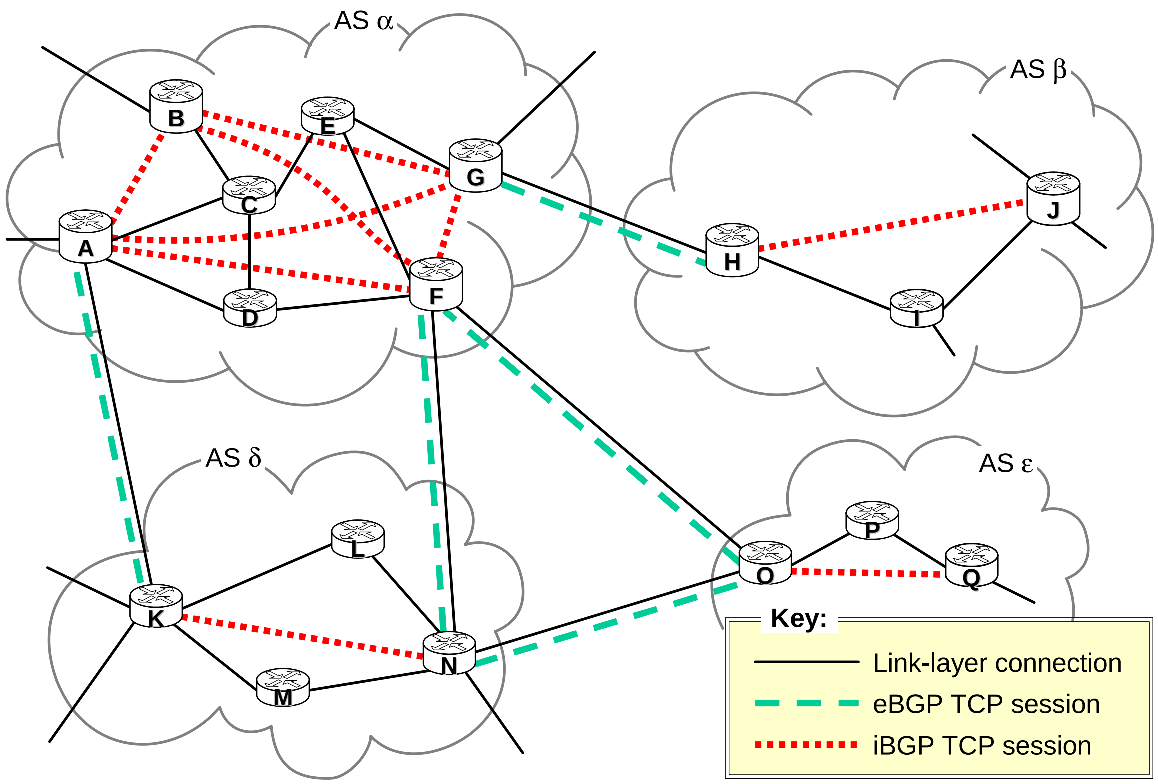

Figure 8-10: Example of eBGP and iBGP sessions.

session that spans two ASs is called an **external BGP (eBGP) session**; a BGP session between two routers within the same AS is called an **internal BGP (iBGP) session**. Recall from the discussion in Section 1.4.5 that the purpose of iBGP sessions is to ensure that network reachability information is consistent among the BGP speakers in the same AS. A BGP speaker can easily decide whether to open eBGP or iBGP session by comparing the Autonomous System Number (ASN) of the other router with its own. Figure 8-10 shows a detail from Figure 1-51 with eBGP and iBGP sessions.

As seen in Figure 8-10 for ASα, all iBGP peers must be fully connected to one another because each TCP session connects only one pair of endpoints. Full mesh connectivity (where everyone speaks to everyone directly) ensures that all the BGP routers in the same AS to exchange routing information and ensure that network reachability information is consistent among them. Given n BGP routers, the full mesh connectivity requires $n/2 \times (n - 1)$ iBGP sessions, which may be large for a large n. In addition, each BGP router will run at least one eBGP session (routers F, N, and O in Figure 8-10 run two each). Large number of sessions may degrade performance of routers, due either to high memory or processing requirements. Methods such as *confederations* (RFC-5065) and *route reflectors* (RFC-4456) help improve the scalability of iBGP sessions.

The BGP finite state machine (FSM) consists of six states (Figure 8-11): Idle, Connect, Active, OpenSent, OpenConfirm, and Established. To start participating in a *BGP session* with another router, a router first sets up a TCP connection on the BGP port 179 (states Connect and Active). If successful, the routers next exchange `OPEN` messages (states OpenSent and OpenConfirm). During the `OPEN` exchanges, BGP routers negotiate optional capabilities of the session, including multiprotocol extensions and various recovery modes.

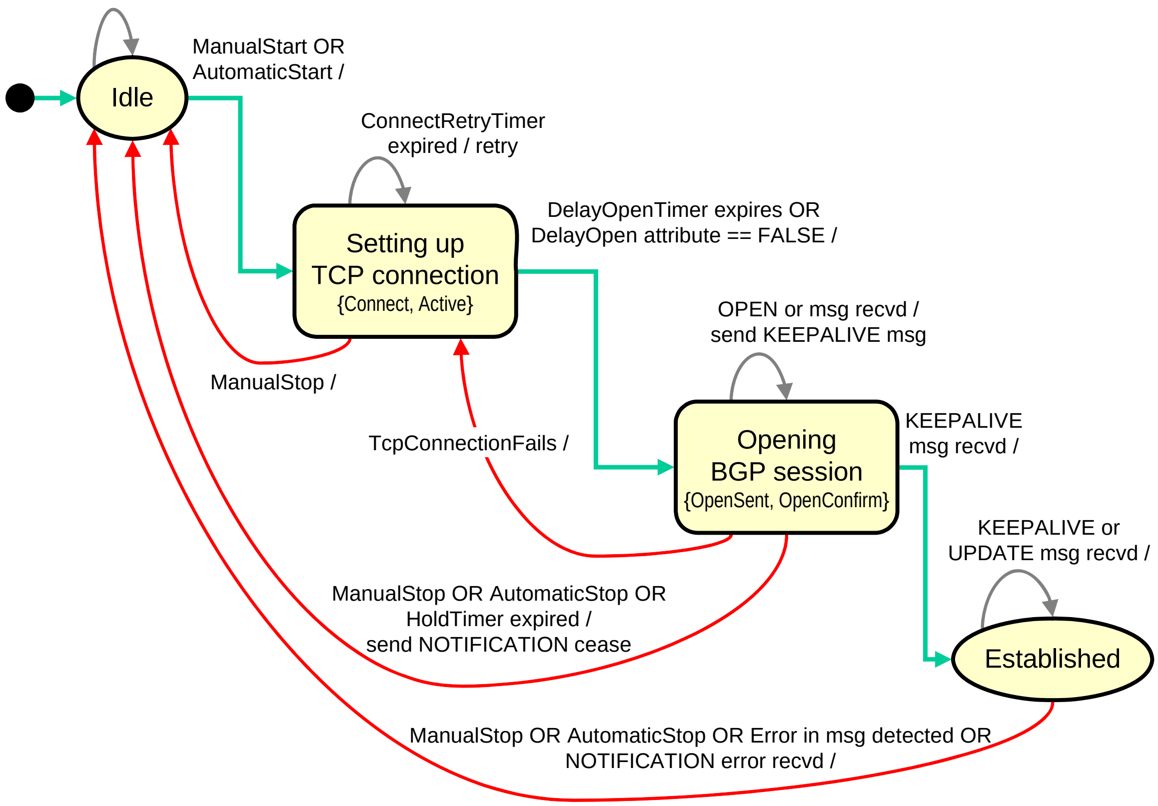

Figure 8-11: Finite state machine of the BGP4 protocol (highly simplified).

After the `OPEN` is completed and a BGP session is running, the BGP speakers transition to the Established state. They exchange `UPDATE` messages about destinations to which they offer connectivity (routing tables to all active routes). All subsequent `UPDATE` messages incrementally announce only routing entries that have changed since the last update. There are two kinds of updates: reachability *announcements*, which are changes to existing routes or new routes, and *withdrawals* of prefixes to which the speaker no longer offers connectivity (because of network failure or policy change). Both positive and negative reachability information can be carried in the same `UPDATE` message. In the protocol, the basic CIDR route description is called Network Layer Reachability Information (NLRI). NLRI includes the destination prefix, prefix length, path vector of the traversed autonomous systems and next hop in attributes, which can carry a wide range of additional information that affects the import policy of the receiving router.

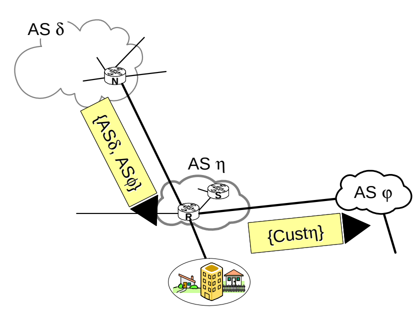

The exchanged routing tables are not necessarily the exact copies of their actual routing tables, because each router first applies the logic rules that implement its **export policy**. If a BGP speaker has a choice of several different routes to a destination, it will choose the best one according to its own local policies, and then that will be the route it advertises. Of course, a BGP speaker is not obliged to advertise any route to a destination, even if it knows one. For example, in Figure 1-51 ASη refuses to provide transit service to its peers and does not readvertise the destination in ASϕ that it learned about from ASδ.

Each router integrates the received information to its routing table according to its **import policy** (or, acceptance policy). The rules defining the import policy are specified using attributes such as LOCAL_PREF and WEIGHT. These attributes are locally exchanged by routers in an AS (using iBGP), but are not disclosed to other ASs (using eBGP). A BGP speaker calculates the degree of preference for each external route based on the locally-configured policy, and includes the degree of preference when advertising a route to its internal peers. A receiving BGP speaker uses the degree of preference learned via LOCAL_PREF in its decision process and favors the route with the highest degree of preference. The rules for BGP route selection are summarized at the end of this section in Table 8-3.

BGP Messages

All BGP messages begin with a fixed-size header, 19-bytes long, that identifies the message type (Figure 8-12(a)). A description of the header fields is as follows:

Marker: This 16-byte field is included for backwards compatibility; it *must* be set to all ones, unless when used for security purposes.

Length: This 2-byte unsigned integer indicates the total length of the message, including the header, in bytes (or, octets). The value of the Length field *must* always be between 19 and 4096, and may be further constrained, depending on the message type.

Type: This 1-byte unsigned integer indicates the type code of the message. BGP defines four type codes: {1 = OPEN; 2 = UPDATE; 3 = NOTIFICATION; 4 = KEEPALIVE}.

The BGP message types are discussed next.

- **OPEN Messages (Figure 8-12(b))**

After a TCP connection is established, the first message sent by each side is an OPEN message (Figure 8-12(b)). If the OPEN message is acceptable, a KEEPALIVE message confirming the OPEN is sent back. A description of the message fields is as follows:

Version: Indicates the BGP protocol version number of the message; currently it is 4.

My Autonomous System: Indicates the Autonomous System number of the router that sent this message.

Hold Time: Indicates the proposed interval between the successive KEEPALIVE messages (in seconds). The receiving router calculates the value of the Hold Timer by using the smaller of its configured Hold Time and the Hold Time received in this OPEN message. A Hold Time value of zero indicates that KEEPALIVE messages will not be exchanged at all; otherwise, the minimum value is three seconds.

BGP Identifier: Identifies the sending BGP router. The value of the BGP Identifier is determined upon startup and is the same for every local interface and BGP peer.

Optional Parameters Length: Indicates the total length of the Optional Parameters field in bytes. If the value of this field is zero, no Optional Parameters are present.

Optional Parameters: Contains a list of optional parameters, in which each parameter is encoded in TLV format <*Type*, *Length*, *Value*>. Parameter Type is a 1-byte field that unambiguously identifies individual parameters. Parameter Length is a 1-byte field that contains

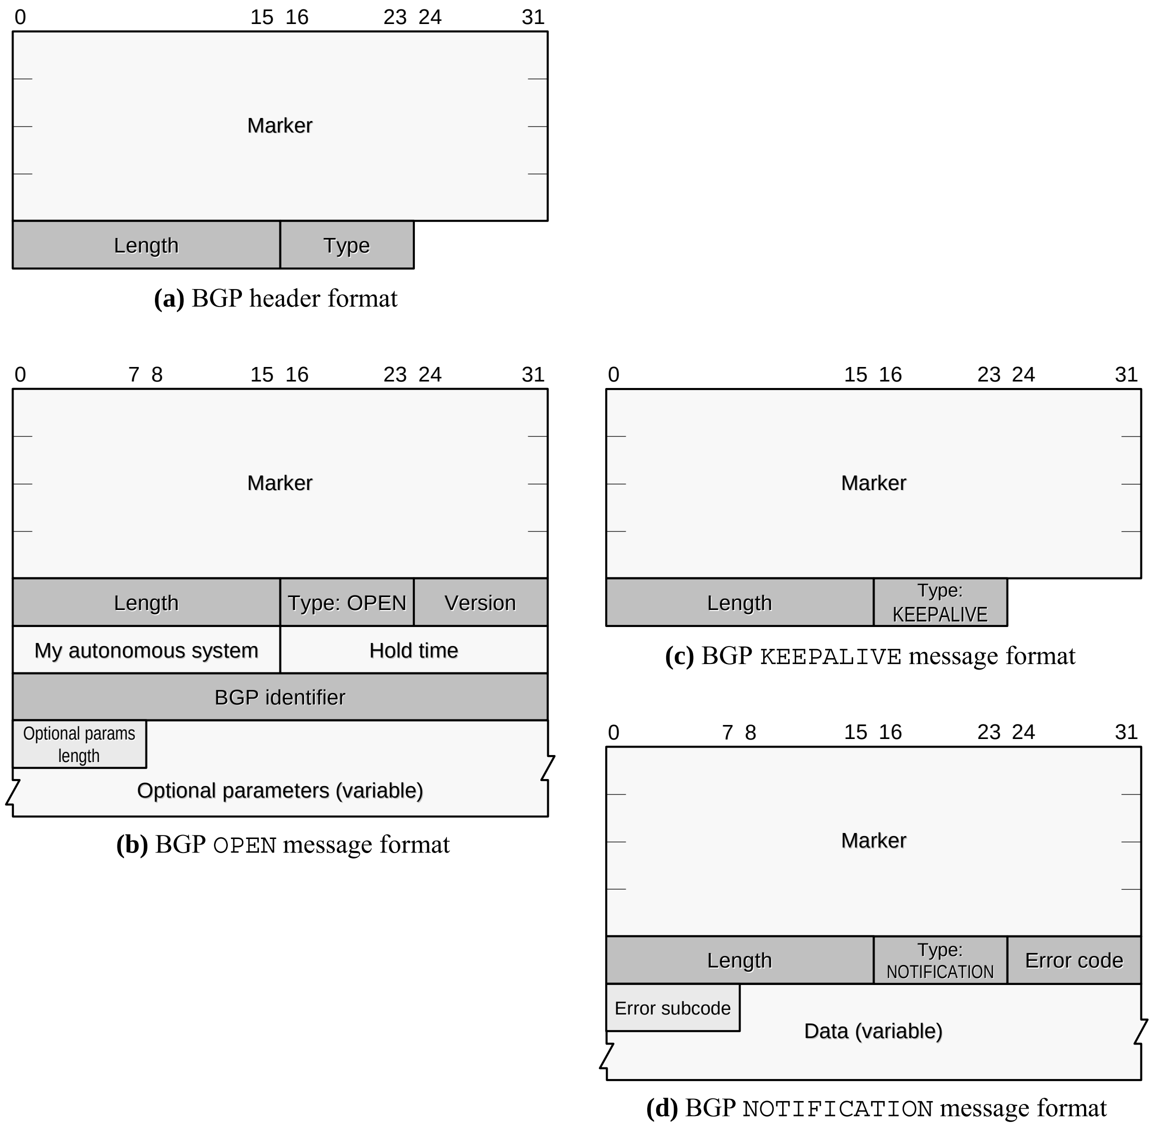

Figure 8-12: Format of BGP headers and messages, except for `UPDATE` (Figure 8-13).

the length of the Parameter Value field in bytes. Parameter Value is a variable length field that is interpreted according to the value of the Parameter Type field.

The minimum length of an `OPEN` message is 29 bytes (including the message header).

- **`KEEPALIVE` Messages (Figure 8-12(c))**

BGP does not use any TCP-based, keep-alive mechanism to determine if peers are reachable. Instead, `KEEPALIVE` messages are exchanged between peers at a rate that prevents the Hold Timer from expiring. A recommended time between successive `KEEPALIVE` messages is one-third of the Hold Time interval. `KEEPALIVE` messages must not be sent more frequently than one per second. If the negotiated Hold Time interval is zero, then periodic `KEEPALIVE` messages will not be sent.

A `KEEPALIVE` message (Figure 8-12(c)) consists of only the message header and has a total length of 19 bytes.

Table 8-2: BGP NOTIFICATION message error codes and subcodes.

Code	Description	Subcodes (if present)	
1	Message header error	1 – Connection not synchronized	3 – Bad message type
		2 – Bad message length	
2	OPEN message error	1 – Unsupported version number	4 – Unsupported optional parameter
		2 – Bad peer AS	5 – Deprecated
		3 – Bad BGP identifier	6 – Unacceptable Hold Time
3	UPDATE message error	1 – Malformed attribute list	7 – Deprecated
		2 – Unrecognized well-known attribute	
		3 – Missing well-known attribute	8 – Invalid NEXT_HOP attribute
		4 – Attribute flags error	9 – Optional attribute error
		5 – Attribute length error	10 – Invalid Network field
		6 – Invalid ORIGIN attribute	11 – Malformed AS_PATH
4	Hold timer expired		
5	Finite state machine error		
6	Cease		

- **NOTIFICATION Messages (Figure 8-12(d))**

A NOTIFICATION message is sent when an error condition is detected. The BGP connection is closed immediately after it is sent. In addition to the fixed-size BGP header, the NOTIFICATION message contains the following fields (Figure 8-12(d)): Error Code, Error Subcode, and Data of variable length. The Error Code indicates the type of error condition, while the Error Subcode provides more specific information about the nature of the reported error (Table 8-2). Each Error Code may have one or more Error Subcodes associated with it. If no appropriate Error Subcode is defined, then a zero (Unspecific) value is used for the Error Subcode field.

The variable-length Data field is used to diagnose the reason for the NOTIFICATION. The minimum length of the NOTIFICATION message is 21 bytes (including message header).

- **UPDATE Messages (Figure 8-13)**

After the connection is established, BGP peers exchange routing information by using the UPDATE messages. The information in the UPDATE messages is used by the path-vector routing algorithm (Section 1.4.5) to construct a graph that describes the connectivity of the Autonomous Systems. By applying logical rules, routing information loops and some other anomalies may be detected and removed from inter-AS routing.

An UPDATE message is used to advertise feasible routes that share common path attributes to a peer, and to withdraw multiple unfeasible routes from service. The UPDATE message always includes the fixed-size BGP header, and other fields, some of which may not be present in every UPDATE message (Figure 8-13(a)).

Withdrawn Routes Length indicates the total length of the Withdrawn Routes field in bytes. A value of zero indicates that no routes are being withdrawn from service, and that the Withdrawn Routes field is not present in this UPDATE message.

The Withdrawn Routes field contains a variable-length list of IP-address prefixes for the routes that are being withdrawn from BGP routing tables. Each prefix is encoded as a 2-tuple of the form <*length*, *prefix*>. The *Length* field indicates the length (in bits) of the prefix. A length of

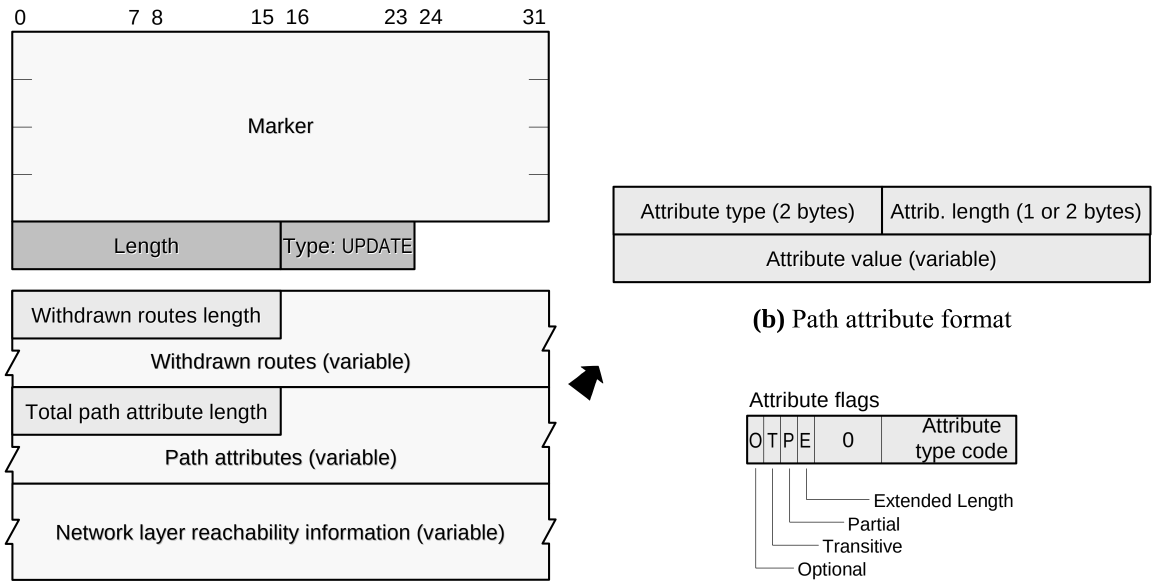

(a) BGP `UPDATE` message format

(b) Path attribute format

(c) Attribute type format

Figure 8-13: Format of BGP `UPDATE` message.

zero indicates a prefix that matches all IP addresses. The *Prefix* field contains an IP-address prefix, possibly followed by padding bits to make the field length a multiple of 8 bits.

Total Path Attribute Length indicates the total length of the Path Attributes field in bytes. A value of zero indicates that neither the Network Layer Reachability Information field nor the Path Attribute field is present in this `UPDATE` message.

A BGP router uses the *Path Attributes* and *Network Layer Reachability Information* (NLRI) fields to advertise a route. The NLRI field contains a list of IP-address prefixes that can be reached by this route. The NLRI is encoded as one or more 2-tuples of the form <*length*, *prefix*>. This is the path-vector information used by the path-vector routing algorithm (Section 1.4.5).

A variable-length sequence of Path Attributes is present in every `UPDATE` message, except for an `UPDATE` message that carries only the withdrawn routes. Each path attribute is a triple <*type*, *length*, *value*> of variable length (Figure 8-13(b)).

Attribute Type field that consists of the Attribute Flags byte, followed by the Attribute Type Code byte (Figure 8-13(c)). The high-order bit (bit 0) of Attribute Flags is the Optional bit. It defines whether the attribute is optional (if set to 1) or well-known (if set to 0). The second bit is the Transitive bit. It defines whether an optional attribute is transitive (if set to 1) or non-transitive (if set to 0). For well-known attributes, the Transitive bit must be set to 1. The third bit is the Partial bit that defines whether the information contained in the optional transitive attribute is partial (if set to 1) or complete (if set to 0). For well-known attributes and for optional non-transitive attributes, the Partial bit must be set to 0. The fourth bit of Attribute Flags is the Extended Length bit. It defines whether the following Attribute Length field is one byte (if set to 0) or two bytes (if set to 1). The lower-order four bits of Attribute Flags are unused. They are set to zero by the sender, and ignored by the receiver.

The Attribute Type Code field contains the Attribute Type Code. Attribute Type Codes defined in RFC-4271 are discussed next.

BGP Path Attributes

This section discusses the path attributes of the `UPDATE` message (Figure 8-13). A BGP route announcement (`UPDATE` message) has a set of attributes associated with each destination prefix. Path attributes can be classified as: (1) well-known mandatory; (2) well-known discretionary; (3) optional transitive; and, (4) optional non-transitive. A BGP router must recognize all well-known attributes. Some of these attributes are mandatory and must be included in every `UPDATE` message that contains Network Layer Reachability Information (NLRI). Others are discretionary and may or may not be sent in a particular `UPDATE` message. Once a BGP peer has updated any well-known attributes, it must pass these attributes to its peers in any updates it transmits.

In addition to well-known attributes, each path may contain one or more optional attributes. It is not required or expected that all BGP implementations support all optional attributes. The handling of an unrecognized optional attribute is determined by the value of the Transitive flag (Figure 8-13(c)).

• **`ORIGIN` (type code 1)** is a well-known mandatory attribute. The `ORIGIN` attribute describes how BGP at the origin AS came to know about destination addresses aggregated in a prefix. The allowed values are: code 1 (IGP) means that the prefix was learned from an interior gateway protocol (IGP); code 2 (EGP) means that the prefix was learned from an exterior gateway protocol; and, code 3 (INCOMPLETE) represents another source, usually a manually configured static route. The value of `ORIGIN` should not be changed by any other speaker.

• **`AS_PATH` (type code 2)** is a well-known mandatory attribute. This attribute lists the autonomous systems through which this `UPDATE` message has traversed (in reverse order). This list is called a **path vector** and hence BGP is a *path vector protocol* (Section 1.4.5). Every router through the message passes prepends its own AS number (ASN) to `AS_PATH` and propagates the `UPDATE` message on (subject to its route filtering rules). An example is shown in Figure 1-51 and a detail in Figure 8-14. When a given BGP speaker advertises the route to an internal peer (over iBGP), the advertising speaker does not modify the `AS_PATH` attribute.

BGP uses the `AS_PATH` attribute to detect a potential routing loop. When an external `UPDATE` message is received (over eBGP), if the ASN of this BGP speaker is already contained in the path vector, then the route should be rejected.

The speaker that modifies `AS_PATH` may prepend more than one instance of its own ASN in the `AS_PATH` attribute. This is controlled via local configuration.

• **`NEXT_HOP` (type code 3)** is a well-known mandatory attribute. It defines the IP address of the next-hop speaker to the destinations listed in the `UPDATE` message (in NLRI). As the `UPDATE` message propagates across an AS boundary, the `NEXT_HOP` attribute is changed to the IP address of the speaker from which this announcement was received (Figure 8-14). (The reader should check RFC-4271 about detailed rules for modifying the `NEXT_HOP` attribute.)

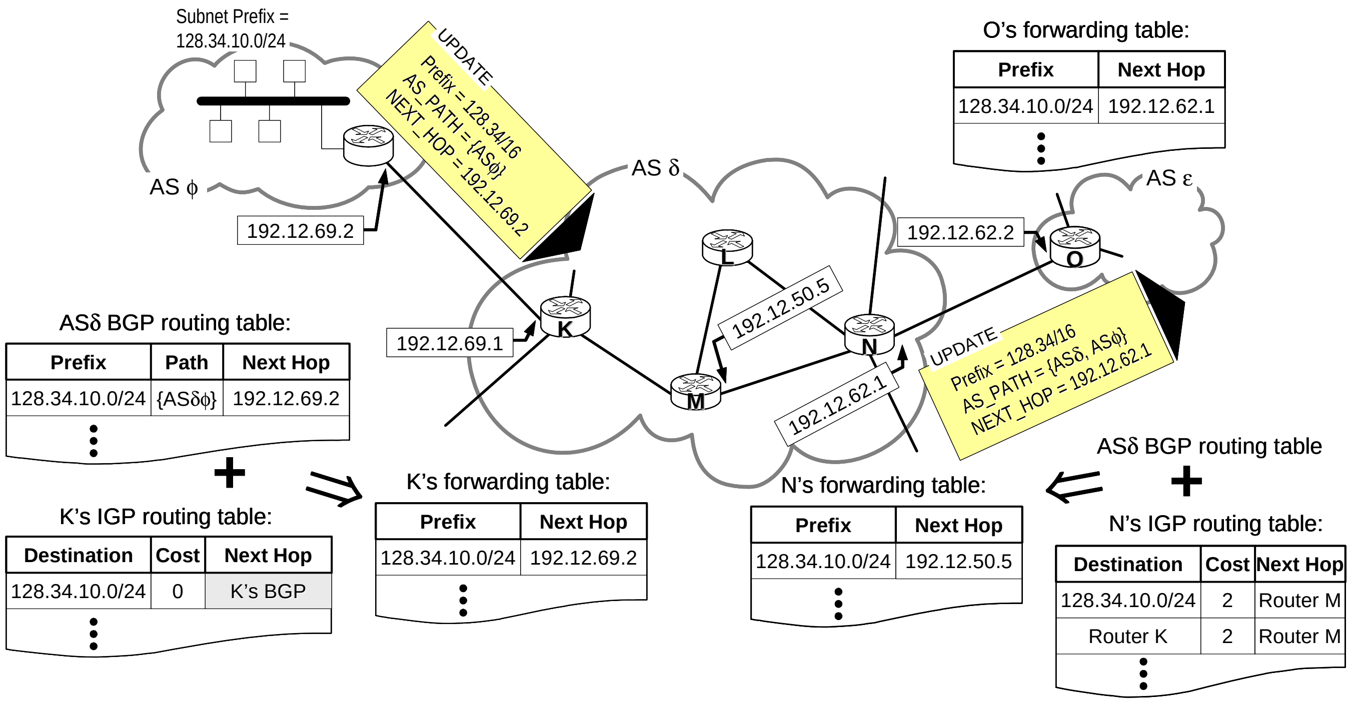

Figure 8-14: Example of BGP `UPDATE` message propagation, originating from ASϕ.

When sending `UPDATE` to an internal peer, if the route is originated externally to this AS, the BGP speaker should not modify the `NEXT_HOP` attribute unless it has been explicitly configured to announce its own IP address as the `NEXT_HOP`. This is because `UPDATE` messages to internal peers are sent by iBGP over the TCP connection, which runs on top of an IGP. On the other hand, when announcing a route to a local destination within this AS, the BGP speaker should use as the `NEXT_HOP` the IP address of the first router that announced the route.

As Figure 8-14 shows, all BGP speakers in an AS have the same BGP routing table (ASδ BGP routing table is shown). The forwarding table is created based on the AS BGP routing table and the IGP routing table (which is, naturally, different for each router).

The immediate next-hop address is determined by performing a recursive route lookup operation for the IP address in the `NEXT_HOP` attribute, using the contents of the Routing Table, selecting one entry if multiple entries of equal cost exist. The Routing Table entry that resolves the IP address in the `NEXT_HOP` attribute will always specify the outbound interface. If the entry specifies an attached subnet, but does not specify a next-hop address, then the address in the `NEXT_HOP` attribute should be used as the immediate next-hop address. If the entry also specifies the next-hop address, this address should be used as the immediate next-hop address for packet forwarding.

• **`MULTI_EXIT_DISC` (type code 4)** is an optional non-transitive attribute that is intended to be used on external (inter-AS) links to discriminate among multiple exit or entry points to the same neighboring AS. To motivate the need for this attribute, consider the example in Figure 8-15 (extracted from Figure 1-51), where Autonomous Systems ASα and ASδ are linked at multiple points. Suppose that router *A* in ASα receives a data packet from ASχ that is destined for ASη. ASα would prefer to get rid of the packet in a hurry ("hot-potato" routing) and simply forward it

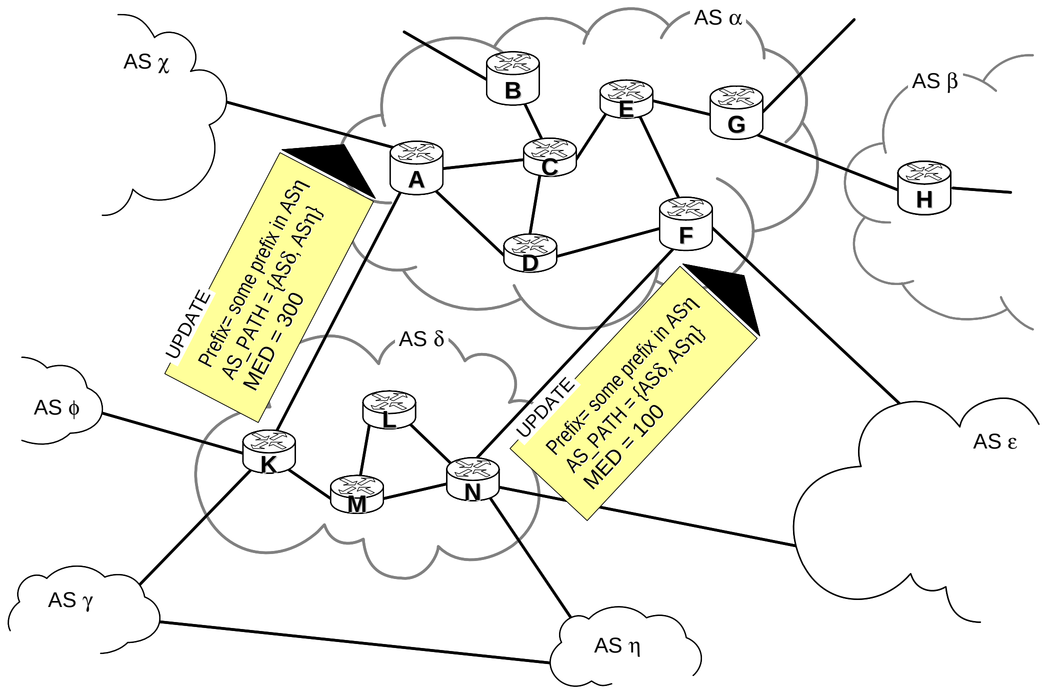

Figure 8-15: Example for BGP `MULTI_EXIT_DISC` (`MED`) attribute.

to router *K* in ASδ. However, ASδ would prefer to receive packets for ASη on the other point (router *N*), because this is less expensive for ASδ—the packet would be immediately forwarded by *N* to ASη instead of traversing ASδ's interior routers. If there were no financial settlement involved, ASα would simply implement its own preferences. However, because ASδ pays ASα for transit service (Figure 1-49), ASα has to honor ASδ's preferences. Earlier we described how a BGP router uses the attribute `LOCAL_PREF` to integrate the received routing advertisement into its routing table. However, because `LOCAL_PREF` expresses the local preferences, it is not useful to express another AS's preferences. For this purposed the `MULTI_EXIT_DISCRIMINATOR` (`MED`) attribute is used.

The value of the `MULTI_EXIT_DISC` (`MED`) attribute is a four-byte integer, called a *metric*, and the standard does not prescribe how to choose the `MED` metric. A common practice is to derive the `MED` metric from the local IGP metric. All other factors being equal, the exit point with the lower metric should be preferred. In Figure 8-15, router *N* advertises a prefix in ASη with `MED` = 100, but router *K* advertises the same prefix with `MED` = 300. Based on this, ASα should deliver packets destined for ASη to router *N*. One the other hand, ASδ would prefer to receive destined packets for ASγ or ASϕ on router *K*, so it will choose opposite values of `MED` attribute when advertising prefixes in ASγ or ASϕ.

In peering relationships between ASs, (Section 1.4.5), the `MED` attribute is usually ignored. The usage of the `MED` attribute becomes complicated when a third AS advertises the same route, because the IGP metrics used by different ASs can be different. In such a case, comparing a `MED` metric received from one AS with another `MED` metric received from another AS makes no sense.

There are three more path attributes (`LOCAL_PREF`, `ATOMIC_AGGREGATE`, and `AGGREGATOR`), and the interested reader should check RFC-4271. Table 8-3 summarizes how a BGP router that learned about more than route to a prefix selects one. The router will select the

Table 8-3: Priority of rules by which BGP speaker selects routes from multiple choices.

Priority	Rule	Comments
1	`LOCAL_PREF`	E.g., `LOCAL_PREF` specifies the order of preference as customer > peer > provider If more than one route remains after this step, go to the next step.
2	`AS_PATH`	Select shortest `AS_PATH` length (i.e., the list with the smallest number of ASNs, *not* smallest number of hops or lowest delay!)
3	`MED`	Select the route with the lowest `MULTI_EXIT_DISC` value, if there is financial incentive involved.
4	IGP path	Select the route for which the `NEXT_HOP` attribute, for which the cost in the IGP routing table is lowest, i.e., use hot-potato routing.
5	eBGP > iBGP	Select the route which is learned from eBGP over the one learned by iBGP (i.e., prefer the route learned first hand)
6	Router ID	Select the BGP router with the smallest IP address as the next hop.

next hop along the first route that meets the criteria of the logical rule, starting with the highest priority (1) and going down to the lowest priority (6).

8.2.4 Multicast Routing Protocols

Multicast Group Management

Internet Group Management Protocol (IGMP) is used by an end-system to declare membership in particular multicast group to the nearest router(s). IGMP v3 (current) is defined by RFC-3376.

Version 1: Timed-out Leave (Joining Host send IGMP Report; Leaving Host does nothing; Router periodically polls hosts on subnet using IGMP Query; Hosts respond to Query in a randomized fashion)

Version 2: Fast, Explicit Leave (ADDS to Version 1: Group Specific Queries; Leave Group Message; Host sends Leave Group message if it was the one to respond to most recent query; Router receiving Leave Group message queries group.)

Version 3: Per-Source Join (ADDS to Version 2: Group-Source Specific Queries, Reports and Leaves; Inclusion/Exclusion of sources)

Multicast Route Establishment

Protocol Independent Multicast (PIM), RFC-4601 defines Protocol Independent Multicast – Sparse Mode (PIM-SM); RFC-3973 defines Protocol Independent Multicast – Dense Mode (PIM-DM). PIM operates independently of the underlying unicast protocol, such as IS-IS or OSPF. It supports applications that operate with fewer servers transmitting to multiple destinations (called the *dense mode*) or numerous small workgroups operating in different multicast groups (called the *sparse mode*).

Distance Vector Multicast Routing Protocol (DVMRP), defined in RFC-1075. DVMRP is an enhancement of Reverse Path Forwarding (RPF, Section 3.3.2) that: Uses Distance Vector routing packets for building tree; Prunes broadcast tree links that are not used (non-membership reports); Allows for Broadcast links (LANs).

Multicast Forwarding in DVMRP: 1. check incoming interface: discard if not on shortest path to source; 2. forward to all outgoing interfaces; 3. do not forward if interface has been pruned; 4. prunes timeout every minute.

Source-Specific Multicast (SSM), defined in RFC-3569 and RFC-4607.

Multicast Open Shortest Path First (MOSPF); RFC-1584 defines multicast extensions to OSPF.

For independent Autonomous Systems, Border Gateway Multicast Protocol (BGMP) was abandoned due to the lack of support within the Internet service provider community. Multiprotocol Extensions for BGP4, defined in RFC-2858 defines the codes for Network Layer Reachability Information (NLRI) that allow BGP to carry multicast information. This information is used by other (i.e., multicast) protocols for multicast forwarding. Multicast Source Discovery Protocol (MSDP) defined in RFC-3618, can be used to connect together rendezvous points in different PIM sparse mode domains (see RFC-4611 for MSDP deployment scenarios).

Additional information: see RFC-3170 – "IP Multicast Applications: Challenges & Solutions"

An excellent overview of the current state of multicast routing in the Internet is RFC-5110 (January 2008).

8.3 Address Translation Protocols

8.3.1 Address Resolution Protocol (ARP)

In Chapter 1 we saw that different protocol layers use different addressing systems. The network-layer Internet Protocol uses IPv4 or IPv6 addresses that are assigned by an Internet authority. Link-layer protocols (Ethernet and Wi-Fi) use MAC addresses that are assigned by the hardware manufacturer. One may wonder, why cannot we use a single addressing system, e.g., MAC addresses that are assigned to all network interface cards? Computer hardware and software are abstract so everything may seem possible to make. To understand better why we need two (or more) addressing systems, let us look at real-world physical objects, such as vehicles. (See also Sidebar 1.2 in Section 1.4.4.) As shown in Figure 8-16, every vehicle comes with a *vehicle identification number* (VIN) that is assigned by the manufacturer and engraved at several locations in the vehicle. Every vehicle also has a registration plate with a unique number. Both numbers can be considered "addresses" of the car. So, why not to use the VIN number for the registration plates, as well? Because VINs are assigned by the manufacturer and vehicles of the same manufacturer are bought by customers all around the world, it is impossible to embed into

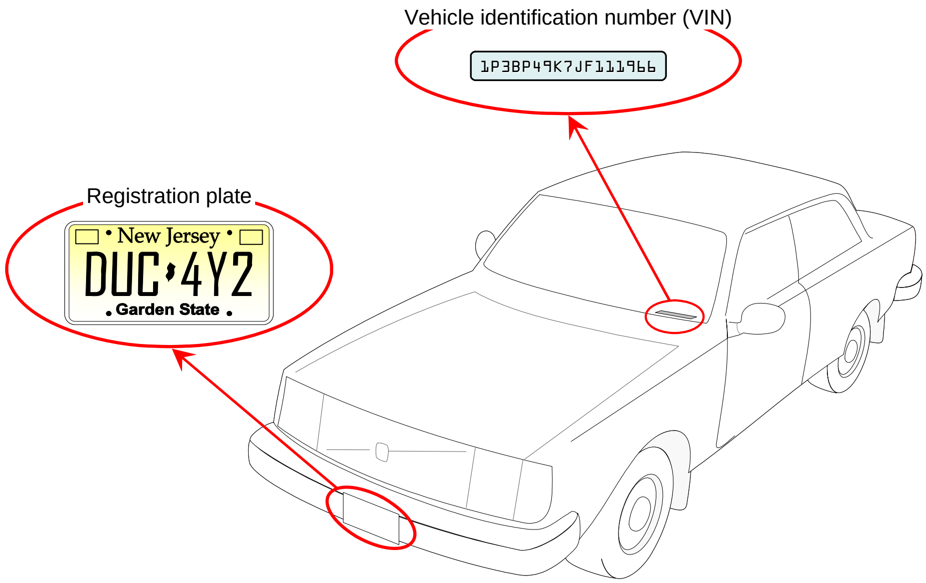

Figure 8-16: Why multiple addressing conventions are necessary.

the address anything specific to a geographic locale or organization that owns cars. Having registration plate number assigned by a local authority makes possible to have for location-specific addresses. Therefore, we need both: the manufacturer needs to be able to distinguish their different products and organizations need location-specific addresses.

A key benefit of location-specific network addresses is the ability to aggregate many addresses by address prefixes, which reduces the amount of information carried in routing advertisements or stored in the forwarding tables of routers.

Let us assume that a node has a packet for a certain destination. The network-layer protocol (IP) at the node looks up the forwarding table and determines the IP address of the next-hop node. Before calling the `send()` method of the link-layer protocol (see Listing 1-1 in Section 1.1.4), the network-layer `send()` must translate the next-hop IP address to the link-layer address of the next-hop node. Recall that Point-to-Point Protocol (PPP, Section 1.5.1) does not use link-layer addresses because it operates over a point-to-point link directly connecting two nodes, one on each end of the link. However, broadcast-based link-layer protocols, such as Ethernet (Section 1.5.2) or Wi-Fi (Section 1.5.3) must use link-layer addresses because many nodes are simultaneously listening on the channel. Because broadcast-based link-layer protocols implement Medium Access Control (MAC), these addresses are called **MAC addresses**. To send the packet to the next-hop node, the node needs a mechanism to translate from a (network-layer) IP address to a (link-layer) MAC address. This is the task of the Address Resolution Protocol.

Address Resolution Protocol (ARP) translates an IP address to a MAC address for a node that is on the same broadcast local-area network (or, subnet). ARP cannot translate addresses for hosts that are not on the same subnet; if such attempt is made, ARP returns an error. When a sender wants a translation, it looks up an **ARP table** on its node, which contains mappings of network-layer IP addresses to MAC addresses.

If the ARP table does not contain an entry for the given IP address, it broadcasts a query **ARP packet** on the LAN (Figure 8-17). (The MAC broadcast address in hexadecimal notation is

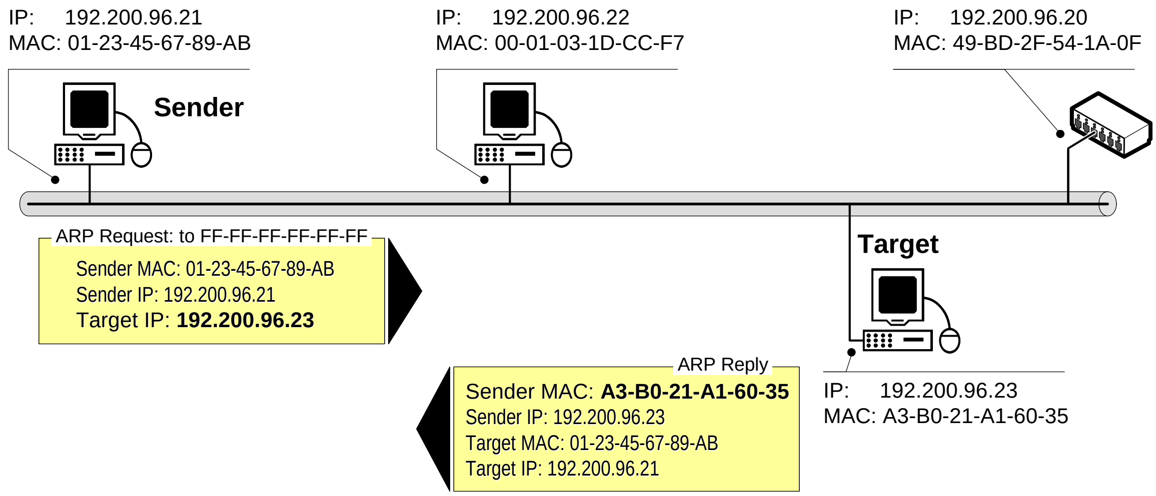

Figure 8-17: ARP request and response example.

FF-FF-FF-FF-FF-FF.) The node that owns the IP address replies with an ARP packet containing the responder's MAC address. The reply is sent to the querier's MAC address, available from the ARP request.

Figure 8-18 shows the ARP packet format for mapping IPv4 addresses to MAC addresses. In this case, the ARP packet size is 28 bytes. ARP can be used for other kinds of mappings, with different address sizes, as standardized by the IANA (http://iana.org/). The packet fields are:

• Hardware Type: specifies the link-layer protocol type. For example, the code for Ethernet is 1.

• Protocol Type: specifies the upper layer protocol for which the ARP request is intended. For example, IPv4 is encoded as 0x0800.

• Hardware Length: length (in bytes) of a hardware address. Ethernet addresses size is 6 bytes (48 bits).

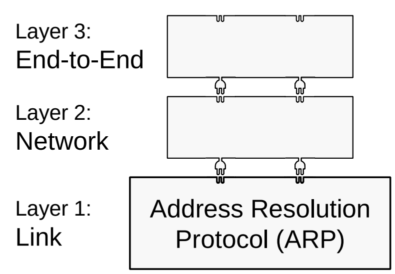

• Protocol Length: length (in bytes) of a logical address of the network-layer protocol. IPv4 address size is 4 bytes (32 bits).

• Operation: specifies the operation what the sender is performing: 1 for request, 2 for reply.

• Sender Hardware Address: hardware (MAC) address of the sender.

• Sender Protocol Address: upper-layer protocol address of the sender, e.g. IP.

• Target Hardware Address: hardware (MAC) address of the intended receiver. This field is ignored in request operations.

• Target Protocol Address: upper layer protocol address of the intended receiver.

The graphic on the right shows ARP as a link-layer protocol. This is somewhat controversial because ARP uses another link-layer protocol for sending ARP packets, and ARP deals with network-layer, IP addresses. The reason to classify ARP as a link-layer protocol is that it operates over a single link connecting nodes on the same local-area network. Unlike network-layer protocols, it does not span multiple hops and does not send packets across intermediate nodes.

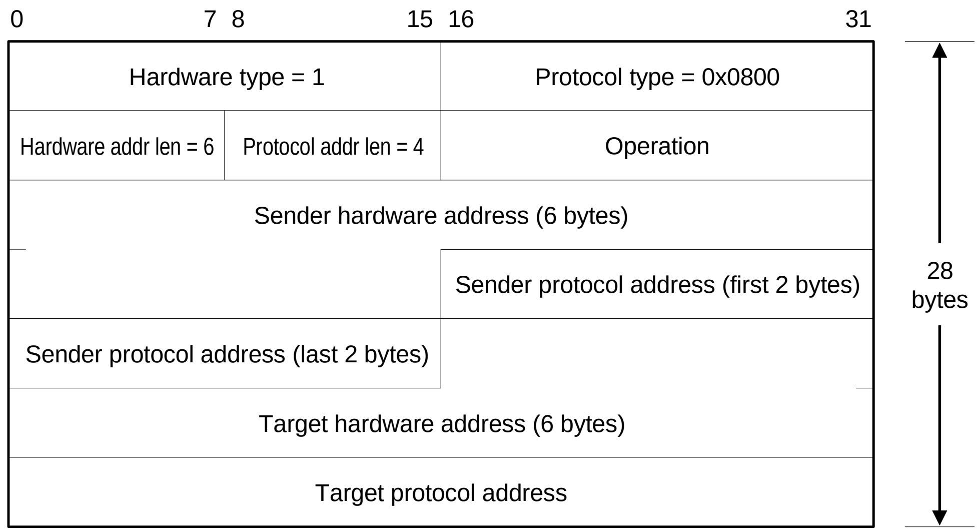

Figure 8-18: ARP packet format for mapping IPv4 addresses into MAC addresses.

ARP solves the problem of determining which MAC address corresponds to a given IP address. Sometimes the reverse problem has to be solved: Given a MAC address, determine the corresponding IP address. An early solution was to use *Reverse ARP* (*RARP*) for this task. RARP is now obsolete and it is replaced by Dynamic Host Configuration Protocol (DHCP), reviewed in the next section.

In the next generation Internet Protocol, IPv6, ARP's functionality is provided by the Neighbor Discovery Protocol (NDP).

8.3.2 Dynamic Host Configuration Protocol (DHCP)

To send a packet on the Internet, a computer must have a network address (IP address). This address is associated with the location of the computer, specifically with the network to which the computer is attached (Section 1.4.4). Dynamic Host Configuration Protocol (DHCP) supports automatic assignment of IP addresses to new hosts, known as plug-and-play. When a new computer is attached to a local-area network, the computer broadcasts a DHCP message asking "Router, give me a network address." The router maintains a pool of free network addresses and assigns one to this computer with a specified time-to-live (say one hour). The computer can then start using the standard Internet applications. As the time-to-live becomes close to zero, the computer asks the router for an extension, which is normally granted. If the user unplugs the computer, there will be no message asking for extension, so the router will return this network address to the pool of free addresses.

8.3.3 Network Address Translation (NAT)

Network Address Translation (NAT) is an Internet Engineering Task Force (IETF) standard used to allow multiple computers on a private network (using private address ranges such as 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16) to share a single, globally routable IPv4 address.

NATs are often deployed because public IPv4 addresses are becoming scarce. Companies often use NAT devices to share a single public IPv4 address among dozens or hundreds of systems that use private, often duplicated IPv4 addresses. These private IPv4 addresses cause problems if inadvertently leaked across the public Internet by private IP-based networks.

The NAT router translates traffic coming into and leaving the private network. NAT allows a single device, such as a router, to act as agent between the Internet (or, “public network”) and a local (or, “private”) network. This means that only a single unique IP address is required to represent an entire group of computers to anything outside their network.

NAT is an immediate but temporary solution to the IPv4 address exhaustion problem that will eventually be rendered unnecessary with IPv6 deployment. However, the shortage of IP addresses in IPv4 is only one reason to use NAT. Two other reasons are:

· Security

· Administration

Implementing dynamic NAT automatically creates a firewall between your internal network and outside networks or the Internet. *Dynamic NAT* allows only connections that originate inside the stub domain. Essentially, this means that a computer on an external network cannot connect to your computer unless your computer has initiated the contact. NAT provides a simple packet filtering function by forwarding only solicited traffic to private network hosts. Solicited traffic is traffic that was requested by a private network host. For example, when a private host computer accesses a Web page, the private host computer requests the page contents from the Web server. The traffic for the Web page contests is solicited traffic. By default, a NAT does not forward unsolicited traffic to private network hosts. Therefore, you can browse the Internet and connect to a site, even download a file. However, somebody else cannot simply latch onto your IP address and use it to connect to a port on your computer.

Static NAT, also called *inbound mapping*, allows connections initiated by external devices to computers on the stub domain to take place in specific circumstances. For instance, you may wish to map an inside global address to a specific inside local address that is assigned to your Web server. Static NAT (inbound mapping) allows a computer on the stub domain to maintain a specific address when communicating with devices outside the network.

8.3.4 Mobile IP

Mobility is the quality of being capable of movement or moving readily from place to place. Wireless devices provide this kind of untethered freedom, but mobility means more than the lack of a network cable. Many terms describe mobility, but this chapter uses the terms mobility and roaming to describe the act of moving between access points.

Defining or characterizing the behavior of roaming stations involves two forms:

* Seamless roaming

* Nomadic roaming

Seamless roaming is best analogized to a cellular phone call. For example, suppose you are using your cellular phone as you drive your car on the freeway. A typical global system for mobile

(GSM) communications or time-division multiple access (TDMA) cell provides a few miles of coverage area, so it is safe to assume that you are roaming between cellular base stations as you drive. Yet as you roam, you do not hear any degradation to the voice call (that is what the cellular providers keep telling us). There is no noticeable period of network unavailability because of roaming. This type of roaming is deemed seamless because the network application requires constant network connectivity during the roaming process.

Nomadic roaming is different from seamless roaming. Nomadic roaming is best described as the use of an 802.11-enabled laptop in an office environment. As an example, suppose a user of this laptop has network connectivity while seated at his desk and maintains connectivity to a single AP. When the user decides to roam, he undocks his laptop and walks over to a conference room. Once in the conference room, he resumes his work. In the background, the 802.11 client has roamed from the AP near the user's desk to an AP near the conference room. This type of roaming is deemed nomadic because the user is not using network services when he roams, but only when he reach his destination.

The Nomadic Mobility Problem

Internet hosts are widely known by their IP addresses. We also know that routers use the CIDR scheme to aggregate sets of contiguous IP addresses and, therefore, simplify the task of routing messages (Section 1.4.4). Imagine you are traveling and you need a means to allow your friends to send you messages while you are away. If this is a short travel, your visiting address(es) will not enter the "infrastructure" records, such as those of government agencies or public registries. You could designate a care-of agent to whom the post office will deliver your mail, which you will collect when returning back. However, if you want your mail forwarded to your visiting location, you need to let the postal office know your visiting address, which may be difficult if you do not know where you will be staying or will be staying at different locations only briefly and unpredictably. One option is that you explicitly notify your care-of agent every time you arrive at a visiting address. Notice that there is a built-in inefficiency when relying on a care-of agent instead of the infrastructure registries, because every communication must first travel to your home address (known to the infrastructure registries), and then be redirected by your care-of agent to your visiting address. In the worst case, you may be located near the sender and far away from home, so the mail needs to make an unnecessary trip to your home and back. We will see that the Internet functions in a similar manner.

The Mobile IP Protocol

Mobile IP is a network-layer protocol (or Layer-3 protocol in the OSI architecture). We have already seen in Section 1.5.3 how extended service set (ESS) specification supports mobility of Wi-Fi hosts in IEEE 802.11 networks (Figure 1-69). ESS is a link-layer (or Layer-2 protocol in the OSI architecture) mechanism for mobility support. Mobile IP allows location-independent routing of datagrams to a mobile host that is identified by its home address. The home address will not change no matter which visiting network the mobile terminal is connected to. When the mobile terminal roams to a visiting network, the visiting network will assign a care-of address to the mobile terminal. The information of this care-of address is sent back to the home agent (HA) in the home network (Figure 8-19). The home agent keeps the association of the care-of address and the mobile terminal's home address. The IP tunnel may be built to connect the mobile

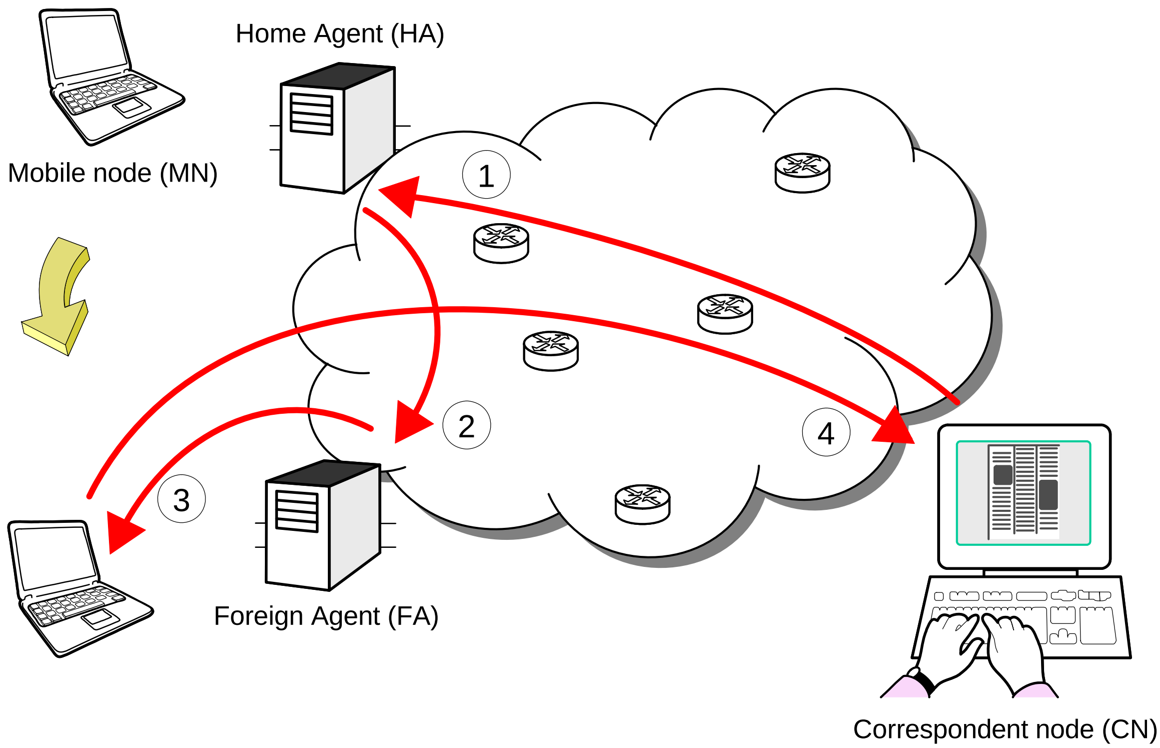

Figure 8-19: Mobile IP.

terminal and the home agent through the Internet cloud. For any packets received in the home agent, the home agent will forward them to the mobile terminal through the tunnel. Mobile IP provides an efficient mechanism for roaming within the Internet. Using Mobile IP, nodes may keep its connection to the Internet without changing their home IP address. Therefore, the location changing in the network is transparent to the correspondent node (CN). Node mobility is realized without the need to propagate the changed location on the network.

8.4 Domain Name System (DNS)

To facilitate network management and operations, the Internet Community has defined the Domain Name System (DNS). An application that wants to send a message to another application on remote computer uses DNS to find the remote computer's IP address given its name. A *domain* refers to a subdivision of a wide area network. Names are hierarchical and a major subdivision is the *top-level domain*, which is broken into organizational and geographic domains. The geographic or country domains use two letters to identify a country and there are 225 country domain labels. For example, *.us* stands for the United States, *.ru* stands for Russia, and *.it* stands for Italy. A name like `university.edu` is found registered with a `.edu` registrar, and within the associated network other names like `mylab.department.university.edu` can be defined, with obvious hierarchy. Security extensions allow a registry to sign the records it contains and in this way demonstrate their authenticity.

Domain Name System is a kind of an address translation protocol (Section 8.3). As shown in Figure 1-37, it translates from the computer name (application-layer address) to a network-layer address. We consider it separately because of its importance and complexity. Unlike the protocols described in Section 8.3, which translate addresses for nodes on the same subnet, DNS resolves host names for hosts anywhere in the Internet.

8.5 Network Management Protocols

Network management tools allow network administrators to monitor network performance, failures, security, and help with accounting management. A basic requirement is to support isolating, diagnosing, and reporting problems to facilitate quick repair and recovery. More advanced features include support for data analytics to predict potential problems, so the network manager can take action before the problem occurs. A number of communication protocols exist for gathering information from network devices. This section reviews some of them.

8.5.1 Internet Control Message Protocol (ICMP)

Internet Control Message Protocol (ICMP) provides a mechanism for communicating control messages and error reports. Both routers and hosts use ICMP to transmit problem reports about datagrams back to the datagram source. In addition, ICMP includes an echo request/reply that can be used to determine if a destination is reachable and if so, is responding. ICMP specifies the format of control messages and when routers should send them. ICMP messages are delivered by IP, as the payload of IP datagrams.

8.5.2 Simple Network Management Protocol (SNMP)

SNMP is the protocol that provides the query language for gathering the information and for sending it to the console. The current version is SNMPv3. In general, the SNMP management system will discover the topology of the network automatically and will display it on the management console in the form of a graph. From this display, the human network manager can select a particular segment of the network to view its status in greater detail.

Each network device hosts a software *agent* that gathers information about the status of that device into a Management Information Base (MIB) and sends it to the *network management system* (NMS), as shown in Figure 8-20(a). SNMP defines seven message types for accessing management information in a client-server relationship (Figure 8-20(b)). Here the NMS is the client and the agent is the server. The message types are as follows:

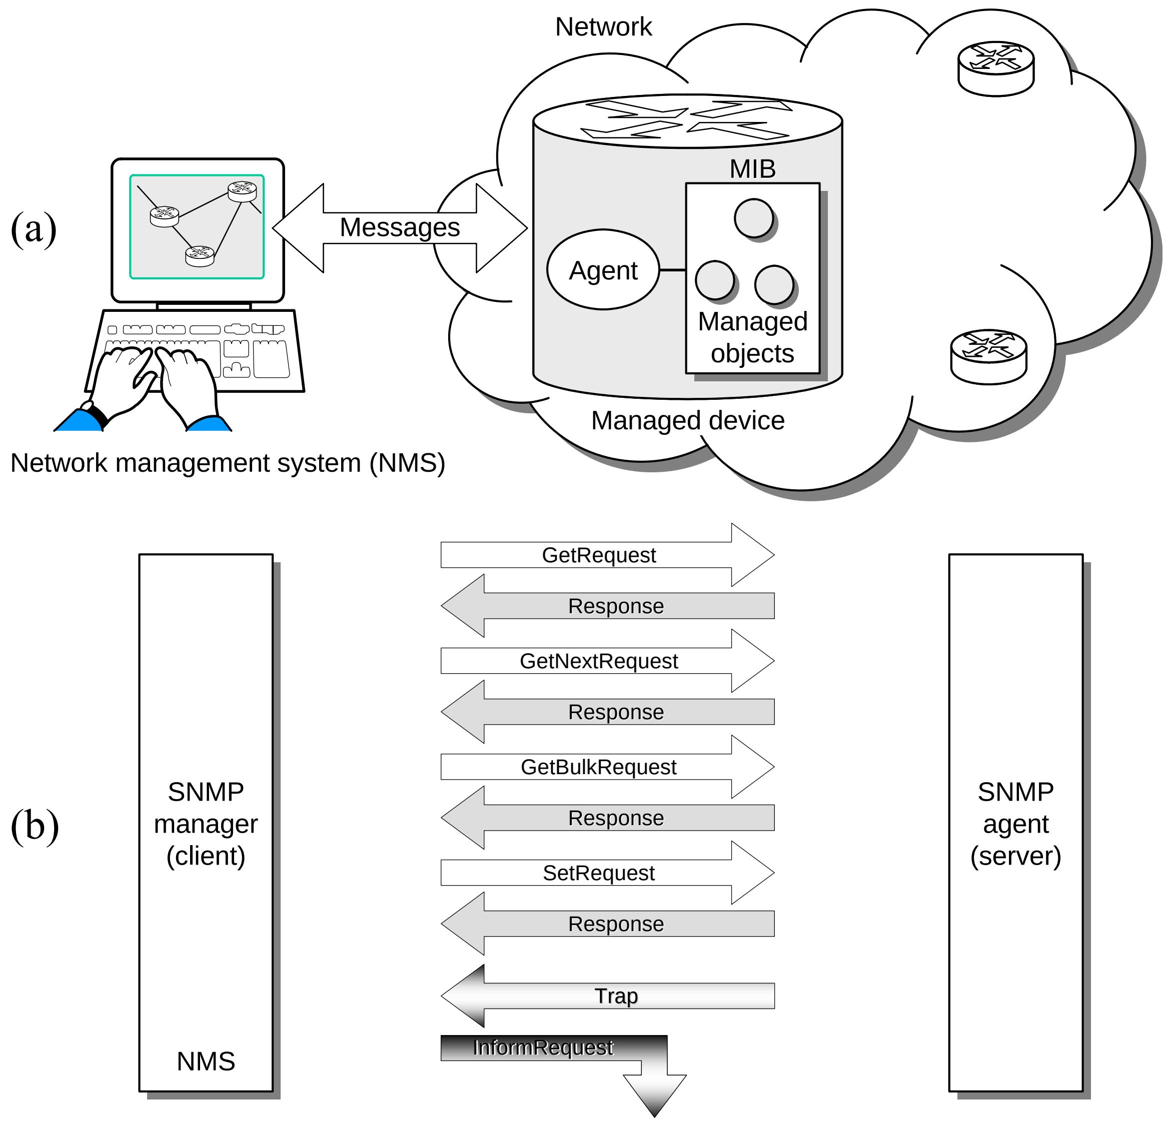

Figure 8-20: (a) SNMP architecture. (b) SNMPv3 message types and their flows.

• *GetRequest*: This is the most commonly used SNMP message and is sent from the manager (NMS) to the agent to retrieve the value of a specific management variable. The manager must send out one GetRequest for each value of a variable that needs to be retrieved.

• *GetNextRequest*: This message is used by an NMS for requesting the next variable in a list or table of variables. It is used mostly to retrieve the values of entries in the MIB if the network manager does not know how many variables there are in the MIB for a certain event.

• *GetBulkRequest*: (not in SNMPv1) This message is sent by an NMS to retrieve large blocks of data, such as multiple rows in a table.

• *Response*: An agent sends this message to a manger/NMS in response to GetRequest or GetNextRequest messages. It contains the value of the variable requested by the manager.

• *SetRequest*: The manager/NMS sends a SetRequest message to an agent to create, store, or modify an information variable. The agent must reply using a Response message.

• *Trap*: An agent sends a Trap message to an NMS to report an event when a certain set of circumstances arises. This is done autonomously, without any request from the manager. For example, if the agent resides in a router that is rebooted, the agent informs the manager of this event and reports the time of the rebooting.

• *InformRequest*: (not in SNMPv1, supports MoM architectures) This message is sent by an NMS to notify another NMS of information in a MIB table that is remote to the receiving manager.

SNMPv2 and SNMPv3 enhanced the original version of SNMP (SNMPv1) with additional message types for establishing multiple manager entities within a single network to support distributed management. *Distributed management* means that NMSs and agents are spread out across the internetwork. A hierarchical distributed arrangement can be used, whereby distributed NMSs send data to more powerful central NMSs using a Manager-of-Managers (MoM) architecture. A centralized system that manages distributed NMSs is sometimes called an *umbrella NMS*. The advantages of distributed management include resilience to failures and the ability to reduce network-management overhead by filtering data at distributed NMSs before sending them to the central stations. On the downside, distributed management is complex and hard to operate, and more susceptible to security breaches.

SNMP is an application-layer protocol and it operates over the user datagram protocol (UDP). As a result, SNMP is a connectionless protocol. No continuous connections exist between a management console and its agent so that each message between them is a separate transaction.

8.6 Multimedia Application Protocols

Multimedia application protocols include RTP and RTCP, which are described in Section 3.3.1.

8.6.1 Session Description Protocol (SDP)

Peer ends of a multimedia application use the Session Description Protocol (SDP), to offer and accept (or not) codecs, decide the port number and IP address for where each endpoint wants to receive their RTP packets (Section 3.3.1). SDP packets are transported by SIP.

8.6.2 Session Initiation Protocol (SIP)

The Session Initiation Protocol is an application-layer control (signaling) protocol for creating, modifying and terminating multimedia sessions on the Internet, meant to be more scalable than H.323. Multimedia sessions can be voice, video, instant messaging, shared data, and/or subscriptions of events. SIP can run on top of TCP, UDP, SCTP, or TLS over TCP. SIP is independent of the transport layer, and independent of the underlying IPv4/v6 version. In fact, the transport protocol used can change as the SIP message traverses SIP entities from source to destination. SIP itself does not choose whether a session is voice or video—the SDP does it.

8.7 Summary and Bibliographical Notes

The best source of information about the Internet protocols are Requests for Comments (RFCs) published by the Internet Engineering Task Force (IETF), which can be found online at: http://www.ietf.org/rfc.html. For a complete listing of all protocols in the IP stack, their standards statuses, and reference to their RFC documents, see http://www.rfc-editor.org/rfcxx00.html, which is updated daily.

Section 8.1: Internet Protocol Version 6 (IPv6)

The IETF solicited proposals for a next generation Internet Protocol (IPng) in July of 1992. A number of proposals were received and by 1994 the design and development of a suite of protocols and standards now known as Internet Protocol Version 6 (IPv6) was initiated. A major milestone was reached in 1995 with the publication of RFC-1752 ("The Recommendation for the IP Next Generation Protocol"). Overall specification of IPv6 is defined in RFC-2460 [Deering & Hinden, 1998]. The address structure of IPv6 is defined in RFC-2373. The format of IPv6 global unicast addresses is defined in RFC-3587 (obsoletes RFC-2374). Great deal of information about IPv6 can be found at the IPv6.com webpage, at http://www.ipv6.com. Also [Huitema, 1998]

At the time of this writing (2010), IPv6 is not widely adopted. One of the greatest obstacles to wider adoption of IPv6 is that it lacks backwards compatibility with IPv4. However, fewer than 10% of IPv4 addresses remain unallocated and industry experts predict the rest of the IPv4 address supply will run out in 2012. Therefore, it can be expected that the adoption rate will grow rapidly. The IPv6 Forum (http://www.ipv6forum.com/) verifies protocol implementation and validates interoperability of IPv6 products. The IPv6 Forum has a service called IPv6 Ready Logo (http://www.ipv6ready.org/), which is a qualification program that assures devices they test are IPv6 capable. Actual testing in the U.S. is performed by the IPv6 Testing Consortium at the University of New Hampshire (http://www.iol.unh.edu/services/testing/ipv6/), which is a pioneer in IPv6 testing.

Section 8.2: Routing Protocols

The Routing Information Protocol (RIP) was the initial routing protocol in the ARPAnet. RIP was originally designed for Xerox PARC Universal Protocol (where it was called GWINFO) and used in the Xerox Network Systems (XNS) protocol suite. RIP became associated with both UNIX and TCP/IP in 1982 when the Berkeley Software Distribution (BSD) version of UNIX began shipping with a RIP implementation referred to as `routed` (pronounced "route dee"). RIP, which is still a very popular routing protocol in the Internet community, is formally defined in the XNS Internet Transport Protocols publication (1981) and in RFC-1058 (1988). RIP version 2 (for IPv4) is defined in RFC-1723. This document does not change the RIP protocol per se; rather, it provides extensions to the message format that allows routers to share important additional information.

Open Shortest Path First (OSPF) is a link-state routing protocol defined in RFC-2328. OSPF was designed to advertise the subnet mask with the network. OSPF supports variable-length subnet masks (VLSM), disjointed subnets, and supernetting.

Border Gateway Protocol (BGP) was designed largely to handle the transition from a single administrative entrity (NSFNet in the US in 1980s) to multiple backbone networks run by competitive commercial entities. Border Gateway Protocol version 4 (BGP4) is defined in RFC-4271 [Rekhter *et al.*, 2006], which obsoletes RFC-1771. See also RFC-4274 and RFC-4276. [Stewart, 1999] provides a concise overview of BGP4, although it does not include the latest updates. BGP4 and CIDR (Section 1.4.4) have played a key role in enabling the Internet to scale to its current size. Large ISPs use path aggregation in their BGP advertisements to other Autonomous Systems. An ISP can use CIDR to aggregate the addresses of many customers into a single advertisement, and thus reduce the amount of information required to provide routing to customers.

IETF specified routing protocols that work with IPv6 include RIP for IPv6 [RFC-2080], IS-IS for IPv6 [RFC-5308], OSPF for IPv6 [RFC-5340], and BGP-4 for IPv6 [RFC-2545].

Section 8.3: Address Translation Protocols

Address Resolution Protocol (ARP) is defined in RFC-826. Reverse Address Resolution Protocol (RARP) is defined in RFC-903. Inverse Address Resolution Protocol (InARP) is defined in RFC-2390. RARP is now obsolete (succeeded by DHCP), and InARP is primarily used in Frame Relay and ATM networks. Neighbor Discovery (NDP) which is used for discovery of other nodes on the link and determining their link layer addresses for IP version 6 (IPv6) is described in RFC-4861.

Network Address Translation (NAT) is described in RFC-2663 [Srisuresh & Holdrege, 1999] and RFC-3022 [Srisuresh & Egevang, 2001].

Section 8.4: Domain Name System (DNS)

The Domain Name System (DNS) is defined in RFC-1034 and RFC-1035.

Section 8.5: Network Management Protocols

The first version of Simple Network Management Protocol, SNMPv1, was defined by RFC-1067 in 1988. SNMPv2 was introduced in 1993 and updated in 1996. SNMPv1 and SNMPv2 supported monitoring network statistics, but had no security features. SNMPv3 is the current standard version (see RFC-3411). SNMPv3 offers security features and is expected to displace the earlier versions. SNMPv3 supports user authentication to prevent unauthorized users from executing network management functions, and message encryption to prevent eavesdropping on the messages exchanged between NMSs and managed devices. SNMP is supported by most commercial network management systems (NMSs) and many networking devices, including switches, routers, servers, and workstations.

The current version of management information base (MIB) for SNMP is defined in RFC-3418.

The statistics of RMON that should be collected are standardized in RFC-1757 and RFC-2021.

Section 8.6: Multimedia Application Protocols

The Session Initiation Protocol is defined in RFC-3261, RFC-3265, RFC-3853, RFC-4320, RFC-4916, RFC-5393, and RFC-5621.

Problems

Chapter 9
Technologies and Future Trends

Contents

burst.com Technology-Burst vs. HTTP streaming

http://www.burst.com/new/technology/versus.htm

9.1 Network Technologies

The often mentioned virtues inherent to wireless networks are in the low cost-to-solution and rapid deployment possibilities. Already in many emerging markets, cellular phones are far more prevalent than computers, driving many in networking technology to consider non-traditional computing deployments that leverage cellular technology instead of optical, cable, DSL or other data circuits. Even in more mature markets, one has to wonder what role smartphones can play in keeping people connected to data sources, documents, and communications media.

The evolution to away from cellular technology to an all IP-based mobile connection also opens up completely new spheres of functionality for roaming employees who need access to network resources. The improved security of IP networking combined with the increased bandwidth of LTE / 4G will allow users to work more efficiently and increase their productivity even if equipped with nothing more than their mobile phone.

9.1.1 Mobile Wi-Fi

Microsoft's ViFi project uses smarter networking to eliminate Internet outages during travel.

Users are likely to see increasing discrepancy between peak (the marketing numbers) and realized throughput. Lucent ORiNICO 802.11b outdoors, no obstruction—these are practically ideal conditions!

Transmission rate	% of coverage area
11 Mbps	8 %
5.5 Mbps	
2 Mbps	
1 Mbps	47 %

Thus, a low probability of having good link!!

One challenge in making shared Wi-Fi LANs deterministic, in part, has been in figuring out whether a client is operating at a low throughput rate because it only has a small amount of data to send or because it's encountering congestion and is out of bandwidth.

IEEE 802.11n (Section 6.3.1) offers tenfold throughput increase and extended range. Where will it be used? In the traditional application: of wireless Internet delivery. But, 802.11n also offers new opportunities for wireless applications:

- Voice over IP (VoIP)
- Professional and personal multimedia delivery
- Wireless storage devices

9.1.2 Wireless Broadband

4G: WiMAX and LTE

WiMAX

“WiMAX 2” coming in 2011: 802.16m standard promises faster data rates, backward compatibility with WiMAX.

LTE (Long Term Evolution)

Cellular Network Backhaul

Backhaul is the connection between infrastructure nodes, most commonly cell sites, and the rest of the network, like a mobile switching center or similar elements. As is the case in enterprise networks, 1.544 Mbps T1 and 2.048 Mbps E1 connections have dominated this space for many years, because (a) they were easy, if not cheap, to get from wireline carriers, and (b) they are well-suited to telephony because they are telephony. However, these data rates no longer appear to be adequate for the carrier needs. Carriers have difficulty with offering data services of the future in part because they lack the backhaul capacity necessary to offer the megabits of service to their individual data users that are becoming feasible with EV-DO Rev A, WiMAX, HSPA, and LTE. Backhaul has become a bottleneck, which will have to change if 3.5G/4G systems are to have sustainable multi-megabit throughput.

Recently, some companies are offering high-capacity point-to-point wireless systems designed for use in carrier networks as backhaul links. However, although the carriers sell wireless, wireless is not regularly the solution of choice for backhaul. Proponents of wireless-for-backhaul point that there are great regions of spectrum, including unlicensed spectrum around 5.8 GHz and bands at 60 and 80 GHz, that are broadly available and that would work adequately in this application.

9.1.3 Ethernet

Network technologies have come and gone, but the Ethernet switch has survived and continues to evolve, currently making its move to 40 Gbps and 100 Gbps datarates and there is a talk of Terabit Ethernet. Although 3.2Tbps and 6.4Tbps speeds were demonstrated in test environments by Siemens/WorldCom and NEC/Nortel, respectively, starting in 2001, the first set of viable solutions are just now taking shape.

Data centers drive Ethernet switch market in the near term. Today, the largest data centers contain over 100,000 servers. Ideally, one would like to have the flexibility to run any application on any server while minimizing the amount of required network configuration and state.

How to build a 100,000-port Ethernet switch:

R. N. Mysore, A. Pamboris, N. Farrington, N. Huang, P. Miri, S. Radhakrishnan, V. Subramanya, and A. Vahdat, "PortLand: A scalable fault-tolerant layer 2 data center network fabric," *Proceedings of ACM SIGCOMM '09*, Barcelona, Spain, August 2009.

[Mysore, *et al*., 2009] describe software that could make data center networks massively scalable. Ttheir PortLand software will enable Layer 2 data center network fabrics scalable to 100,000 ports and beyond. The goal is to allow data center operators to manage their network as a single fabric. They observe that data center networks are often managed as a single logical network fabric with a known baseline topology and growth model. They leverage this observation in the design and implementation of PortLand, a scalable, fault tolerant layer 2 routing and forwarding protocol for data center environments. Key to this is the development of system for servers to find one another without broadcasting their requests across an entire network. Under PortLand, switches use what are called Pseudo MAC addresses and a directory service to locate servers they need to connect, including new virtual servers. The researchers say this setup can eliminate much of the manual labor required to build a Layer 3 network. The software will work with existing hardware and routing protocols.

Also see: http://www.eurekalert.org/pub_releases/2009-08/uoc--css081709.php

The IEEE802.3at Power over Ethernet (PoE) standard

9.1.4 Routers and Switches

Cisco is reportedly soon announcing a new carrier core router as a next-generation follow on to the 6-year-old CRS-1, and a better competitor to Juniper's T1600. Cisco MSC120 will bring 120G per slot to CRS-1, to better compete with Juniper T1600.

Juniper's answer to Cisco in the data center: Stratus Project

Juniper's Stratus Project is a year old and comprises six elements: a data center manager, storage, compute, Layer 4-7 switching, appliances and networking. It is intended to be a flat, non-blocking, lossless fabric supporting tens of thousands of Gigabit Ethernet ports, an order of magnitude reduction in latency, no single point of failure, and with security tightly integrated and virtualized. Stratus will be managed like a large JUNOS-based switch.

9.2 Multimedia Communications

One way to think about future developments is to consider the application needs for computing and communication resources (Table 9-1).

Video is having a profound effect on the way we consume information. It is estimated that at the time of this writing (2010) video represents approximately one-quarter of all consumer Internet traffic. Thirteen hours of video are uploaded each minute on YouTube alone.

9.2.1 Internet Telephony and VoIP

Voice over IP (VoIP) is an umbrella term used for all forms of packetized voice, whether it is Internet telephony, such as Skype.com, or Internet telephony services provided by cable operators. Voice over IP is also used interchangeably with IP telephony, which is very much enterprise focused. And, there the problems with service quality are very real.

IP telephony is really a LAN-based system and as an application inside the enterprise, it is going to be a pervasive application. Voice over IP is growing in popularity because companies are attracted to its potential for saving money on long distance and international calls. The evolution from a LAN-based system to the broader context of the Internet is not straightforward. Integrating the Voice over IP that may be on a LAN and the Voice over IP that is going to be Internet-based is going to become a reality.

Traditional telephone services have typically gained a reputation of providing excellent voice

Table 9-1: The application hierarchy:

Data Type / Resource	Text	Audio	Video	?
Transmission Rate	10 bits/sec	10 Kbits/sec	10 Mbits/sec	10 Gbits/sec
Processing	KIPS	MIPS	GIPS	TIPS
Memory	KB	MB	GB	TB

quality and superior reliability. Consequently, users take for granted that their phone systems will provide high quality with virtually no downtime. Yet many VoIP installations fail to meet these expectations, primarily because organizations have not adequately evaluated their network infrastructure to determine whether it can adequately support applications that are very sensitive to latency, packet loss, jitter and other similar performance factors.

In fact, call quality and availability are expected to vary between services within certain acceptable boundaries. Consumers make trade-offs based on price, accessibility, and mobility, and it is important to understand that mix. If you were using a home telephony product from your cable company, they would offer a different grade of service than a free service like Skype. Consumers put up with a little call quality degradation for cheapness. As companies increasingly adopt VoIP to replace the traditional PSTN, there are several major concerns:

Reliability Concerns

The traditional public switched telephone network (PSTN) provides very high reliability and users have come to expect that when they pick up the phone, they get a dial tone. Computers sand computer networks are still lacking this degree of reliability. This is particularly critical for companies, because they depend on the phones to stay in contact with customers, partners, and vendors, as well as within the company for employee communications. A phone outage can have significant impact on company operations. In addition, the regular phone lines in buildings are independent of electric lines, so phone work even during power outages. VoIP depends on both electrical power and Internet service. Interruption of either means losing phone service. The problem can be mitigated by having redundant Internet connections and power backup such as a generator, but this adds to the cost.

Network Quality-of-Service

The reader knows by now that delays in transmission or dropped packets cause a disrupted phone call, which the call participants will surely notice and complain about. To help prevent such problems, the IP network must support quality-of-service (QoS) mechanisms that allow administrators to give priority to VoIP packets. This means a VoIP network is more trouble to manage than a data network, and it requires a higher level of expertise—or at least an additional skill set—on the part of network administrators.

Although many switch and router vendors will advertise that their products can handle a certain level of throughput, few talk about the volume of packets that can be processed during periods of peak utilization. For example, even though a switch might be able to accommodate a full line rate traffic stream when all packets are nearly maximum size, it may not be able to manage the same aggregate throughput when the stream is composed of many more minimum-sized packets. Because most Real-Time Protocol (RTP) audio packets are relatively small (just over 200 bytes for G.711), a device's ability to process packets of that size at full line rate must be assured. Understanding how a device reacts to traffic streams characterized by many short bursts of many packets is also important.

VoIP monitoring and management solutions are available that make it easier to optimize voice services, but that adds to the cost of deployment. It also negates some of the cost savings that motivate the move to VoIP in the first place.

Other factors, such as external microphones and speakers, Internet connection speeds, and operating systems, also can affect call quality and should be taken into account before writing off a service provider's performance as poor. In addition, the analog-to-digital conversion process can affect VoIP call quality, causing users to experience unpleasant distortion or echo effects. Another culprit is signal level problems, which can cause excessive background noise that interferes with conversations. It does not tend to be as much of a service problem as it is an access or device problem for the consumer.

Complexity and Confusion

The complexity and unfamiliar terrain of VoIP communications presents another big obstacle for many companies. Network administrators experienced with running a data network may not know much about how VoIP works, what equipment is necessary, or how to set up and maintain that equipment.

In addition, VoIP terminology quickly gets confusing—media gateways, analog telephone adapter (ATA), audio response unit (ARU), interactive voice response (IVR), etc. Company managers and IT personnel hear about different VoIP protocols—H.323, SIP, IAX—and do not understand the differences or know which one they need.

Already overworked IT staffs may not be eager to undertake the task of learning a completely new specialty nor the added burden of ongoing maintenance of the components of a VoIP system. They may not be sure how to integrate the VoIP network into the existing data network.

Of course, there are answers to these problems. Consultants with the requisite knowledge can help set up a VoIP network, or companies can use hosted VoIP services to reduce both the complication and the upfront expenses of buying VoIP servers. However, once again, this ups the price tag of going to VoIP and eats into the cost savings that are one of VoIP's main advantages.

Security

There is also an issue of securing IP telephony environments. The risk of intercepted calls and eavesdropping are a concern. Although possible, it is difficult to tap traditional telephone lines. Traditional phone communications travel over dedicated circuits controlled by one entity—the phone company. But when VoIP packets go out there into the "Internet cloud," they go through numerous routers and servers at many different points. Potential types of threats:

- *Toll Fraud*: the use of corporate resources by inside or outside individuals for making unauthorized toll calls.
- *Denial of Service Attacks*: attacks typically aimed at the data network and its components that can have a severe impact on voice calling capabilities.
- *Impersonation Exploits*: where a caller changes call parameters, such as caller id, to make the call look like it is originating from a different user. The caller id may be used to gain a level of trust about the caller, who may then proceed to get private information that might not have been otherwise given.
- *Eavesdropping*: the ability for a hacker to sniff packets relating to a voice call and replay the packets to hear the conversation.

Encryption and other security mechanisms can make VoIP as secure as or even more secure than PSTN. We should encrypt our voice inside the LAN, and the same applies to data and video in the long run. Mixing data and voice (or, multimedia) raises another concern. It is not an IP telephony or voice over IP issue; it is an IP issue, one should not be lulled into the suspicion that IP or the layers above it are secure. We have already seen vulnerabilities against PBXs, against handsets, so it is only a matter of time before we see execution against these vulnerabilities. ... attacks at the server level or at massive denial-of-service attack at the desktop level ...

However, extra security mechanisms mean extra cost. Moreover, a simple perception that data networks are inherently insecure may hold back VoIP adoption. Addressing all the above issues, while keeping VoIP costs lower than the costs of traditional phone service, will be a challenge.

9.2.2 Unified Communications

Unified communications (UC) is an umbrella term for integrated, multi-media communications controlled by an individual user for both business and social purposes. UC is supposed to offer the benefits of seamless communication by integrating voice, data, video, and presence into a single environment. UC refers to a real-time delivery of communications based on the preferred method and location of the recipient. Its purpose is to optimize business processes and enhance human communications by reducing latency, managing flows, and eliminating device and media dependencies.

With an increasingly mobile workforce, businesses are rarely centralized in one location. Unified communications facilitates this on-the-go, always-available style of communication. In addition, unified communications technology can be tailored to each person's specific job or to a particular section of a company.

9.2.3 Wireless Multimedia

A study from Pew Internet & American Life Project predicts the mobile phone will be the primary point of Internet access, while technologies such as touch-screen interfaces and voice recognition will become more prevalent by the year 2020 (http://www.pewinternet.org/PPF/r/270/report_display.asp). The participants predict telephony will be offered under a set of universal standards and protocols accepted by most operators internationally, making for reasonably effortless movement from one part of the world to another. At this point, the "bottom" three-quarters of the world's population account for at least 50 percent of all people with Internet access, up from 30 percent in 2005.

For the small-to-medium-size business (SMB) owner, the results of the survey suggest international transactions and growth will be made easier by a more internationally flexible mobile infrastructure, while the prevalence of the Web on mobile devices and smartphones, which the survey predicts will have considerable computing power by 2020, will allow SMB owners access to their business dealings nearly anytime and anywhere.

Video phones

9.2.4 Videoconferencing

Videoconferencing systems have recently become popular for reasons including cutting down travel expenses (economic), as well as "going green" (environmental). A recent study by Ferran and Watts [2008] highlights the importance and diverse aspects of perceived "quality-of-service." The study surveyed 282 physicians who attended grand rounds (presentations on complex cases) in person or by video. The video attendees were twice as likely to base their evaluation of the meetings on the speaker rather than the content. They were also more likely to say the speaker was hard to follow. The authors speculate that our brains gather data about people before turning to what they say. In person, we do this quickly. However, speakers are harder to "read" on screen, so we focus more on them.

Videoconferencing More Confusing for Decision-makers than Face-to-face Meetings: http://www.sciencedaily.com/releases/2008/10/081028184748.htm

Study Finds Videoconferences Distort Decisions: http://www.wtop.com/?nid=108&sid=1506873

Immersive Videoconferencing

New "telepresence" video systems, which create the illusion of sitting in the same room—even allowing eye contact between participants—may help solve the problem. Products in this category include Cisco WebEx, Citrix GotoMeeting, Microsoft Office Live Meeting and IBM Lotus Sametime Unyte—all of which let users conduct meetings via the Internet. However, some of these systems cost up to $300,000.

9.2.5 Augmented Reality

Expert Discusses Importance Of Increasingly Accessible 3-D CAD Data.

http://www.industryweek.com/ReadArticle.aspx?ArticleID=17986

3-D CAD Data Grows More Accessible -- For manufacturers, that means less wasted time.

IndustryWeek (1/1, Jusko) reports, "3-D CAD is widely used in organizations' product design function. In fact 60% of respondents to a recent survey say they use it 81% to 100% of the time." According to David Prawel, founder and president of Longview Advisors, "the potential usefulness of 3-D CAD data is enormous in downstream functions." Prawel said that, "not only would it help the machine operator to have a view of the part in 3-D, but the operators could add more value to the process," as the operators "could see the part in front of them and make their own course corrections and feed that information back to the designer." Prawel also predicted that "use of 3-D data on the shop-floor will grow ... not 3-D CAD itself, but the use of 3-D representations." He said that "the 'democratization' of 3-D CAD data is beginning to occur" by "moving out of the CAD department and into greater use throughout the organization." Prawel added, "The important thing is to get it to the people who need it in a language they understand."

9.3 The Internet of Things

The *Internet of Things* refers to a network of objects, such as household appliances. It is often a self-configuring wireless network. The idea is that if cans, books, shoes, or parts of cars were equipped with miniature identifying devices, our daily life would undergo a transformation. No longer would supermarkets run out of stock or waste products because they will know exactly what is being consumed at any specific location. If everyday objects, from yogurt to an airplane, are equipped with radio identification (RFID) tags, they can be identified and managed by computers in the same way humans can. The next generation of Internet Protocol (IPv6) has sufficiently large address space to be able to identify any kind of object. RFID tag can contain information on anything from retail prices to washing instructions to person's medical records.

A movement is underway to add any imaginable physical object into the Internet of Things. In Japan, for example, many cows have IP addresses embedded onto RFID chips implanted into their skin, enabling farmers to track each animal through the entire production and distribution process. This points to a future where pretty much everything is online. Put simply, the Internet of Things in essence means building of a global infrastructure for RFID tags. You could think of it as a wireless layer on top of the Internet where millions of things from razor blades to grocery products to car tires are constantly being tracked and accounted for. It is a network where it is possible for computers to identify any object anywhere in the world instantly.

* Emerging applications and interaction paradigms

o using mobile phones and other mobile devices as gateways to services for citizens

o integrating existing infrastructure in homes (digital picture frames, smart metering of energy...)

o embedding virtual services into physical artifacts

o the electronic product code (EPC) network aims to replace the global barcode with a universal system that can provide a unique number for every object in the world.

Of course, as with any technology, there is potential for misuse, such as privacy invasion. The Internet of Things could be the ultimate surveillance tool.

9.3.1 Smart Grid

"Smart Grid" is the term used for modernization of the electricity grid that involves supporting real-time, two-way digital communications between electric utilities and their increasingly energy-conscious customers. It will provide electric utilities with real-time visibility and control of the electricity used by customers. Having a Smart Grid is considered vital to the development of renewable energy sources such as solar and wind as well as plug-in electric hybrid vehicles. Smart Grid is expected to enable buildings with solar panels and windmills to inject power into the grid. This means adding all kinds of information technology, such as sensors, digital meters and a communications network akin to the Internet, to the dumb wires. Among other things, a smart grid would be able to avoid outages, save energy and help other green undertakings, such as electric cars and distributed generation.

It is expected that the Smart Grid will enable integrating renewable energy and making more efficient use of energy. The part of Smart Grid that most users would see is at the distribution level, where a resident has a smart meter that is their building's interface into the grid. With Smart Grid, customer has the ability for a smart metering infrastructure to send near real-time measurements of his or her energy usage. This infrastructure will also to be able to signal when there are overload conditions on the grid and there needs to be some demand response to adjust the load. Sensors on transmission lines and smart meters on customers' premises tell the utility where the fault is and smart switches then route power around it. That is similar to the Internet, which redirects data packets around failed nodes or links.

On the other hand, bulk power generation plants at the heart of the grid need to have real-time controls for supervisory control and data acquisition (SCADA) systems—computer systems monitoring and controlling the process. These have very different communications requirements, where reliability and security are key, and quality of service attributes such as latency are also important. A relay needs to close in milliseconds, which is a very different requirement than gathering electricity usage information from a resident every 15 minutes.

Electric grid already implements what is called "demand response," where some big companies have agreed to throttle back their consumption at times of peak demand. With a Smart Grid, all consumers would be able to do the same. In a basic version, they would get real-time information about their usage and could then turn off the tumble dryer or other energy-hungry appliances. If prices also varied with a grid's load, rising when demand was heavy, customers would cut back their consumption during peak hours. That reduction would increase if smart meters could turn appliances off automatically should rates rise above a certain point. With peak demand lower, utilities would no longer have to hold as much expensive backup capacity.

More intelligence in the grid would also help integrate renewable sources of electricity, such as solar panels or wind turbines. Currently, the problem is that their output is highly variable, because it is tightly coupled with the weather conditions. A standard grid becomes hard to manage if too many of renewable sources are connected to it; supply and demand on electricity-transmission systems must always be in balance. A Smart Grid could turn on appliances should, for instance, the wind blow more strongly. Added intelligence would also make it much easier to cope with the demand from electric cars by making sure that not all of a neighborhood's vehicles are being charged at the same time.

Because there are great profits to be made, the Smart Grid market has attracted the attention of every major networking vendor. Cisco expects that the underlying communications network will be 100 or 1,000 times larger than the Internet. These vendors are pushing for Smart Grid to adopt common network standards rather than special-purpose protocols.

A key characteristic of the electricity grid in the U.S. is that it is highly fragmented: 80% is owned and operated by private companies, including about 3,100 electric utilities. This is really a system of systems, which is highly complex, and therefore reliability is a great concern. It is also well known that systems transmitting and distributing electricity are exceedingly wasteful and vulnerable. Smart grids increase the connectivity, automation and coordination between these suppliers, consumers and networks that perform either long distance transmission or local distribution tasks.

The Smart Grid has to have a very robust communications infrastructure underlying it. Standards are highly important because they provide a common set of network protocols that can run end-to-end over a variety of underlying physical and link layer technologies. In terms of the protocol stack, the functionality is at the application layer to support features such as consumer energy management, electrical vehicles, etc. There is consensus that IP and the Internet standards will be a protocol of choice in the Smart Grid.

However, some high-performance command and control applications could require special-purpose network protocols. SCADA systems are different because the requirement is real-time control of a critical asset (response times are in milliseconds), rather than routing data around a network. Therefore, a specialized protocol historically has been used and may still have a role. On the other hand, to support communications with the hundreds of millions of devices that will interact with the Smart Grid ("smart appliances"), IP has great advantages in terms of ubiquity, implementation and its ability to create interoperable infrastructure as a low cost.

An important requirement for the Smart Grid needs to account for the fact that the environments are so varied. The electric grid in an urban area like New York City is very different from that in a rural environment in Montana. These impose communications requirements, such as the ability to use different physical and link layer technologies to move data around effectively. IP provides that, as well as the ability to evolve and have an open architecture in which the industry can find better ways to help customers manage energy.

It is believed that the Smart Grid will be one of the drivers for greater adoption of IPv6 (Section 8.1). Given that millions of appliances and devices need to be addressable, IPv6 is necessary because the IPv4 address space is exhausted. The security issues about smart meters are very important. The meters themselves are not consumer devices. They are part of the utility's infrastructure. There are about 150 million meters in the U.S. associated with buildings. One of the functions of those meters is to connect and disconnect to a building so a utility does not have to send maintenance crews. Of course, this presents a danger of having an architecture that that is vulnerable to cyber criminals. A poorly designed system would allow a hacker to remotely disconnect 150 million meters from the grid. This is why that needs to be a locked-down architecture.

Another argument for IPv6 in the Smart Grid meter interface is that utilities will likely use wireless networks to communicate with thousands of meters through a management gateway or router. This requires a protocol where all these devices will be put in a subnet. Although this is possible with IPv4, it is easy with IPv6.

Advanced Metering Infrastructure

Key tasks of the Smart Grid are evaluating congestion and grid stability, monitoring equipment health, energy theft prevention, and control strategies support. To support these tasks, advanced sensing and measurement infrastructure is at the heart of every smart grid. Technologies include: advanced microprocessor meters ("smart meter") and meter reading equipment, wide-area monitoring systems, dynamic line rating (typically based on online readings by distributed temperature sensing combined with real-time-thermal-rating (RTTR) systems), electromagnetic signature measurement/analysis, time-of-use and real-time pricing tools, advanced switches and cables, backscatter radio technology, and digital protective relays.

Smart meters resemble smart-phones: they have a powerful chip and a display, and are connected to a communications network. The main task of a metering system is to get information reliably into and out of meters—for example, how much power is being used, when and at what price. In Europe, this is mostly done by using power lines to communicate. However, in America this would be too costly. The grid's architecture does not allow it to be turned into a data network easily. Using a public cellular network would also be hard. A meter cannot move to get better reception, for instance. The best approach is to use wireless mesh networks (Section 6.1), in which data are handed from one meter to the next. Such networks automatically reconfigure themselves when new meters are added.

A key component will be the software that makes the Smart Grid work, and the applications that run on it. Power system automation enables rapid diagnosis of and precise solutions to specific grid disruptions or outages. These technologies rely on and contribute to each of the other four key areas. Three technology categories for advanced control methods are: distributed intelligent agents (control systems using artificial intelligence programming techniques), analytical tools (software algorithms and high-speed computers), and operational applications (SCADA, substation automation, demand response, etc).

Usage Data Management and Demand-Driven Rate Adjustment

A key technology that an electric utility needs will allow it to manage the usage data, combine the data with other information, and set rates depending on demand. This will require large databases for data repositories and data-mining software for detecting trends and usage patterns. Information systems that reduce complexity so that operators and managers have tools to effectively and efficiently operate a grid with an increasing number of variables. Technologies include visualization techniques that reduce large quantities of data into easily understood visual formats, software systems that provide multiple options when systems operator actions are required, and simulators for operational training and "what-if" analysis.

Home Area Networks (HANs)

Home Area Network (HAN) covers all the smart-grid technology in the home, behind the meter. HAN will include things such as wireless displays that show the household's power consumption at that instant, thermostats that are connected to the meter and smart appliances that can be switched on and off remotely. The big question is how all these devices will be connected and controlled. Will the HAN be dedicated to regulating electricity consumptions, for instance, or will it also control home security or stream music through the rooms? Figure 9-1 illustrates a typical smart home network environment.

One option is to have an integrated HAN that allows consumers to control almost everything in a house that runs on electricity. This would include systems for secure home access used to keep burglars out (cameras, sensors, etc.), as well as systems managing energy consumption.

Google and Microsoft have launched web-based services, called PowerMeter and Hohm respectively, that allow households to track their power usage—and, at a future point, their operators to sell more advertising.

ITU standard G.hn for home networks

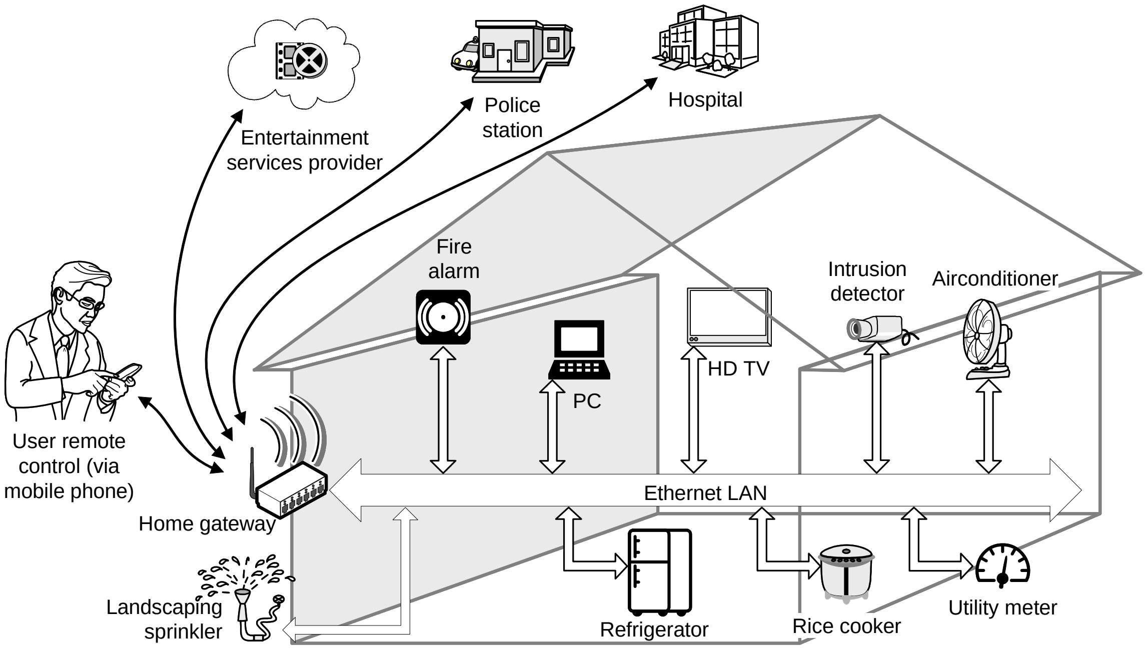

Figure 9-1: A typical smart home network environment.

9.3.2 The Web of Things

The *Web of Things* is a vision inspired from the Internet of Things where everyday devices and objects are connected by fully integrating them to the World Wide Web. Unlike in the many systems that exist for the Internet of Things, the Web of Things is about reusing the Web standards to connect the quickly expanding eco-system of embedded computers built into everyday smart objects. Widely adopted and well-understood standards (such as URI, HTTP, RSS, etc.) are used to access the functionality of the smart objects.

Given a large body of data continuously collected by the networked devices, the idea has arisen that the Web can learn inferentially from the data. The Web would eventually begin to "understand" things without developers having to explicitly explain to it.

9.4 Cloud Computing

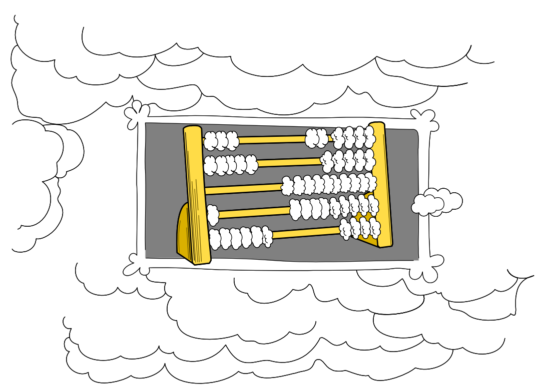

Cloud computing is the latest incarnation of an old idea that in 1960s was called "timesharing" with dumb terminals and mainframe computers, and General Electric opened the first commercial timesharing service in 1965. In 1990s was called network computing with thin clients, and in 2000s has also been called "utility computing," or "grid computing." Cloud computing can be loosely defined as using minimal terminals to access computing resources provided as a service from outside the

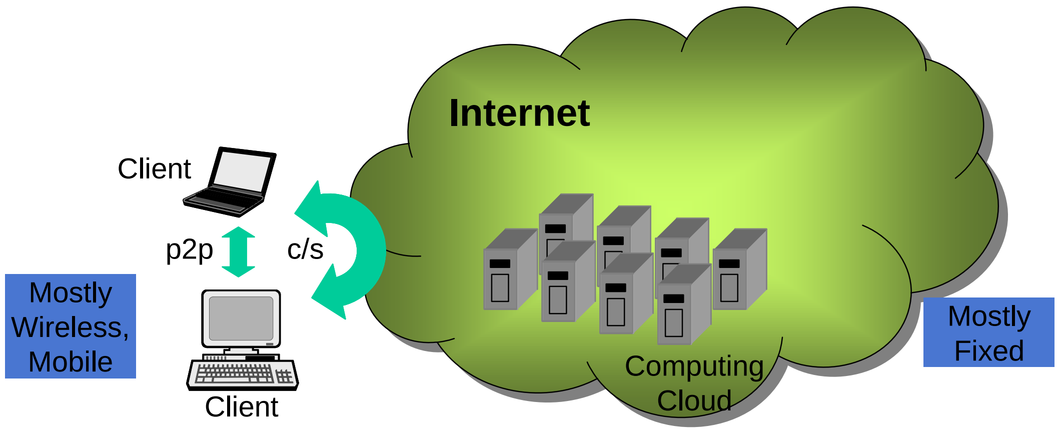

Figure 9-2: Current shift in the industry toward cloud computing and large-scale networks.

environment on a pay-per-use basis. One uses only what one needs, and pay for only what one uses. One can access any of the resources that live in the "cloud" at any time, and from anywhere across the Internet. One does not have to care about how things are being maintained behind the scenes in the cloud. The cloud is responsible for being highly available and responsive to the needs of the application. The name derives from the common depiction of a network as cloud in network architecture diagrams.

The idea is that companies, instead of spending time and resources building their own infrastructure to gain a competitive advantage, they would outsource these tasks to companies that specialize in cloud computing infrastructure. The proprietary approach created large tracts of unused computing capacity that took up space in big data centers and required large personnel to maintain the servers. It also has had associated energy costs. The unused computing power wasted away, with no way to push it out to other companies or users who might be willing to pay for additional compute cycles. With cloud computing, excess computing capacity can be put to use and be profitably sold to consumers. The reader might recall that similar themes emerged when justifying statistical multiplexing in Section 1.1.3. This transformation of computing and IT infrastructure into a utility, which is available to all, somewhat levels the playing field. It forces competition based on ideas rather than computing resources.

Figure 9-2: Compute-intensive applications run in the pools of data-processing capacity, i.e., the Computing Cloud at the Internet core.

At the Internet periphery: mostly document editing and management, visualization, browsing, querying, …

As of early 2010, enterprises are not moving their data massively to the cloud. Many of the current cloud services (Gmail, Google Docs, etc.) do not have much latency in the U.S. Network bandwidth may be a concern, particularly in parts of the world where bandwidth is not fast or cheap (e.g., if the customer incurs per-megabyte costs). As such, there are still good reasons to keep data in a local data center (known as the "private cloud"), and there are more technologies coming that will make this storage better and cheaper. There are many factors to consider, but price is the one that everyone sees on the bottom line. Data storage is not cheap, but it is getting less expensive per megabyte over time. However, no matter how much the $/MB ratio drops, it seems that we are seeing bigger and bigger data storage needs every day.

Security attacks on the cloud could cause major global outages. In principle, a single vulnerability in any part of the various software elements on which a cloud provider bases its services could compromise not just a single application but also the entire virtualized cloud service and all its customers.

9.5 Network Neutrality vs. Tiered Services

Section 1.6.2 briefly reviewed the current debate on *network neutrality* that refers to efforts to keep the Internet open, accessible and "neutral" to all users, application providers and network carriers. In theory, this means, for example, that one carrier would not be allowed to discriminate against an application written by a third party (such as Google Voice) by requiring its users to rely on the carrier's own proprietary voice applications. The public Internet is shared by both businesses and consumers, smart public policy and management of consumer-centric Internet access will profoundly affect business use of the Internet in the future. The U.S. Federal Communications Commission (FCC) is currently (2010) considering rules that would prevent carriers from favoring certain types of content or applications over others or from degrading traffic of Internet companies that offer services similar to those of the carriers.

Who are the biggest players in the Net neutrality debate? On one side, neutrality proponents claim that telecom companies seek to impose a tiered service model in order to control the pipeline and thereby remove competition, create artificial scarcity, and compel subscribers to buy their otherwise uncompetitive services. This bloc includes an array of citizen actions groups loosely aligned with Google and other companies that want to offer new and different uses for the Web but do not generally run networks carrying Internet data. Google communicates mainly on its official blog, where it announced in February 2010 its experimental fiber network.

On the other side of the debate, a group of traditional cable, wireless and telecommunications providers has taken an active role in the debate. Netcompetition.org has posted a list of its members on its e-forum site. The group claims that the Internet is working just fine without any Net neutrality rules.

Enterprises are willing to pay different rates for connectivity based on the quality-of-service needed for business purposes. The cost-benefit model is especially relevant to mobile Internet access because limited radio spectrum precludes unlimited wireless Internet access capacity, even if service costs did not matter. It is also important to point out that not all networks that use IP are part of the Internet. IPTV networks such as AT&T's U-Verse service are isolated from the Internet, and are therefore not subject to network neutrality agreements.

"Best effort" Internet service means that if everyone in the neighborhood streams a film video at the same time, all the neighbors will suffer service interruptions. One option is to build sufficient Internet access capacity for all users to stream uninterrupted video, but many consumers may find that the subscription rate increase needed to pay for such an upgrade is unaffordable. Another option is to enforce policies ensuring that very heavy users do not crowd everyone else out.

The notion of a dichotomy between that network providers like AT&T and application/content providers like Google assumes that the providers should be considered “dumb pipes” whose sole job is to neutrally push traffic from content providers. This may leave the carriers without an incentive to improve network services, and increase capacity and efficiency. The carrier/content-provider dichotomy resonates with the protocol-layering model (Section 1.1.4), in that carriers are assumed to provide layer-1 (link and physical layers) services, whereas content providers provide application-layer services. End-to-end argument places emphasis on the endpoints to best manage their specific needs and assumes that lower layers and intermediary nodes essentially provide “dumb pipes” infrastructure.

It is instructive that in real-world vehicular traffic, congestion problems can be solved by widening highways and building additional lanes. But, more often these problems are solved by introducing policies, such as carpooling lanes, tiered pricing for different roads, and congestion pricing.

There are several interpretations of the network neutrality principle, ranging from absolute non-discrimination to different degrees of allowable discrimination, mainly for providing different quality-of-service. Purist supporters of network neutrality state that lack of neutrality will involve leveraging quality of service to extract remuneration from websites that want to avoid being slowed down. They do not believe that carriers would invest to increase capacity and efficiency. I believe that these arguments are hard to discuss without deep understanding of human psychology and functioning of free markets.

A more technical argument, allegedly based on the end-to-end principle argues that net neutrality means simply that all like Internet content must be treated alike and move at the same speed over the network. Notice that this also requires a list of criteria based on which the content likeness can be decided. Even the proponents of net neutrality admit that that the current Internet is not neutral as, given the space of possible networking applications, the Internet’s best effort generally favors file transfer and other non-time sensitive traffic over real-time communications.

Internet service providers and some networking technology companies argue that providers should have the ability to offer preferential treatment in the form of a tiered services, for example by giving businesses willing to pay an option to transfer their data packets faster than other Internet traffic. The added revenue from such services could be used to pay for the building of increased broadband access to more consumers. Opponents to net neutrality have also argued that a net neutrality law would discourage providers from innovation and competition.

9.6 Summary and Bibliographical Notes

Smart Grid

IEC TC57 has created a family of international standards that can be used as part of the Smart Grid. These standards include IEC61850 which is an architecture for substation automation, and

IEC 61970/61968—the Common Information Model (CIM). The CIM provides for common semantics to be used for turning data into information. Online at: http://tc57.iec.ch/

Office of the National Coordinator for Smart Grid Interoperability, "NIST Framework and Roadmap for Smart Grid Interoperability Standards," Release 1.0 (Draft), National Institute of Standards and Technology (NIST) Draft Publication, September 2009. Online at: http://www.nist.gov/public_affairs/releases/smartgrid_interoperability.pdf

Former IETF Chair and Cisco Fellow Fred Baker has written a document that identifies the core protocols in the IP suite that Smart Grid projects should consider using:

F. Baker, "Core Protocols in the Internet Protocol Suite," Internet-Draft, October 23, 2009. Online at: http://tools.ietf.org/html/draft-baker-ietf-core-04 (Expires: April 26, 2010)

A great deal of useful information on the Smart Grid can be found in the Wikipedia article: http://en.wikipedia.org/wiki/Smart_grid

Cloud Computing

UC Berkeley – "Above the Clouds: A Berkeley View of Cloud Computing," 2009. Online at: http://berkeleyclouds.blogspot.com/

Published by the UC Berkeley Reliable Adaptive Distributed Systems Laboratory (a.k.a. RAD Lab), this report is an excellent overview of the move to cloud computing. It identifies some key trends, addresses the top obstacles to cloud use, and makes some excellent points about cloud economics.

Network Neutrality

A great source of information on network neutrality is the Wikipedia article: http://en.wikipedia.org/wiki/Net_neutrality.

Programming Assignments

The following assignments are designed to illustrate how a simple model can allow studying individual aspects of a complex system. In this case, we study the congestion control in TCP.

The assignments are based on the reference example software available at this book's web site (given in Preface); follow the link "Team Projects." This software implements a simple TCP simulator for Example 2.1 (Section 2.2) in the Java programming language. Only the Tahoe version of the TCP sender is implemented. This software is given only as a reference, to show how to build a simple TCP simulator. You can take it as is and only modify or extend the parts that are required for your programming assignments. Alternatively, you can write your own software anew, using the programming language of your choosing. Instead of Java, you can use C, C++, C#, Visual Basic, or another programming language.

The action sequence in Figure P-1 illustrates how the simulator works. It consists of four Java objects, of which Simulator is the main class that orchestrates the work of others. It repeatedly cycles around visiting in turn Sender, Router, and Receiver. Simulator passes two arrays to each (segments and acknowledgements array) and gets them back after each object does its work.

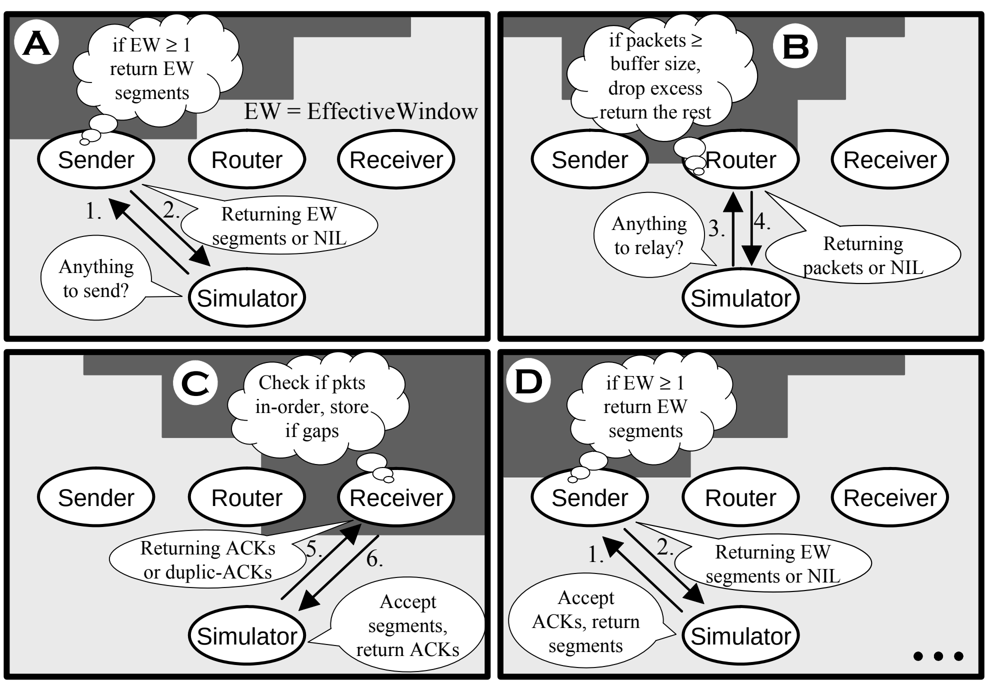

Figure P-1: Action sequence illustrating how the Simulator object orchestrates the work of other software objects (shown are only first four steps of a simulation).

Project Report Preparation

When you get your program working, run it and plot the relevant charts similar to those provided for Example 2.1. Calculate the sender utilization, where applicable, and provide explanations and comments on the system performance. Also, calculate the latency for transmitting a 1 MB file.

Each chart/table should have a caption and the results should be discussed. Explain any results that you find non-obvious or surprising. Use manually drawn diagrams (using a graphics program such as PowerPoint), similar to Figure 2-14 and Figure 2-15, where necessary to support your arguments and explain the detailed behavior.

Assignment 1: TCP Reno

Implement the Reno version of the TCP sender which simulates Example 2.1. You can use the Java classes from the TCP-Tahoe example, in which case you only need to extend from `TCPSender` and implement `TCPSenderReno`, fashioned after the existing `TCPSenderTahoe`. Alternatively, you can implement everything anew. Use RFC 2581 as the primary reference, to be found here http://www.apps.ietf.org/rfc/rfc2581.html.

(a) Use the same network parameters as in Example 2.1.

(b) Modify the router to randomly drop up to one packet during every transmission round, as follows. The router should draw a random number from a uniform distribution between 0 and `bufferSize`, which in our case equals to 7. Use this number as the index of the segment to delete from the `segments` array. (Note that the given index may point to a `null` element if the array is not filled up with segments, in which case do nothing.)

1. Compare the sender utilization for the TCP Reno sender with that of a Tahoe sender (given in Figure 2-16). Explain any difference that you may observe.
2. Compare the sender utilization for case (b) with random dropping of packets at the router.
3. Show the detail of slow-start and additive increase phases diagrams. Compare them to Figure 2-14 and Figure 2-15. Explain any differences that you may find.

Include these findings in your project report, which should be prepared as described for all projects at the beginning of this chapter.

Assignment 2: TCP Tahoe with Bandwidth Bottleneck

Consider the network configuration as in the reference example, but with the router buffer size set to a large number, say 10,000 packets. Assume that RTT = 0.5 s and that at the end of every RTT period the sender receives a cumulative ACK for all segments *relayed by the router* during that period. Also, set `RcvWindow` = 1 MByte instead of the default value of 65535 bytes.

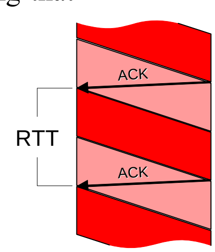

Due to the large router buffer size there will be no packet loss, but the bandwidth mismatch between the router's input and output lines still remains. Because of this, the router may not manage to relay all the packets from the current round before the arrival of the packets from the subsequent round. This behavior is illustrated in Figure P-2 (also see the inset figure on the right, which illustrates how the pattern repeats for each RTT). The

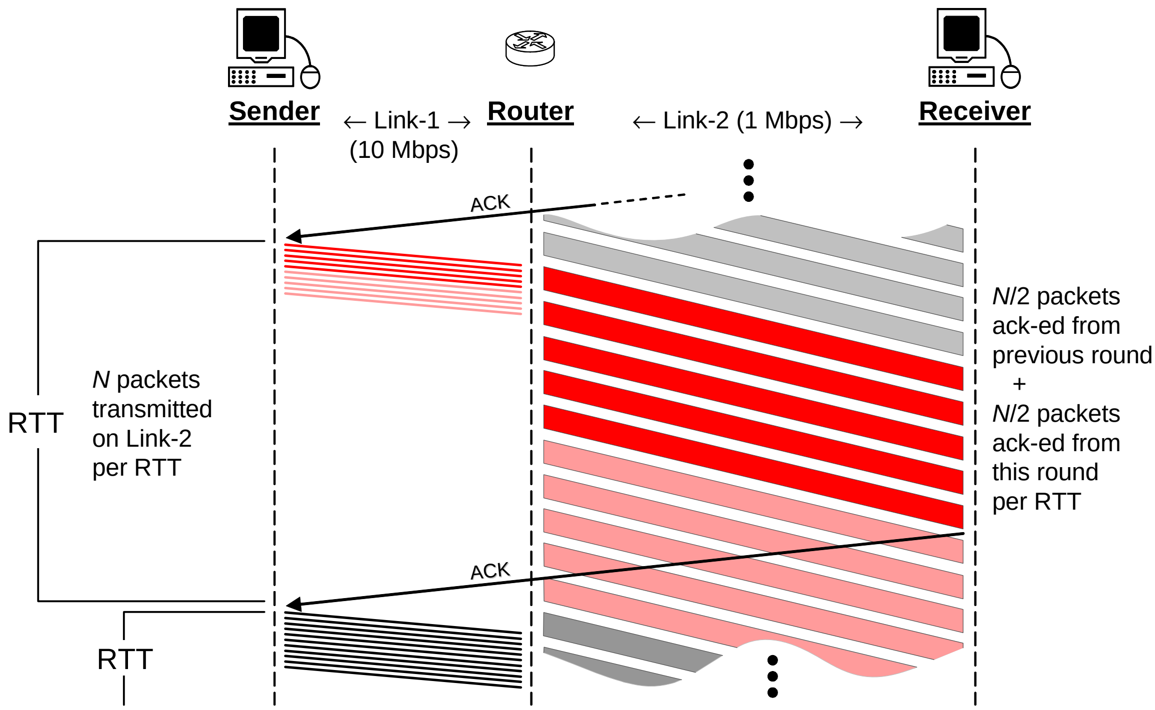

Figure P-2: This time diagram for Assignment 2 illustrates how the number of packets transmitted in one round is limited by the bandwidth of Link-2.

remaining packets are carried over to the next round and accumulate in the router's queue. As the sender's congestion window size grows, there will be a queue buildup at the router. There will be no loss because of the large buffer size, but packets will experience delays. The packets carried over from a previous round will be sent first, before the newly arrived packets. Thus, the delays still may not trigger the RTO timer at the sender because the packets may clear out before the timeout time expires.

The key code modifications are to the router code, which must be able to carry over the packets that were not transmitted within one RTT. In addition, you need to modify `TCPSimulator.java` to increase the size of the arrays `segments` and `acks`, because these currently hold only up to 100 elements.

1. Determine the average queuing delay per packet once the system stabilizes. Explain why buffer occupancy will never reach its total capacity. Are there any retransmissions (quantify, how many) due to large delays, although packets are never lost. Use manually drawn diagrams to support your arguments.

2. In addition to the regular charts, plot the two charts shown in the following figure:

   The chart on the left should show the number of the packets that remain buffered in the router at the end of each transmission round, which is why the time axis is shown in RTT units. (Assume the original `TimeoutInterval` = 3×RTT = 3 sec.)

   To generate the chart on the right, make the variable `TimeoutInterval` an input parameter to your program. Then input different values of `TimeoutInterval` and measure the corresponding utilization of the TCP sender. Of course, RTT remains constant at 1 sec. (Notice the logarithmic scale on the horizontal axis. Also, although it may appear

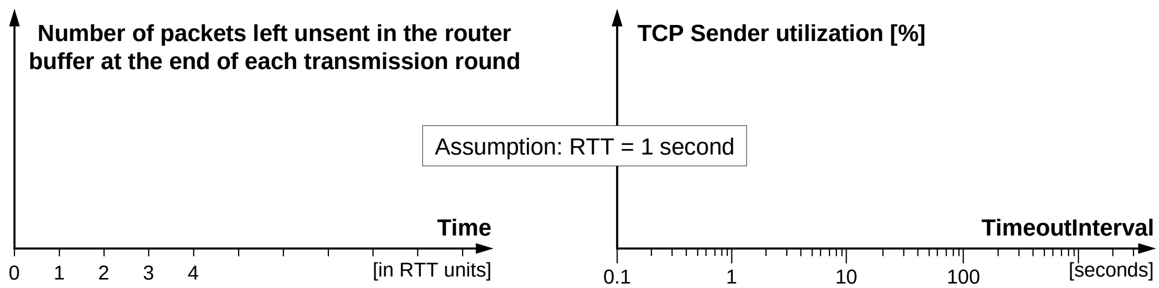

Figure P-3: Results charts for Assignment 2.

strange to set `TimeoutInterval` smaller than RTT, this illustrates the scenario where RTT may be unknown, at least initially.) Provide explanation for any surprising observations or anomalies.

A more ambitious team should consider the following problem. While a long queue is less likely to overflow during a traffic burst (thus reducing packet loss probability), it potentially increases the queuing delay for non-dropped packets. A short queue reduces this delay, but conversely increases the probability of packet loss for bursty traffic. Experiment by adjusting the router buffer size and study the tradeoff between packet loss and queuing delay.

Prepare the project report as described for all projects at the beginning of this chapter.

Assignment 3: TCP Tahoe with More Realistic Time Simulation and Packet Reordering

In the reference example implementation, the packet sending times are clocked to the integer multiples of RTT (see Figure 2-10). For example, Packets #2 and 3 are sent together at the time = 1 × RTT; packets #4, 5, 6, and 7 are sent together at the time = 2 × RTT; and so on. Obviously, this does not reflect the reality, as illustrated in Figure 2-6. For example, in the current implementation, duplicate ACKs are generated only because of packets dropped in the router, never because packets are reordered. Your simulation will show what happens when packets are reordered (rather than only dropped).

A simple way to implement this modification is explained next. The key change in your code should be in the class `Router.java`, as follows. For every newly received packet, the router assigns a random amount of delay (integer number ≥0). For every invocation of the method `Router.relay()`, router decrements all delays by one and returns only the packets with the delay equal zero. To maintain the record of delays for individual packets, we could modify `TCPSegment.java` and add a new field (`public int delay = 0;`). This field will be used *only* by `Router.java` and will be *ignored* by all other classes.

The method `Router.relay()` has an argument array `packets_[]` where the packets to be relayed are received. This is also where the packets are returned that are let pass through the router and onwards to TCP Receiver. Add a new array `RouterBuffer[]` to the Router, to store the packets that are currently delayed, waiting for their delay counter to reach down to zero. Make the size of this array an input argument to the simulation, to allow also for packets dropped because of buffer overflow. Initialize `RouterBuffer[]` with a nil pointer. The method `Router.relay()` should execute the following steps (remove the old code of this method):

• **STEP 1:** Iterate through the array `RouterBuffer[]` until you find a nil element (the end of packets stored in the buffer). For each `RouterBuffer[i]`, decrement the delay value of the stored TCP segment by "1". Some of the delays may become zero after decremented. Leave them as is, they will be handled in STEP 3 below.

• **STEP 2:** Append all non-nil elements from the argument array `packets_[]` at the end of `RouterBuffer[]`, after the existing packets in `RouterBuffer[]`, if any. Set each `packets_[i]` to nil after you move it to `RouterBuffer[]`. If during this process you run out of space in `RouterBuffer[]`, all remaining packets from `packets_[]` should be dropped. For each packet moved to `RouterBuffer[]`, generate an integer random number with exponential distribution and a small mean value, say 0, 1 or 2. (Make it possible for the user to enter the parameters of the exponential distribution from the command line, similarly to entering the number of iterations to run the simulation). Set the `delay` field of `RouterBuffer[i]` to the generated delay value. (The delay value must be integer ≥ 0.)

• **STEP 3:** Iterate through the array `RouterBuffer[]` until you find a nil element (the end of the stored packets.) If `RouterBuffer[i].delay` equals zero (i.e., zero delay value), then move the packet `RouterBuffer[i]` to the method argument array `packets_[]`. You will need to shift the remaining elements of `RouterBuffer[]` to remove the gap. At the end of `RouterBuffer[]` make sure to put nil pointers to indicate the of the stored packets. After this step, there should be no element left in `RouterBuffer[]` with `delay` value equal zero.

• **STEP 4:** Return from the method `Router.relay()`. The argument array `packets_[]` contains the packets that the router has let pass. This array should be passed on to the Receiver, as is in the current simulator code.

Notice that if you run your simulation say for 100 iterations, there may at the end still remain some packets in the router in `RouterBuffer[]`. You will need to flush the router buffer by invoking `Router.relay()` with the argument array `packets_[]` having all elements set to nil and then passing the returned array `packets_[]` to the TCP Receiver, until no packets remain in `RouterBuffer[]`.

Notice that, unlike Example 2.1 where a segment can arrive at the receiver out-of-sequence only because a previous segment was dropped at the router, in your assignment an additional reason for out-of-sequence segments is that different segments can experience different amounts of delay. Each packet is assigned the delay value *individually* as generated by the random number generator for each packet. So if, say, 3 packets arrive in the input argument array `packets_[]`, then it is possible that packet #1 gets delay 4, so it will have to sit inside the router buffer through four invocations of method `Router.relay()`. On the other hand, packet #2 that arrived in the same iteration could get assigned delay 0 and leave at the end of this method invocation, and packet #3 could get assigned delay value 1 and leave in the next invocation, but still before packet #1. Recall that for every out-of-order segment, the receiver reacts immediately and sends a dupACK (see Figure 2-9). Your simulation will show what happens when packets are reordered. By controlling the size of the `RouterBuffer[]` array, you may cause packets dropped because of the router buffer overflow.

Print some statistics from your new router code for every iteration, such as how many packets are currently left in `RouterBuffer[]` before the method `Router.relay()` is exited, the

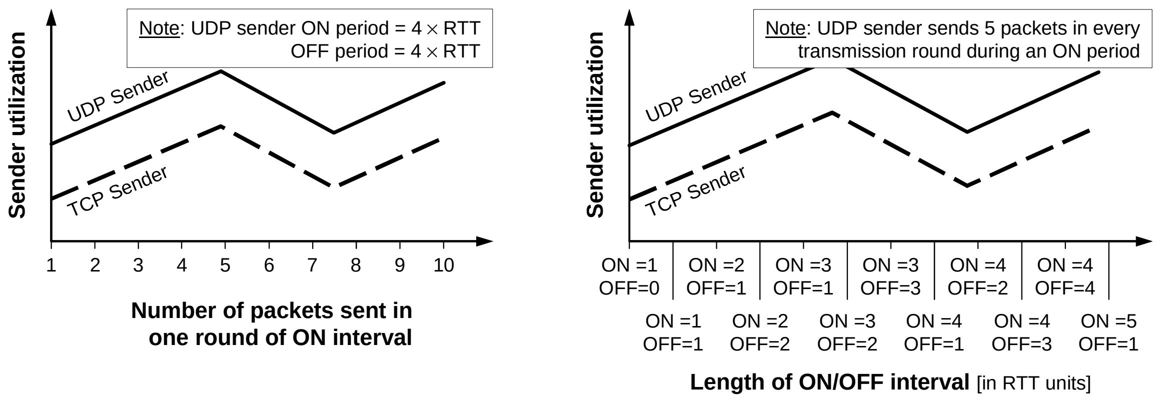

Figure P-4: Results charts for Assignment 4.

histogram of delay values for packets currently left in `RouterBuffer[]`, and how many packets are dropped because of the router buffer overflow.

Prepare the project report as described for all projects at the beginning of this chapter. Average over multiple runs to obtain average sender utilization.

Assignment 4: TCP Tahoe with a Concurrent UDP Flow

In the reference example implementation, there is a single flow of packets, from the sender, via the router, to the receiver. Your task is to add an additional, UDP flow of packets that competes with the TCP flow for the router resources (i.e., the buffering memory space). Modify the router Java class so that it can accept simultaneously input from a TCP sender and an UDP sender, and it should correctly deliver the packets to the respective TCP receiver and UDP receiver. The UDP sender should send packets in an ON-OFF manner. First, the UDP sender enters an ON period for the first four RTT intervals and it sends five packets at every RTT interval. Then the UDP sender enters an OFF period and becomes silent for four RTT intervals. This ON-OFF pattern of activity should be repeated for the duration of the simulation. At the same time, the TCP Tahoe sender is sending a very large file via the same router.

1. In addition to the TCP-related charts, plot also the charts showing the statistics of the packet loss for the UDP flow.
2. How many iterations takes the TCP sender to complete the transmission of a 1 MByte file? (Because randomness is involved, you will need to average over multiple runs.)
3. Perform the experiment of varying the UDP sender regime as shown in Figure P-4. In the diagram on the left, the UDP sender keeps the ON/OFF period duration unchanged and varies the number of packets sent per transmission round. In the diagram on the right in Figure P-4, the UDP sender sends at a constant rate of 5 packets per transmission round, but varies the length of ON/OFF intervals.
4. Based on these two experiments, can you speculate how increasing the load of the competing UDP flow affects the TCP performance? Is the effect linear or non-linear? Can you explain your observation?

Prepare the project report as described for all projects at the beginning of this chapter.

Assignment 5: Competing TCP Tahoe Senders

Suppose that you have a scenario where two TCP Tahoe senders send data segments via the same router to their corresponding receivers. In case the total number of packets arriving from both senders exceeds the router's buffering capacity, the router should discard all the *excess packets* as follows. Discard the packets that are the tail of a group of arrived packets. The number of packets discarded from each flow should be (approximately) proportional to the total number of packets that arrived from the respective flow. That is, if more packets arrive from one sender then proportionally more of its packets will be discarded, and vice versa.

Assume that the second sender starts sending with a delay of three RTT periods after the first sender. Plot the relevant charts for both TCP flows and explain any differences or similarities in the corresponding charts for the two flows. Calculate the total utilization of the router's output line and compare it with the throughputs achieved by individual TCP sessions. *Note*: to test your code, you should swap the start times of the two senders so that now the first sender starts sending with a delay of three RTT periods after the second sender.

In addition to the above charts, perform the experiment of varying the relative delay in transmission start between the two senders. Plot the utilization chart for the two senders as shown in the figure. Are the utilization curves for the two senders different? Provide an explanation for your answer. *Note*: Remember that for fair comparison you should increase the number of iterations by the amount of delay for the delayed flow.

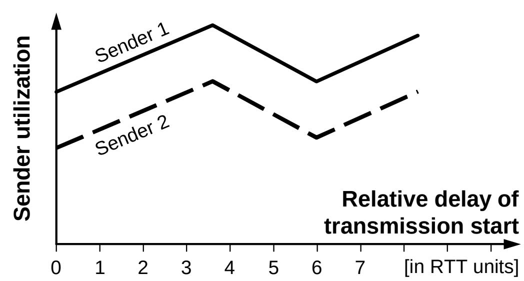

Prepare the project report as described for all projects at the beginning of this chapter.

Assignment 6: Random Loss and Early Congestion Notification

This assignment is intended to simulate Random Early Detection (RED), described in Section 5.3.1. The network configuration is the same as in Example 2.1 with the only difference being in the router's behavior. Because we are dealing with a very small buffer size (i.e., 6), the granularity of TCP bursts is relatively high compared to the buffer size. As a result, it may not make much difference if we worked with the average rather than instantaneous queue length. In our simple approximation of RED, we will consider only the instantaneous queue length when deciding about dropping a packet.

(a) First, assume that, in addition to discarding the packets that exceed the buffering capacity, the router also discards packets randomly. For example, suppose that 14 packets arrive at the router in a given transmission round. Then *in addition* to discarding all packets in excess of 6+1=7 as in the reference example, the router *also* discards some of the six packets in the buffer. (The packet currently in service is *never* considered for being dropped.) For each packet currently in the buffer, the router draws a random number from a normal distribution, with the mean equal to zero and adjustable standard deviation. If the absolute value of the random number exceeds a given threshold, then the corresponding packet is dropped. Otherwise, it is forwarded.

(b) Second, assume that the router considers for discarding only the packets that are located within a certain zone in the buffer. For example, assume that the random-drop zone starts at 2/3 of the total buffer space and runs up to the end of the buffer. Then perform the above dropping procedure only on the packets that are located between 2/3 of the total buffer space and the end of the buffer. (Packets that arrive at a full buffer are automatically dropped!)

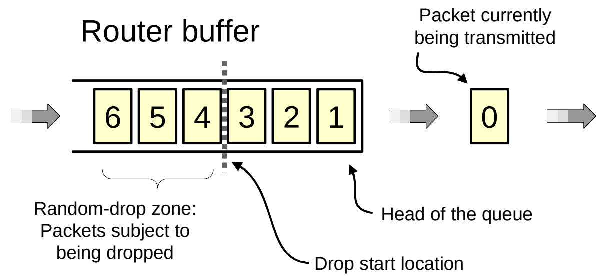

Your program should allow entering different values of parameters for running the simulation, such as: variance of the normal distribution and the threshold for random dropping the packets in (a); and, the start discarding location in (b).

1. In addition to the regular charts, plot the three-dimensional chart shown in the figure. (Use MatLab or a similar tool to draw the 3D graphics.) Because the router drops packets *randomly*, you should repeat the experiment several times (minimum 10) and plot the *average utilization* of the TCP sender.

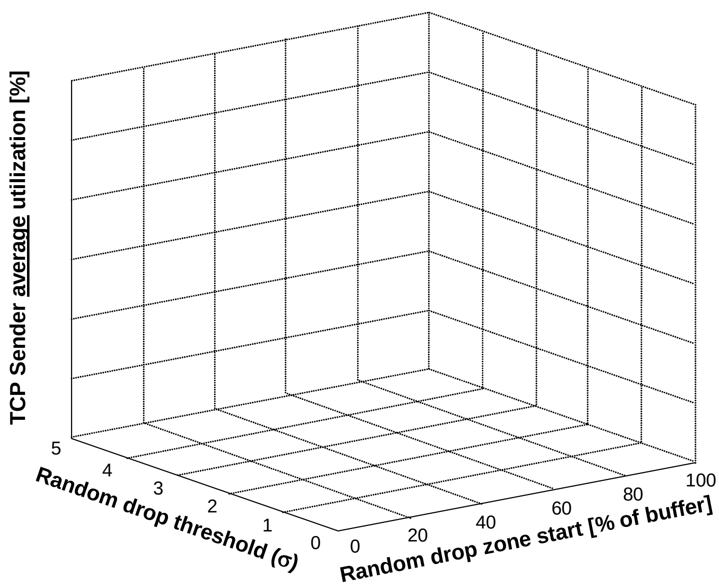

2. Find the regions of maximum and minimum utilization and indicate the corresponding points/regions on the chart. *Explain your findings*: why the system exhibits higher/ lower utilization with certain parameters?

3. You should also present different two-dimensional cross-sections of the 3D graph, if this can help illuminate your discussion.

A more ambitious team should try implementing a more accurate approximation of RED (Section 5.3.1). Prepare the project report as described for all projects at the beginning of this chapter.

Solutions to Selected Problems

Problem 1.1 — Solution

Let t_x denote the transmission time for a packet L bits long and let t_p denote the propagation delay on each link. Each packet crosses three links (Link-1: source to router1; Link-2: router1 to router2; Link-3: router2 to destination).

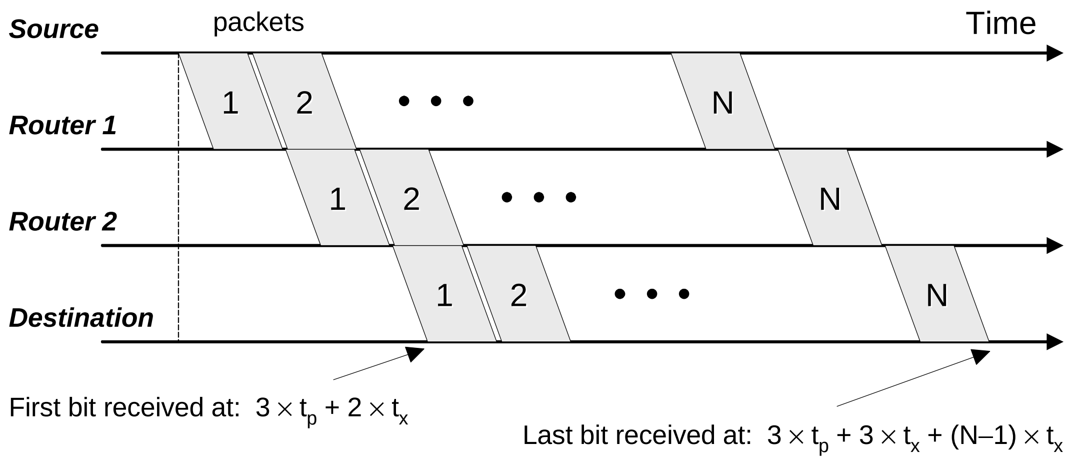

(a)

The total transmission time equals: $3 \times t_p + 3 \times t_x + (N - 1) \times t_x = 3 \times t_p + (N + 2) \times t_x$.

(b)

The transmission time is ½ t_x, the propagation time remains the same, and there are twice more packets, so the total transmission time equals:

$3 \times t_p + (2\times N + 2) \times (t_x/2) = 3 \times t_p + (N + 1) \times t_x$.

(c)

The total delay is smaller in case (b) by t_x because of greater parallelism in transmitting shorter packets. If we use, for example, four times shorter packets (L/4), then the total transmission time equals: $3 \times t_p + (N + 1/2) \times t_x$.

On the other hand, if we use two times longer packets, i.e., N/2 packets each $2\times L$ bits long, then the total transmission time equals:

$3 \times t_p + (N/2 + 2) \times (2\times t_x) = 3 \times t_p + (N + 4) \times t_x$.

which is longer by $2\times t_x$ then in case (a).

Problem 1.2 — Solution

Problem 1.3 — Solution

(a)

Having only a single path ensures that all packets will arrive in order, although some may be lost or damaged due to the non-ideal channel. Assume that *A* sends a packet with SN = 1 to *B* and the packet is lost. Because *A* will not receive ACK within the timeout time, it will retransmit the packet using the same sequence number, SN = 1. Because *B* already received a packet with SN = 1 and it is expecting a packet with SN = 0, it concludes that this is a duplicate packet.

(b)

If there are several alternative paths, the packets can arrive out of order. There are many possible cases where *B* receives duplicate packets and cannot distinguish them. Two scenarios are shown below, where either the retransmitted packet or the original one gets delayed, e.g., by taking a longer path. These counterexamples demonstrate that the alternating-bit protocol cannot work over a general network.

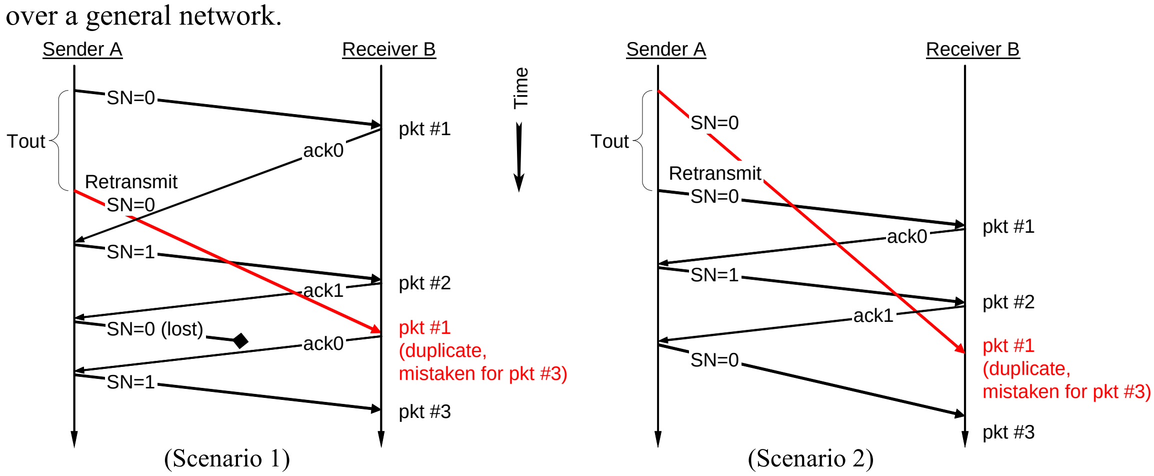

Problem 1.4 — Solution

Problem 1.5 — Solution

Problem 1.6 — Solution

Recall that the utilization of a sender is defined as the fraction of time the sender is actually busy sending bits into the channel. Because we assume errorless communication, the sender is maximally used when it is sending without taking a break to wait for an acknowledgement. This happens if the first packet of the window is acknowledged before the transmission of the last packet in the window is completed. That is,

$$(N-1) \times t_x \geq \text{RTT} = 2 \times t_p$$

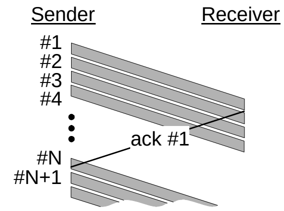

where t_x is the packet transmission delay and t_p is the propagation delay. The left side represents the transmission delay for the remaining $(N-1)$ packets of the window, after the first packet is sent. Hence, $N \geq \left\lceil \frac{2 \times t_p}{t_x} \right\rceil + 1$. The ceiling operation $\lceil \cdot \rceil$ ensures integer number of packets.

In our case, ℓ = 10 km, $v \approx 2 \times 10^8$ m/s, R = 1 Gbps, and L = 512 bytes.

Hence, the packet transmission delay is: $t_x = \frac{L}{R} = \frac{512 \times 8 \text{ (bits)}}{1 \times 10^9 \text{ (bits/sec)}} = 4.096\ \mu\text{s}$

The propagation delay is: $t_p = \frac{\ell}{v} = \frac{10000 \text{ m}}{2 \times 10^8 \text{ m/s}} = 50\ \mu\text{s}$

Finally, $N \geq \lceil 24.41 \rceil + 1 = 26$ packets.

Problem 1.7 — Solution

str. 284

Problem 1.8 — Solution

The solution is shown in the following figure. We are assuming that the retransmission timer is set appropriately, so ack0 is received before the timeout time expires. Notice that host A simply ignores the duplicate acknowledgements of packets 0 and 3, i.e., ack0 and ack3.

There is no need to send source-to-destination acknowledgements (from C to A) in this particular example, because both $\overline{AB}$ and $\overline{BC}$ links are reliable ***and*** there are no alternative paths from A to C but via B. The reader should convince themselves that should alternative routes exist, e.g., via another host D, then we would need source-to-destination acknowledgements in addition to (or instead of) the acknowledgements on individual links.

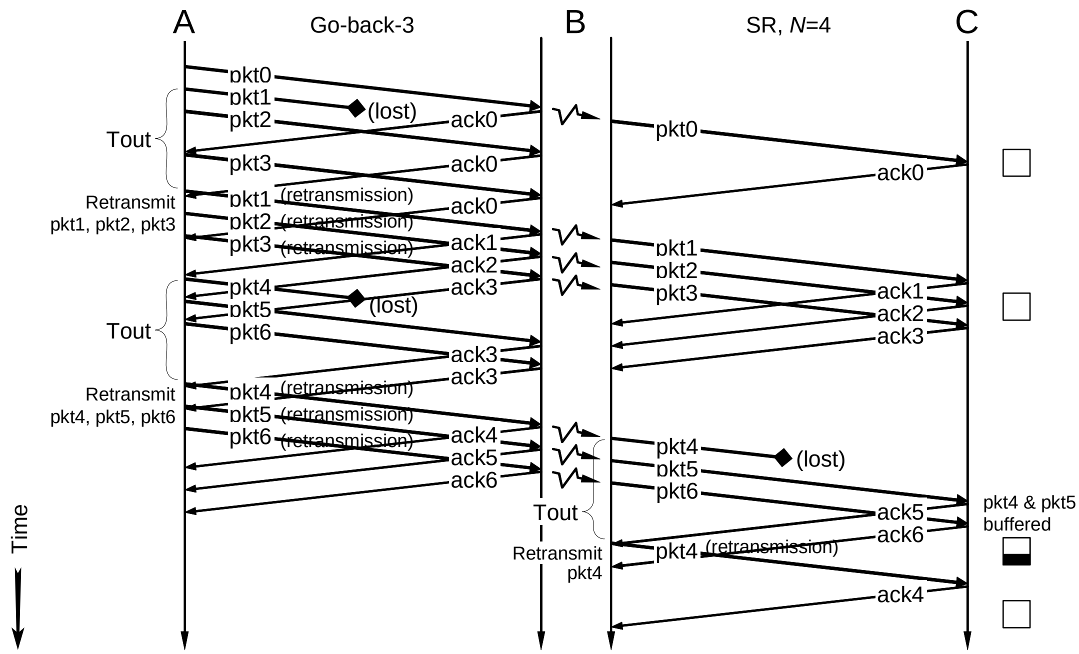

Problem 1.9 — Solution

The solution is shown in the following figure.

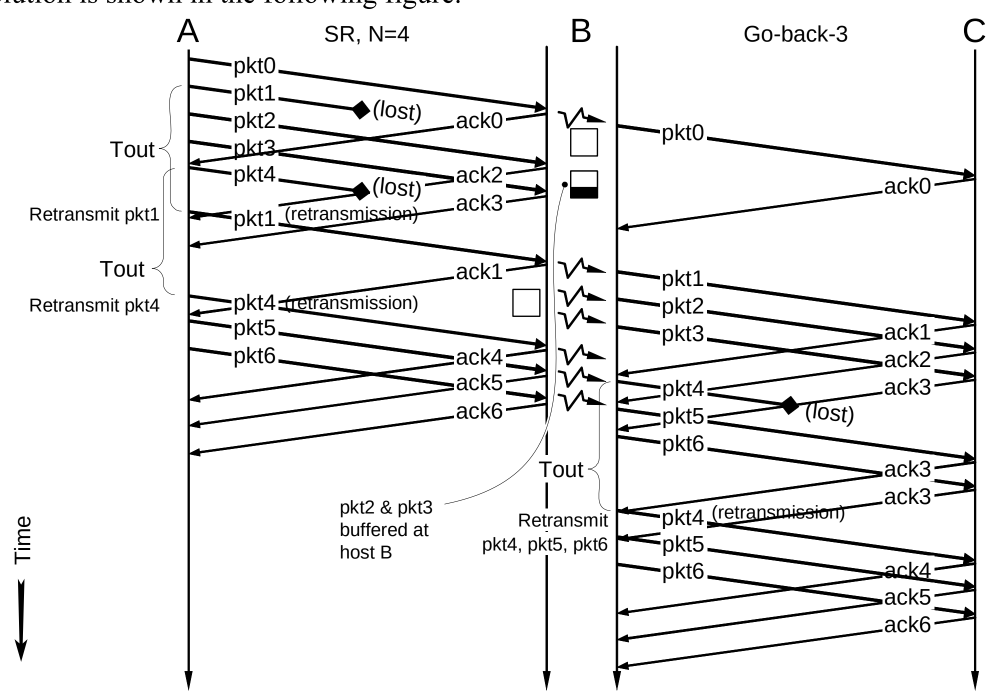

Problem 1.10 — Solution

It is easy to get tricked into believing that the second, (b), configuration would offer better performance, because the router can send in parallel in both directions. However, this is not true, as will be seen below.

(a)

Propagation delay = 300 m / $(2 \times 10^{8}) = 1.5 \times 10^{-6}$ s = 1.5 μs

Transmission delay per data packet = $2048 \times 8 / 10^{6} = 16.384 \times 10^{-3}$ = 16.384 ms

Transmission delay per ACK = $10^{8} / 10^{6} = 0.08 \times 10^{-3}$ = 0.08 ms

Transmission delay for $N = 5$ (window size) packets = $16.384 \times 5 < 82$

$82 + 0.08 + 0.0015 \times 2 = 82.083$

Subtotal time for 100 packets in one direction = $100 \times 82.083 / 5$ = 1641.66 ms

Total time for two ways = 1641.66×2 = 3283.32 ms

(b)

If host A (or B) sends packets after host B (or A) finishes sending, then the situation is similar to (a) and the total time is about 3283.32 ms.

If hosts A and B send a packet each simultaneously, the packets will be buffered in the router and then forwarded. The time needed is roughly **double** (!), as shown in this figure.

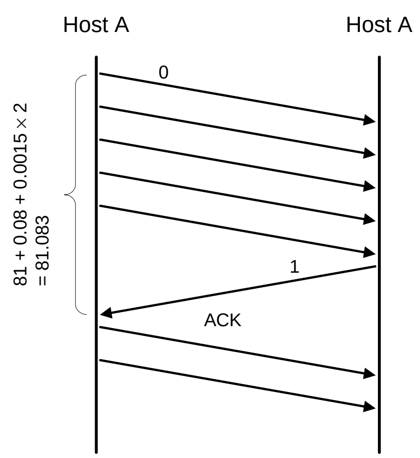

(a)

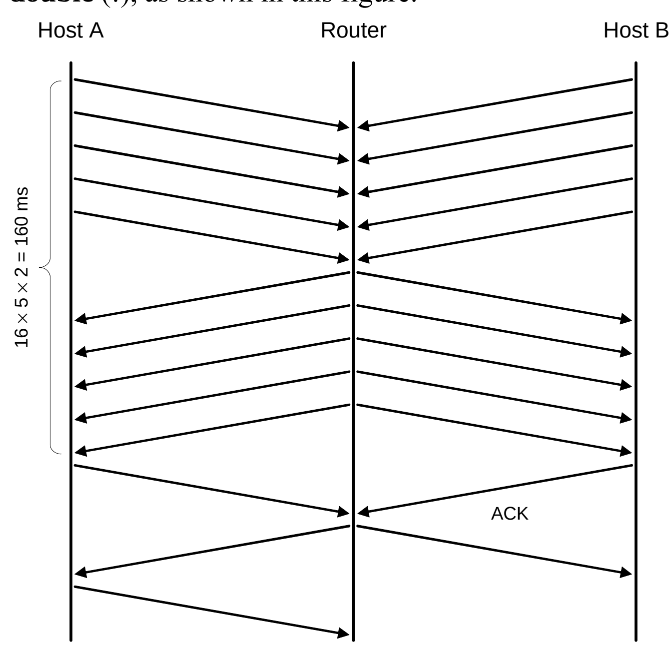

(b)

Problem 1.11 — Solution

Go-back-*N* ARQ.

Packet error probability = p_e for data packets; ACKs error free. Successful receipt of a given packet with the sequence number *k* requires successful receipt of all previous packets in the sliding window. In the worst case, retransmission of frame *k* is always due to corruption of the earliest frame appearing in its sliding window. So, $p_{\text{succ}} = \prod_{i=k-N}^{k} (1 - p_e) = (1 - p_e)^N$, where *N* is the sliding window size. An upper bound estimate of $E\{n\}$ can be obtained easily using Eq. (1.8) as $E\{n\} = \frac{1}{p_{\text{succ}}} = \frac{1}{(1 - p_e)^N}$. On average, however, retransmission of frame *k* will be due to an error in frame (*k*–LAR)/2, where LAR denotes the sequence number of the *Last Acknowledgement Received*.

(a)

Successful transmission of one packet takes a total of	$t_{\text{succ}} = t_x + 2 \times t_p$
The probability of a failed transmission in one round is	$p_{\text{fail}} = 1 - p_{\text{succ}} = 1 - (1 - p_e)^N$
Every failed packet transmission takes a total of	$t_{\text{fail}} = t_x + t_{\text{out}}$

(assuming that the remaining *N*−1 packets in the window will be transmitted before the timeout occurs for the first packet).

Then, using Eq. (1.9) the expected (average) total time per packet transmission is:

$$E\{T_{\text{total}}\} = t_{\text{succ}} + \frac{p_{\text{fail}}}{1 - p_{\text{fail}}} \cdot t_{\text{fail}}$$

(b)

If the sender operates at the maximum utilization (see the solution of Problem 1.6), then the sender waits for *N*−1 packet transmissions for an acknowledgement, $t_{\text{out}} = (N-1) \cdot t_x$, before a packet is retransmitted. Hence, the expected (average) time per packet transmission is:

$$E\{T_{\text{total}}\} = t_{\text{succ}} + \frac{p_{\text{fail}}}{1 - p_{\text{fail}}} \cdot t_{\text{fail}} = 2 \cdot t_p + t_x \cdot \left(1 + (N-1) \cdot \frac{p_{\text{fail}}}{1 - p_{\text{fail}}} \right)$$

Problem 1.12 — Solution

(a)

Packet transmission delay equals = 1024 × 8 / 64000 = 0.128 sec; acknowledgement transmission delay is assumed to be negligible. Therefore, the throughput S_{MAX} = 1 packet per second (pps).

(b)

To evaluate $E\{S\}$, first determine how many times a given packet must be (re-)transmitted for successful receipt, $E\{n\}$. According to Eq. (1.8), $E\{n\} = 1/p \cong 1.053$. Then, the expected throughput is $E\{S\} = \dfrac{S_{\text{MAX}}}{E\{n\}} = 0.95$ pps.

(c)

The fully utilized sender sends 64Kbps, so $S_{\text{MAX}} = \dfrac{64000}{1024 \times 8} = 7.8125$ pps.

(d)

Again, we first determine how many times a given packet must be (re-)transmitted for successful receipt, $E\{n\}$. The sliding window size can be determined as $N = 8$ (see the solution for Problem 1.6). A lower bound estimate of $E\{S\}$ can be obtained easily by recognizing that $E\{n\} \leq 1/p^8 \cong 1.5$ (see the solution for Problem 1.11). Then, $S_{\text{MAX}} \times p^8 \cong (7.8125) \times (0.6634) \cong 5.183$ pps represents a non-trivial lower bound estimate of $E\{S\}$.

Problem 1.13 — Solution

The packet size equals to transmission rate times the slot duration, which is 1500 bps × 0.08333 s = 125 bits. This wireless channel can transmit a maximum of 12 packets per second, assuming that in each slot one packet is transmitted. (Recall that slotted ALOHA achieves maximum throughput when $G = 1$ packet per slot). Of these, some will end up in collision and the effective throughput will be equal to 0.368 × 12 = 4.416 packets per second. Because this is aggregated over 10 stations, each station can effectively transmit 4.416 / 10 = 0.4416 packets per second, or approximately 26 packets per minute, at best.

Problem 1.14 — Solution

(a) Given the channel transmission attempt rate G, the probability of success was derived in Eq. (1.12) as e^{-G}, which is the probability that no other station will transmit during the vulnerable period.

(b) The probability of exactly K collisions and then a success is

$$(1 - e^{-G})^k \cdot e^{-G}$$

which is derived in the same manner as Eq. (1.7) in Section 1.3.1.

Problem 1.15 — Solution

(a)

The solution is given in Figure 1-26(b): there will be $G \cdot P_0$ successfully transmitted packets, and $G \cdot (1-P_0)$ collided packets. Slotted ALOHA operates under maximum efficiency when G=1, that is when on average one packet is transmitted per each slot. This means that, on average, there will be no idle slots—every slot, on average, will be used for a transmission. Of those transmissions, there will be on average $\lambda = 1/e$ fresh packet arrivals as well as $1/e$ slots with successful transmissions (some of the successful transmissions may be retransmissions of backlogged packets), and there will be on average $(1 - 1/e)$ slots with collisions.

(b)

A slotted ALOHA system would operate with a less-than-maximum efficiency if $G \neq 1$. When $G < 1$ there are on average less than one transmission attempts per packet time, so we say that the system is *underloaded*. Conversely, when $G > 1$ there are on average more than one transmission attempts per packet time, so we say that the system is *overloaded*. In both cases, the arrival rate will be $\lambda < 1/e$ (see Figure 1-28).

The system would be *underloaded* when some stations do not have packets in some slots to transmit or retransmit. As a result, there will be a non-zero fraction of idle slots. In the underloaded case, there will be on average $1/e - \lambda > 0$ idle slots, λ successful slots, and the remaining $\leq(1 - 1/e)$ slots with collisions.

The system would be *overloaded* if the arrival rate became temporarily $\lambda > 1/e$, i.e., greater than the optimal value. This will result in many collisions and there will be many stations that become backlogged, trying to retransmit previously collided packets. The backlogged stations will not accept fresh packets, which effectively means that the arrival rate will drop to $\lambda < 1/e$. (Notice that λ becomes small not because users are not generating new packets but because many stations are backlogged and do not accept fresh packets). In the overloaded case, there will be on average no idle slots, $\lambda < 1/e$ successful slots, and the remaining $>(1 - 1/e)$ slots with collisions.

(c)

Both the maximum-efficiency and underloaded cases are stable. In both cases for a steady arrival rate λ, the system will remain in a stable condition.

The overloaded case is *unstable*: initially the arrival rate will be $\lambda > 1/e$, but then many stations will become backlogged and will not accept new packet arrivals so λ will drop to $\lambda < 1/e$. These backlogged stations will keep retransmitting their packets and eventually the backlogged packets will clear. If after this clearing the arrival rate remains $\lambda \leq 1/e$, then the system will remain maximally efficient ($\lambda = 1/e$) or it will remain underloaded ($\lambda < 1/e$). Otherwise, it will again become temporarily overloaded.

Problem 1.16 — Solution

Problem 1.17 — Solution

The stations will collide only if both select the same backoff value, i.e., either both select 0 or both select 1. The "tree diagram" of possible outcomes is shown in the figure below. To obtain the probabilities of different outcomes, we just need to multiply the probabilities along the path to each outcome.

(a)

Probability of transmission $p = ½$, the success happens if either station transmits alone:

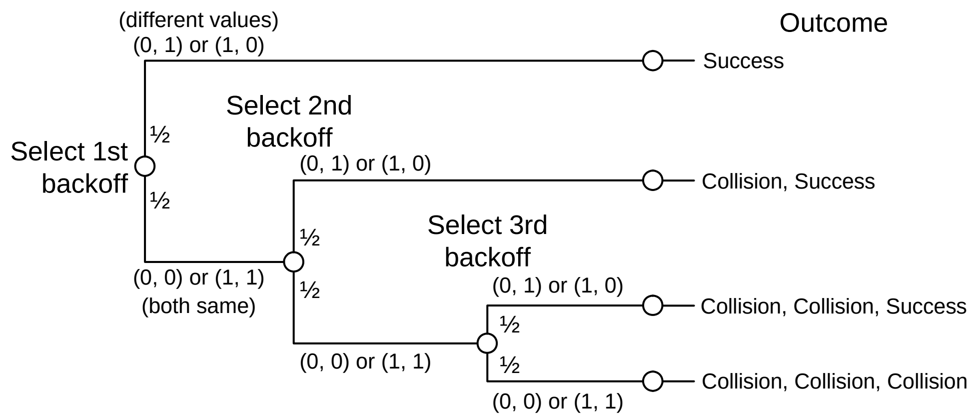

$$P_{\text{success}} = \binom{2}{1} \cdot (1-p)^{2-1} \cdot p = 2 \times ½ \times ½ = 0.5$$

(b)

The first transmission ends in collision if both stations transmit simultaneously

$$P_{\text{collision}} = 1 - P_{\text{success}} = 0.5$$

Because the waiting times are selected *independent* of the number of previous collisions (i.e., the successive events are independent of each other), the probability of contention ending on the second round of retransmissions is

$$P_{\text{collision}} \times P_{\text{success}} = 0.5 \times 0.5 = 0.25$$

(c)

Similarly, the probability of contention ending on the third round of retransmissions is

$$P_{\text{collision}} \times P_{\text{success}} \times P_{\text{success}} = 0.5 \times 0.5 \times 0.5 = 0.125$$

(d)

Regular nonpersistent CSMA with the normal binary exponential backoff algorithm works in such a way that if the channel is busy, the station selects a random period to wait. When the waiting period expires, it senses the channel. If idle, transmit; if busy, wait again for a random period, selected from the *same range*. And so on until the channel becomes idle and transmission occurs.

If the transmission is successful, the station goes back to wait for a new packet.

If the transmission ends in collision, *double the range* from which the waiting period is drawn and repeat the above procedure.

Unlike the regular nonpersistent CSMA, in the modified nonpersistent CSMA there are only two choices for waiting time. Therefore, half of the backlogged stations will choose one way and the other half will choose the other way.

In conclusion, regular nonpersistent CSMA performs better than the modified one under heavy loads and the modified algorithm performs better under light loads.

Problem 1.18 — Solution

The solution is shown in the figure below. Recall that nonpersistent CSMA operates so that if the medium is idle, it transmits; and if the medium is busy, it waits a random amount of time and senses the channel again. The medium is decided idle if there are no transmissions for time duration β, which in our case equals τ. Therefore, although station *B* will find the medium idle at the time its packet arrives, ($\tau/2$), because of τ propagation delay, the medium will become busy during the τ sensing time. The figure below shows the case where station *C* happens to sense the channel idle before *B* so station *C* transmits first.

For CSMA/CD, the smallest frame size is twice the propagation time, i.e., 2τ, which is the duration of the collision.

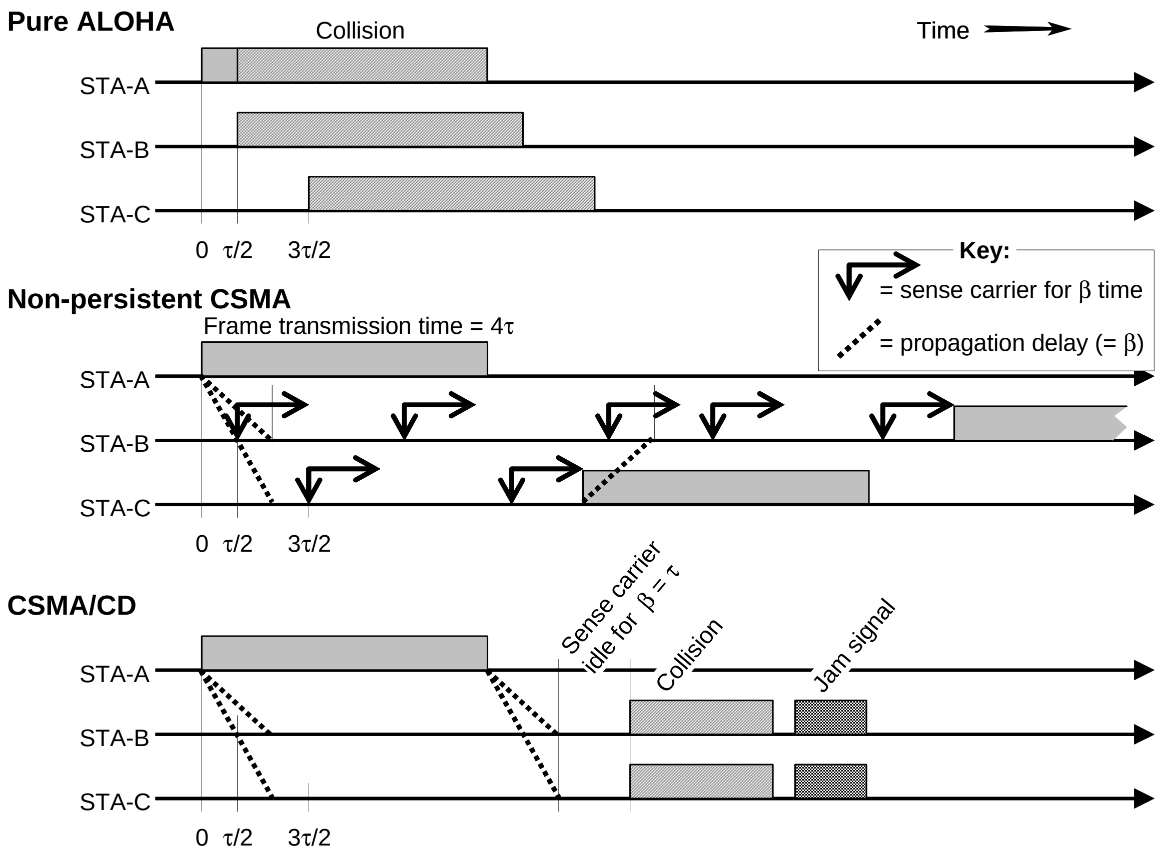

Problem 1.19 — Solution

Problem 1.20 — Solution

The solution for the three stations using the CSMA/CA protocol is shown in the figure below. The third station has the smallest initial backoff and transmits the first. The other two stations will freeze their countdown when they sense the carrier as busy. After the first frame, STA3 randomly chooses the backoff value equal to 4, and the other two stations resume their previous countdown. STA1 reaches zero first and transmits, while STA2 and STA3 freeze their countdown waiting for the carrier to become idle. STA3 transmits next its second frame and randomly chooses the backoff value equal to 1, and the other two stations again resume their previous countdown. STA2 finally transmits its first frame but STA3 simultaneously transmits its third frame and there is a collision.

(Note: The reader may wish to compare this with Figure 1-32 which illustrates a similar scenario for three stations using the CSMA/CD protocol, as well as with Problem 1.33 for IEEE 802.11.)

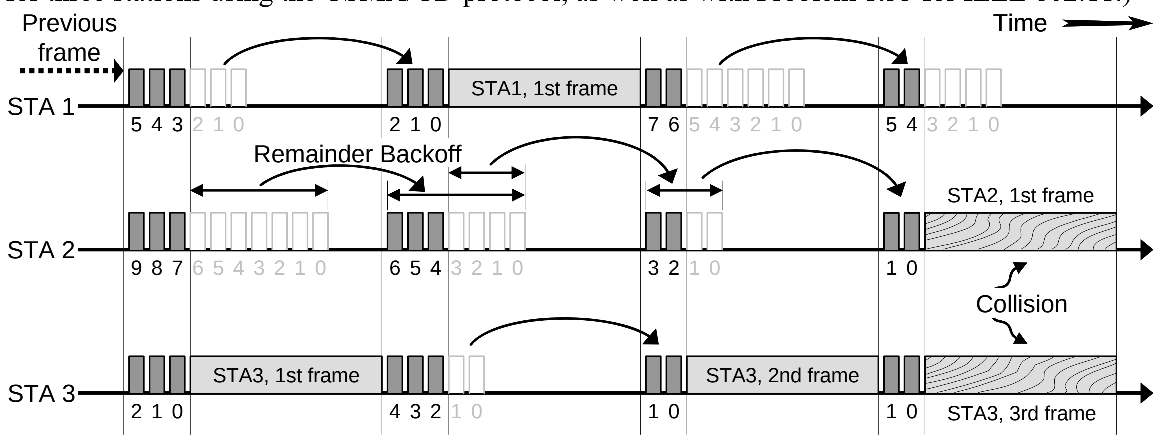

Problem 1.21 — Solution

The analogy with a slide (Figure 1-80) is shown again in the figure below. We define three probabilities for a station to transition between the backoff states. P_{11} represents the probability that the kid will enter the slide on Platform-1, therefore after sliding through Tube-1 he again slides through Tube-1. P_{12} represents the probability that the kid will enter the slide on Platform-2, therefore after sliding through Tube-1 he first slides through Tube-2. Finally, P_{21} represents the probability that the kid will slide through Tube-1 after sliding through Tube-2, and this equals 1 because the kid has no choice. Below I will present three possible ways to solve this problem.

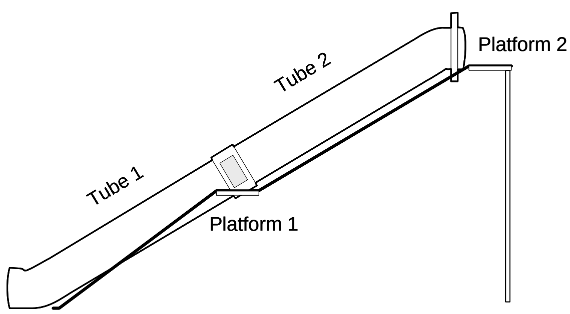

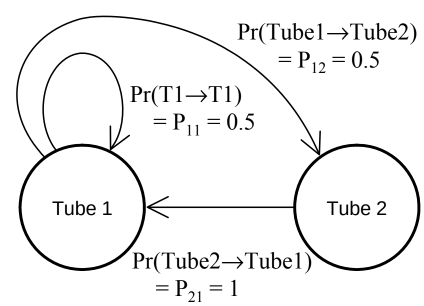

Solution 1 (Statistical)

There are two events:

(a) The kid enters the slide on Platform-1 and slides through Tube-1

(b) The kid enters the slide on Platform-2 and slides through Tube-2, then through Tube-1

Because the probabilities of choosing Platform-1 or proceeding to Platform-2 are equal, these two events will, statistically speaking, happen in a regular alternating sequence:

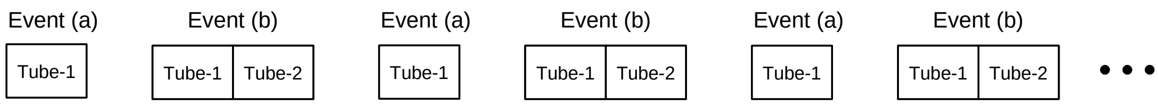

By simple observation we can see that the kid will be found two thirds of time in Tube-1 and one third of time in Tube-2. Therefore, P_{T1} = 2/3 and P_{T2} = 1/3 and these correspond to the distribution of steady-state probabilities of backoff states.

Solution 2 (Algebraic)

We introduce two variables, $x_1(k)$ and $x_2(k)$ to represent the probabilities that at time k the kid will be in Tube-1 or Tube-2, respectively. Then we can write the probabilities that at time $k + 1$ the kid will be in one of the tubes as:

$$\begin{aligned} x_1(k+1) &= P_{11} \cdot x_1(k) + P_{21} x_2(k) &&= \tfrac{1}{2} x_1(k) + x_2(k) \\ x_2(k+1) &= P_{12} \cdot x_1(k) &&= \tfrac{1}{2} x_1(k) \end{aligned}$$

Write the above equations using matrix notation:

$$\mathrm{X}(k+1) = \begin{bmatrix} \frac{1}{2} & 1 \\ \frac{1}{2} & 0 \end{bmatrix} \cdot \mathrm{X}(k) = \mathrm{A} \cdot \mathrm{X}(k)$$

Find the eignevalues of the matrix **A**:

$$\left| \mathrm{A} - \lambda \cdot \mathrm{I} \right| = \begin{vmatrix} \frac{1}{2} - \lambda & 1 \\ \frac{1}{2} & -\lambda \end{vmatrix} = \lambda^2 - \tfrac{1}{2}\lambda - \tfrac{1}{2} = 0 \qquad \Rightarrow \qquad \lambda_1 = 1, \quad \lambda_2 = -\tfrac{1}{2}$$

and eigenvectors of the matrix **A**:

$$\lambda = \lambda_1 = 1 \qquad \Rightarrow \qquad \begin{bmatrix} -\frac{1}{2} & 1 \\ \frac{1}{2} & -1 \end{bmatrix} \cdot [v_1] = 0 \qquad \Rightarrow \qquad v_1 = \begin{bmatrix} 2 \\ 1 \end{bmatrix}$$

$$\lambda = \lambda_2 = -\tfrac{1}{2} \qquad \Rightarrow \qquad \begin{bmatrix} 1 & 1 \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix} \cdot [v_2] = 0 \qquad \Rightarrow \qquad v_2 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$$

So,

$$S = [v_1 \quad v_2] = \begin{bmatrix} 2 & 1 \\ 1 & -1 \end{bmatrix} \quad \text{and} \quad S^{-1} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & -\frac{2}{3} \end{bmatrix}, \qquad \Lambda = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & -\frac{1}{2} \end{bmatrix}$$

Then we solve the backoff-state equation:

$$\begin{aligned} \mathrm{X}(k) = \mathrm{A}^k \cdot \mathrm{X}(0) = S \cdot \Lambda^k \cdot S^{-1} &= \begin{bmatrix} 2 & 1 \\ 1 & -1 \end{bmatrix} \cdot \begin{bmatrix} (1)^k & 0 \\ 0 & (-\frac{1}{2})^k \end{bmatrix} \cdot \begin{bmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & -\frac{2}{3} \end{bmatrix} \cdot \mathrm{X}(0) \\ &= \begin{bmatrix} \frac{2}{3} \cdot (1)^k + \frac{1}{3} \cdot (-\frac{1}{2})^k & \frac{2}{3} \cdot (1)^k - \frac{2}{3} \cdot (-\frac{1}{2})^k \\ \frac{1}{3} \cdot (1)^k - \frac{1}{3} \cdot (-\frac{1}{2})^k & \frac{1}{3} \cdot (1)^k + \frac{2}{3} \cdot (-\frac{1}{2})^k \end{bmatrix} \cdot \mathrm{X}(0) \end{aligned}$$

The steady state solution is obtained for $k \to \infty$:

$$X(k) = \begin{bmatrix} \frac{2}{3} & \frac{2}{3} \\ \frac{1}{3} & \frac{1}{3} \end{bmatrix} \cdot \begin{bmatrix} x_1(0) \\ x_2(0) \end{bmatrix}$$

We obtain that regardless of the initial conditions $x_1(0)$ and $x_2(0)$, the backoff state probabilities are $x_1(k) = 2/3$ and $x_2(k) = 1/3$.

Problem 1.22 — Solution

Problem 1.23 — Solution

(a)

The lower boundary for the vulnerable period for *A*'s transmission is any time before *A*'s start (at t_A) that overlaps with *A*'s transmission. This is equal to packet transmission time (t_x), which for data rate 1 Mbps and packet length of 44 bytes gives 352 μs or 18 backoff slots. To account for the cases when *A* selects its backoff countdown $b_A < 18$ slots, we can write the lower bound of the vulnerable period as: $max\{t_A - t_x, 0\}$.

Similarly, the upper bound is the end of *A*'s packet transmission, which is at $t_A + t_x$. To account for the limited size of the backoff contention window *CW*, we can write the lower bound of the vulnerable period as: $min\{t_A + t_x, CW\}$.

Therefore, if *A* starts transmission at t_A, then the vulnerable period for the reception of this packet at the receiver *B* is $[max\{t_A - t_x, 0\}, min\{t_A + t_x, CW\}]$. In our specific example, the vulnerable period is $[max\{12 \times 20\ \mu\text{s} - 352\ \mu\text{s}, 0\}, max\{12 \times 20\ \mu\text{s} + 352\ \mu\text{s}, 32 \times 20\ \mu\text{s}\}] = [0, 592\ \mu\text{s}]$.

(b)

The timing diagram is shown in the figure below.

Notice that for the first transmission, *A* and *B* will start their backoff countdown at the same time (synchronized with each other). Conversely, for the second transmission, *A* and *B* will start their backoff countdown at different times (desynchronized with each other).

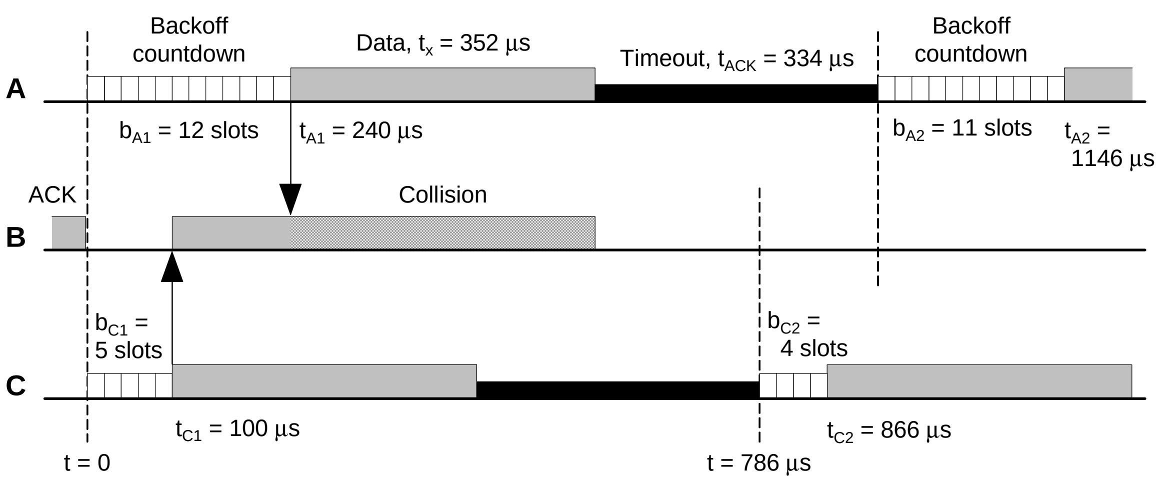

(c)

Problem 1.24 — Solution

(a)

At Kathleen's computer, the TCP layer will slice the 16-Kbytes letter into payload segments of:

TCP segment size = 512 − 20(TCP hdr) − 20 (IP hdr) = 472 bytes

Total number of IP datagrams generated is: 16,384 ÷ 472 = 34 × 472 + 1 × 336 = 35 IP datagrams

Notice that the payload of the first 34 IP datagrams is: 20(TCP hdr) + 472(user data) = 492 bytes and the last one has 356-bytes payload.

(b)

There will be no fragmentation on Link 2, but there will be on Link 3. The IP datagram size allowed by MTU of Link 3 is 256, which means that the payload of each datagram is up to 256 − 20(IP hdr) = 236 bytes. Because the fragment offset is measured in 8-byte chunks, not in bytes, the greatest number divisible by 8 without remainder is 232, which means that each of the 34 incoming datagrams will be split into 3 new fragments: 492 = 232 + 232 +28, and the last one 35th will be split into two fragments: 356 = 232 + 124.

Total number of fragments on Link 3 is: 34 × 3 + 1 × 2 = 104 datagrams, of which 69 have payload of 232 bytes, 34 have payload of 28 bytes, and 1 has payload of 124 bytes (this is the *last* one). The reader should recall that the IP layer does not distinguish any structure in its payload data. Each original IP datagram sent by Kathleen's computer contains a TCP header and user data. The TCP header will be end up in the first fragment and there will be 212 bytes of user data, and the remaining two fragments will contain only user data. Again, IP is not aware of the payload structure and treats the whole payload in the same way.

(c)

First 4 packets that Joe receives	**Last 5 packets that Joe receives**
#1: Length = 232, ID = 672, MF = 1, Offset = 0	#100: Length = 232, ID = 774, MF = 1, Offset = 0
#2: Length = 232, ID = 672, MF = 1, Offset = 29	#101: Length = 232, ID = 774, MF = 1, Offset = 29
#3: Length = 28, ID = 672, MF = 0, Offset = 58	#102: Length = 28, ID = 774, MF = 0, Offset = 58
#4: Length = 232, ID = 673, MF = 1, Offset = 0	#103: Length = 232, ID = 775, MF = 1, Offset = 0
	#104: Length = 124, ID = 775, MF = 0, Offset = 29

(d)

If the very last (104th) fragment is lost, Joe's computer cannot reassemble the last (35th) IP datagram that was sent by Kathleen's computer. Therefore, the TCP sender will not receive an acknowledgement for the last segment that it transmitted and, after the retransmission timer expires, it will resend the last segment of 336 bytes of data + 20 bytes TCP header. The IP layer will create a datagram with the payload of 356 bytes. So, Kathleen's computer will retransmit only 1 IP datagram. This datagram will be fragmented at Link 3 into 2 IP datagrams. So, Joe's computer will receive 2 IP datagrams. The relevant parameters for these 2 fragments are:

#105. Length = 232, ID = 776, MF = 1, Offset = 0
#106. Length = 124, ID = 776, MF = 0, Offset = 29

Notice that the ID of the retransmitted segment is different from the original, to avoid confusion between the fragments of two different datagrams.

Problem 1.25 — Solution

After the steps (a) through (e), the network will look like:

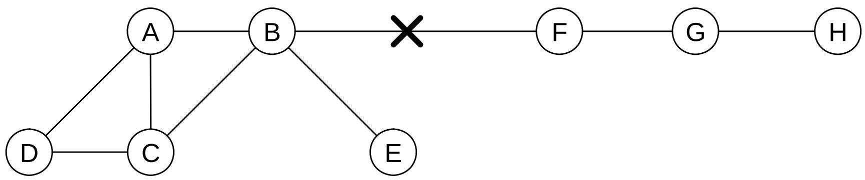

The link-state advertisements flooded after the step (a) are as follows:

Node A:

Seq#	Neighbor	Cost
1	B	
	C	

Node B:

Seq#	Neighbor	Cost
1	A	1
	C	1
	F	∞

Node C:

Seq#	Neighbor	Cost
1	A	1
	B	1

Node F:

Seq#	Neighbor	Cost
1	B	∞
	G	1

Node G:

Seq#	Neighbor	Cost
1	F	1

[We assume that the sequence number starts at 1 in step (a), although a valid assumption would also be that it is 1 before step (a).]

In the remaining steps, we show only those LSAs that are different from the previous step. (Note that the sequence number changes in every step for all LSAs, including the ones not shown.)

Step (b):

Node G:

Seq#	Neighbor	Cost
2	F	1
	H	1

Node H:

Seq#	Neighbor	Cost
2	G	1

Step (c):

Node C:

Seq#	Neighbor	Cost
3	A	1
	B	1
	D	1

Node D:

Seq#	Neighbor	Cost
3	C	1

Step (d):

Node B:

Seq#	Neighbor	Cost
4	A	1
	C	1
	E	1

Node E:

Seq#	Neighbor	Cost
4	B	1

Step (e):

Node A:

Seq#	Neighbor	Cost
5	B	1
	C	1
	D	1

Node D:

Seq#	Neighbor	Cost
5	A	1
	C	1

Step (f):

Node B:

Seq#	Neighbor	Cost
6	A	1
	C	1
	E	1
	F	1

Node F:

Seq#	Neighbor	Cost
6	B	1
	G	1

Problem 1.26 — Solution

Problem 1.27 — Solution

The tables of distance vectors at all nodes after the network stabilizes are shown in the leftmost column of the figure below. Notice that, although, there are two alternative $\overline{AC}$ links, the nodes select the best available, which is $\overline{AC}$ =1.

When the link $\overline{AC}$ with weight equal to 1 is broken, both *A* and *C* detect the new best cost $\overline{AC}$ as 50.

1. *A* computes its new distance vector as
   $D_A(B) = \min\{c(A,B) + D_B(B),\ c(A,C) + D_C(B)\} = \min\{4+0,\ 50+1\} = 4$
   $D_A(C) = \min\{c(A,C) + D_C(C),\ c(A,B) + D_B(C)\} = \min\{50+0,\ 4+1\} = 5$
   Similarly, *C* computes its new distance vector as
   $D_C(A) = \min\{c(C,A) + D_A(A),\ c(C,B) + D_B(A)\} = \min\{50+0,\ 1+2\} = 3$
   $D_C(B) = \min\{c(C,B) + D_B(B),\ c(C,A) + D_A(B)\} = \min\{1+0,\ 50+2\} = 1$
   Having a global view of the network, we can see that the new cost $D_C(A)$ via *B* is *wrong*. Of course, *C* does not know this and therefore a *routing loop* is created.
   Both *B* and *C* send their new distance vectors out, each to their own neighbors, as shown in the second column in the figure (**first exchange**).
2. Upon receiving *C*'s distance vector, *A* is content with its current d.v. and makes no changes. Ditto for node *C*.
   B computes its new distance vector as
   $D_B(A) = \min\{c(B,A) + D_A(A),\ c(B,C) + D_C(A)\} = \min\{4+0,\ 1+3\} = 4$
   $D_B(C) = \min\{c(B,C) + D_C(C),\ c(B,A) + D_A(C)\} = \min\{1+0,\ 4+5\} = 1$
   B sends out its new distance vector to *A* and *C* (**second exchange**).
3. Upon receiving *B*'s distance vector, *A* does not make any changes so it remains silent. Meanwhile, *C* updates its distance vector to the correct value for $D_C(A)$ and sends out its new distance vector to *A* and *B* (**third exchange**).
4. *A* and *B* will update the *C*'s distance vector in their own tables, but will *not* make further updates to their own distance vectors. There will be no further distance vector exchanges related to the $\overline{AC}$ breakdown event.

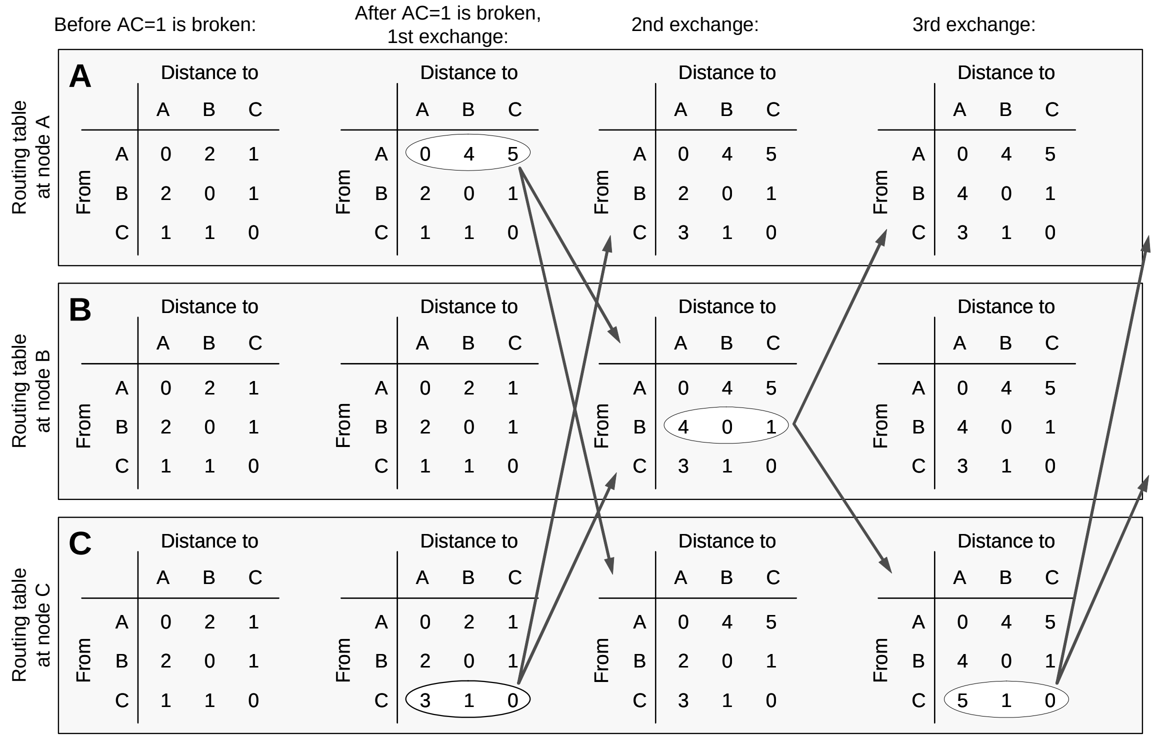

Problem 1.28 — Solution

Problem 1.29 — Solution

(a)

The following figure shows how the routing tables are constructed until they stabilize.

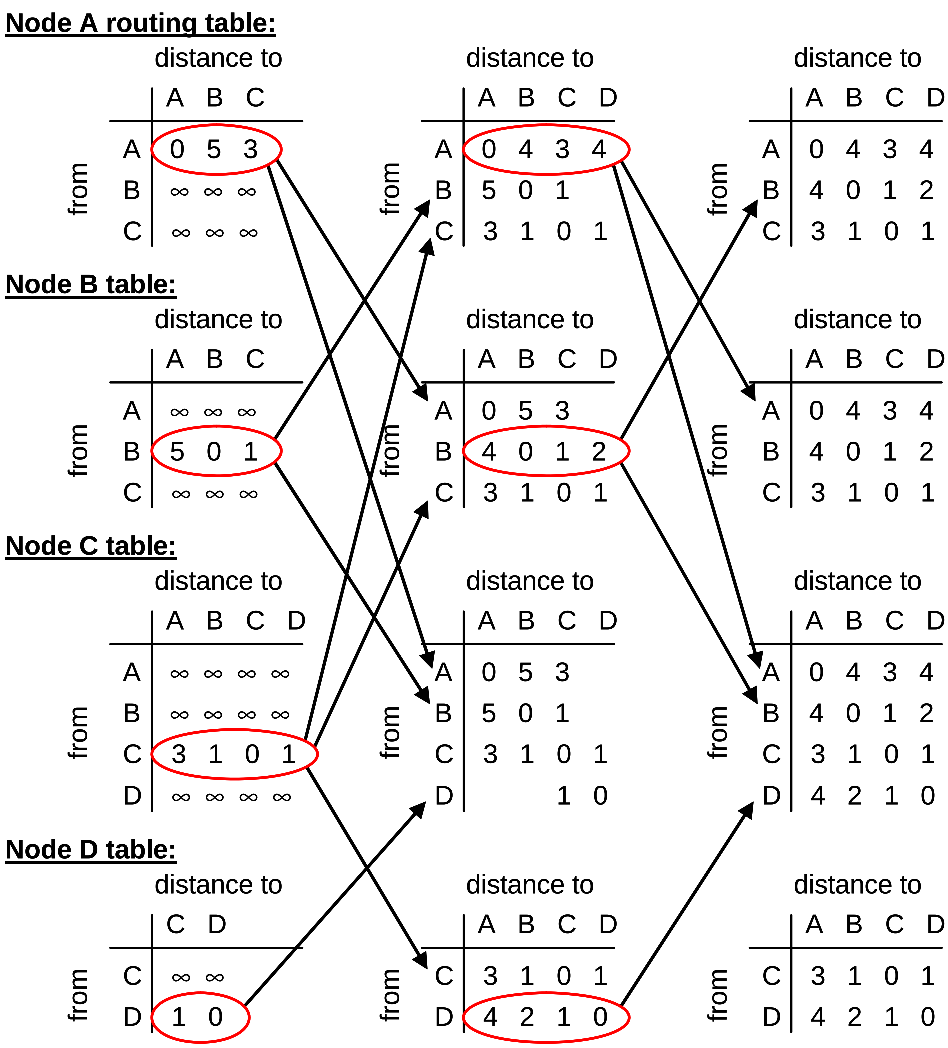

(b)

The forwarding table at node *A* after the routing tables stabilize is shown in the figure below. Notice that the forwarding table at each node is kept separately from the node's routing table.

Node A forwarding table:

destination	interface
A	--
B	AC
C	AC
D	AC

(c)

See also solution of Problem 1.25.

First, consider the figure below. *C* updates its distance vector and thinks that the new shortest distance to *D* is 4 (via *B*). It sends its updated distance vector to its neighbors (*A* and *D*) and they update their distance vectors, as shown in the figure.

A routing loop is formed because a packet sent from *A* to *D* would go to *C* and *C* would send it to *B* (because *C*'s new shortest path to *D* is via *B*). *B* would return the packet to *C* because *B*'s shortest path to *D* is still via *C*. This looping of the packet to *D* would continue forever.

Before link failure

After link failure is detected

C sends its new distance vector

Node A table:

Before link failure (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

After link failure is detected (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

C sends its new distance vector (distance to):

from	A	B	C	D
A	0	4	3	**6**
B	4	0	1	2
C	3	1	0	3

Node B table:

Before link failure (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

After link failure is detected (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

C sends its new distance vector (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	**4**
C	3	1	0	3

Node C table:

Before link failure (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1
D	4	2	1	0

After link failure is detected (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	**3**

C sends its new distance vector (distance to):

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	3

Node D table:

Before link failure (distance to):

from	A	B	C	D
C	3	1	0	1
D	4	2	1	0

After link failure is detected (distance to):

from	D
D	0

C sends its new distance vector (distance to):

from	D
D	0

(d)

If the nodes use split-horizon routing, then neither *A* nor *B* would advertise to *C* their distances to *D*, because *C* is the next hop for both of them on their paths to *D*. Therefore, in principle *D* would not think there is an alternative path to *D*.

Even in this case, it is possible that a routing loop forms. The key to a routing loop formation in this case is the periodic updates that nodes running a distance vector protocol are transmitting. Thus, the cycle shown below is due to the pathological case where *B* transmits its periodic report some time in the interval after the link *CD* crashes, but before *B* receives update information about the outage to *D* from node *C*.

Before link failure	*After link failure is detected*	*B sends its periodic report* *C sends its updated report*

Node A table:

Before link failure (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

After link failure is detected (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

B sends its periodic report, C sends its updated report (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	∞

Node B table:

Before link failure (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

After link failure is detected (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1

B sends its periodic report, C sends its updated report (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	4
C	3	1	0	∞

Node C table:

Before link failure (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	1
D	4	2	1	0

After link failure is detected (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	**3**

B sends its periodic report, C sends its updated report (distance to; rows: from)

from	A	B	C	D
A	0	4	3	4
B	4	0	1	2
C	3	1	0	**3**

Node D table:

Before link failure (distance to; rows: from)

from	A	B	C	D
C	3	1	0	1
D	4	2	1	0

After link failure is detected (distance to; rows: from)

from	D
D	0

B sends its periodic report, C sends its updated report (distance to; rows: from)

from	D
D	0

Problem 1.30 — Solution

Problem 1.31 — Solution

(a)

Example IP address assignment is as shown:

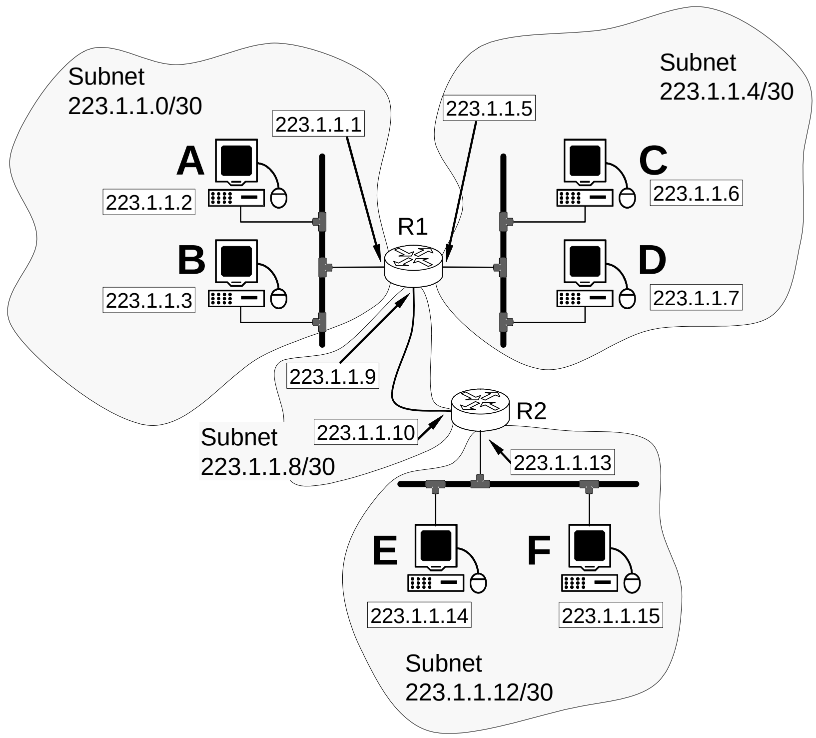

(b)

Routing tables for routers R1 and R2 are:

Router R1	**Destinat. IPaddr / Subnet Mask**	**Next Hop**	**Output Interface**
	223.1.1.2 (A)	223.1.1.2 (A)	223.1.1.1
	223.1.1.3 (B)	223.1.1.3 (B)	223.1.1.1
	223.1.1.6 (C)	223.1.1.6 (C)	223.1.1.5
	223.1.1.7 (D)	223.1.1.7 (D)	223.1.1.5
	223.1.1.12/30	223.1.1.10	223.1.1.9

Router R2	**Destinat. IPaddr / Subnet Mask**	**Next Hop**	**Output Interface**
	223.1.1.14 (E)	223.1.1.14 (E)	223.1.1.13
	223.1.1.15 (F)	223.1.1.15 (F)	223.1.1.13
	223.1.1.0/30	223.1.1.9	223.1.1.10
	223.1.1.4/30	223.1.1.9	223.1.1.10

Problem 1.32 — Solution

Recall that in CIDR the *x* most significant bits of an address of the form *a.b.c.d*/*x* constitute the network portion of the IP address, which is referred to as ***prefix*** (or *network prefix*) of the address. In our case the forwarding table entries are as follows:

Subnet mask	Network prefix	Next hop
223.92.32.0/20	**11011111 01011100 0010**0000 00000000	A
223.81.196.0/12	**11011111 0101**0001 11000100 00000000	B
223.112.0.0/12	**11011111 0111**0000 00000000 00000000	C

223.120.0.0/14	**11011111 011110**00 00000000 00000000	D
128.0.0.0/1	**1**0000000 00000000 00000000 00000000	E
64.0.0.0/2	**01**000000 00000000 00000000 00000000	F
32.0.0.0/3	**001**00000 00000000 00000000 00000000	G

Notice that the network prefix is shown in bold face, whereas the remaining 32−x bits of the address are shown in gray color. When forwarding a packet, the router considers only the leading x bits of the packet's destination IP address, i.e., its network prefix.

	Packet destination IP address			Longest prefix match	Next hop
(a)	195.145.34.2	=	11000011 10010001 00100010 00000010	1	E
(b)	223.95.19.135	=	11011111 01011111 00010011 10000111	11011111 0101	B
(c)	223.95.34.9	=	11011111 01011111 00100010 00001001	11011111 0101	B
(d)	63.67.145.18	=	00111111 01000011 10010001 00010010	001	G
(e)	223.123.59.47	=	11011111 01111011 00111011 00101111	11011111 011110	D
(f)	223.125.49.47	=	11011111 01111101 00110001 00101111	11011111 0111	C

Problem 1.33 — Solution

The packet forwarding is given as follows:

	Destination IP address	Binary representation	Next hop
(a)	128.6.4.2 (cs.rutgers.edu)	10000000 00000110 00000100 00000010	A
(b)	128.6.236.16 (caip.rutgers.edu)	10000000 00000110 11101100 00010000	B
(c)	128.6.29.131 (ece.rutgers.edu)	10000000 00000110 00011101 10000011	C
(d)	128.6.228.43 (toolbox.rutgers.edu)	10000000 00000110 11100100 00001010	D

From this, we can reconstruct the forwarding table as:

Network Prefix	Subnet Mask	Next Hop
10000000 00000110 0000	128.6.0.0 / 20	A
10000000 00000110 11101	128.6.232.0 / 21	B
10000000 00000110 0001	128.6.16.0 / 20	C
10000000 00000110 111	128.6.224.0 / 19	D

Notice that in the last row we could have used the prefix "10000000 00000110 11100" and the subnet mask "128.6.224.0 / 21." However, it suffices to use only 19 most significant bits because the router forwards to the *longest possible match*, as explained next.

If we calculate the range of address associated with the prefix 128.6.224.0/19 (last row in the above forwarding table), we find the range as 128.6.224.0 to 128.6.255.254. The addresses associated with 128.6.232.0/21 (second row) are 128.6.232.1 to 128.6.239.254. As seen, the latter set of addresses is a subset of the former set, so one may think that the router will route all

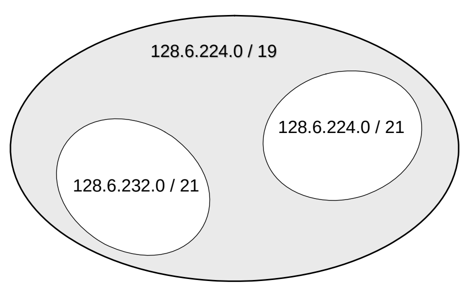

addresses starting with 128.6.224.0/19 to the next hop *D*. This is not the case. When a packet arrives with 128.6.228.43 in its header, the router will find that the longest match is 128.6.224.0/19 (last row in the above table). Hence, this packet will be forwarded to *D*. Alternatively, when a packet arrives with 128.6.236.16 in its header, the router will find that there are two matches in the forwarding table: 128.6.224.0/19 (last row) and 128.6.232.0/21 (second row). Of these two, the latter is longer by two bits, so the router will not get confused and route this packet to *D*. Rather, the router will correctly route the packet to *B*.

Problem 1.34 — Solution

Problem 1.35 — Solution

The observed topology of an internetwork of autonomous systems depends on the vantage point. The view from ASφ's vantage point is shown in the problem statement. When solving the problem for other stub ASs, it is key to keep in mind that ASs do not like to provide transit without being paid. ASs will happily provide transit to their paying customers, or will provide access of their customers to their peers, but will not provide free transit to their peers.

The solutions for the remaining vantage points are as follows:

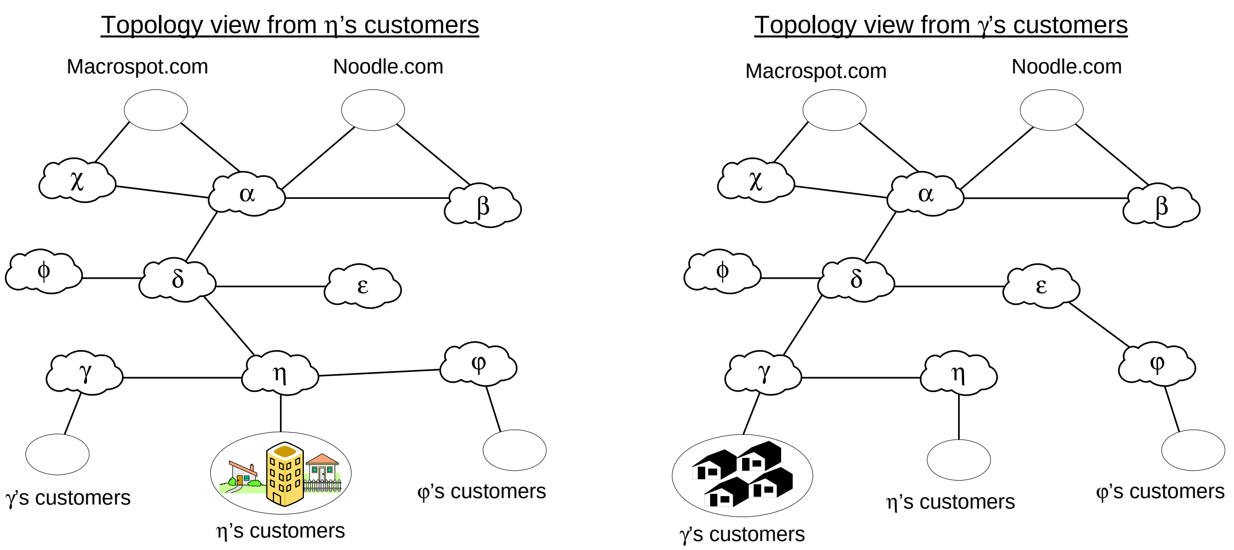

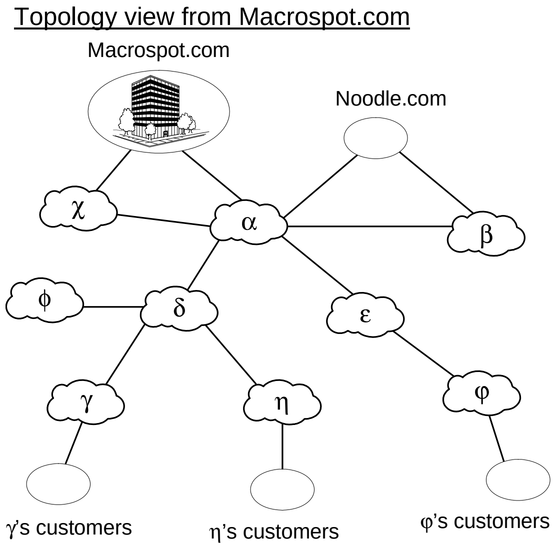

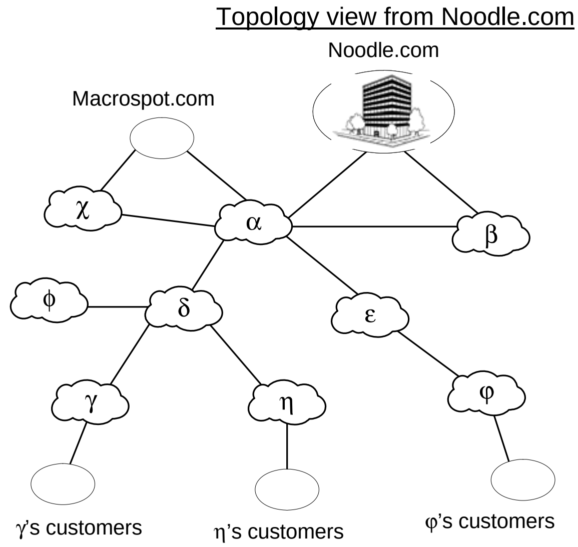

Problem 1.36 — Solution

(a)

Recall that autonomous systems use path vector routing for inter-domain routing (Section 1.4.5). ASα will receive three routes to ASβ, along router links CH, path: $\langle 1 \mid \beta \rangle$; DK, path: $\langle 1 \mid \beta \rangle$; and DE, path: $\langle 2 \mid \gamma, \beta \rangle$.

(b)

$X \rightarrow Y$ traffic will take link CH, because this is the shortest path when crossing ASα. Recall that internally ASs use hot-potato routing, so when a packet from X arrives at router A, the shortest path across ASα to ASβ is AC, then CH. Note that this strategy minimizes cost to the source of the traffic (i.e., ASα) and is not optimal to other ASs along the path, such as ASβ.

By the same reasoning, $Y \rightarrow X$ traffic will take link KD. The resulting paths are $X \rightarrow A \rightarrow C \rightarrow H \rightarrow I \rightarrow J \rightarrow Y$ and $Y \rightarrow J \rightarrow K \rightarrow D \rightarrow B \rightarrow A \rightarrow X$. In both cases, hot-potato routing within the given AS does not look at the overall path length, but only the path length within this AS.

(c) To have all $X \rightarrow Y$ traffic take link DK, the Exterior Gateway Protocol of ASα could simply be configured with a routing policy that prefers link DK in all cases. The Exterior Gateway Protocol of ASα then would simply not tell the Interior Gateway Protocol (IGP) of ASα that it is possible to reach ASβ via the speaker C.

The only general solution, though, is for ASα to accept into its routing tables some of the internal structure of ASβ, so that the IGP protocol of ASα, for example, knows where Y is relative to links CH and DK. (In our example both alternative paths $X \rightarrow Y$ (or $Y \rightarrow X$) are equally long, so both ASs would need to break the tie in the same way.)

Of course, if one of the links CH or DK (or an attached speaker) goes down, then all $X \rightarrow Y$ traffic and $Y \rightarrow X$ traffic will be forced to take the same path.

(d)

If ASα were configured with a routing policy to prefer AS paths through ASγ, or to avoid AS paths involving links direct links to ASβ, then ASα might route to β via γ.

Problem 1.37 — Solution

Problem 1.38 — Solution

Problem 1.39 — Solution

Problem 1.40 — Solution

PPP and Link Control Protocol (LCP)

(a)

If packets are not identified with unique identifiers, then the following sequence may happen, due to delayed acknowledgements from Responder. The delayed ACKs are falsely interpreted as being ACKs for a subsequent disconnect and link configuration. This in turn causes a failure in authentication.

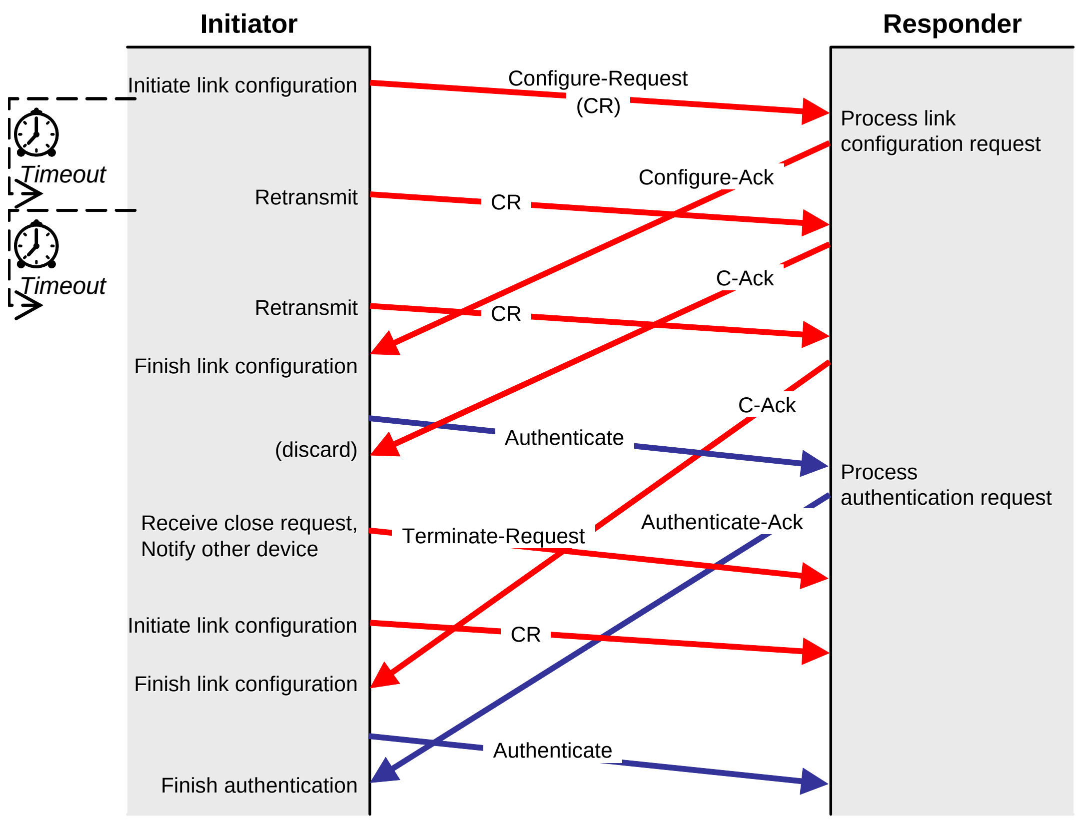

Note: I was not able to find in the literature explicitly stated the rationale for unique identifiers in LCP frames. This example is based on [Bertsekas & Gallagher, 1992], Section 2.7.2, Figure 2.44, which is an example of initialization failure for the HDLC protocol. Because PPP is derived from HDLC, I assume that its designers anticipated this problem. Another example is also from

[Bertsekas & Gallagher, 1992], Section 2.7.4, Figure 2.46, which happens in case of a sequence of node failures.

(b)

If it is not necessary to acknowledge a Terminate-Request, then one side may terminate the link, while the other side still considers the link open and keeps sending data.

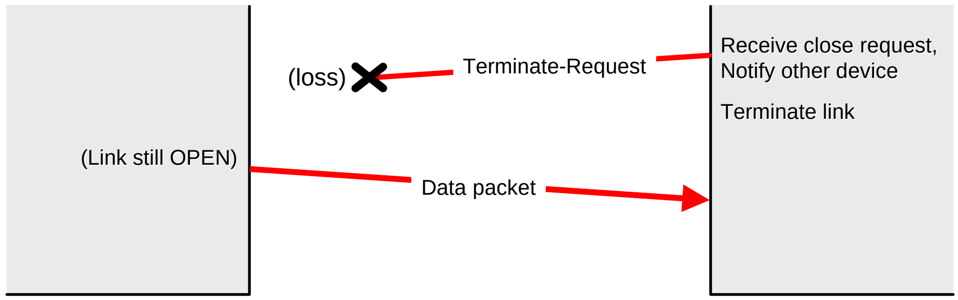

Problem 1.41 — Solution

Problem 1.42 — Solution

(a) and (b)

See the figure below for the configuration messages that are sent and how the ports are labeled. The network converges to its spanning tree in just two iterations.

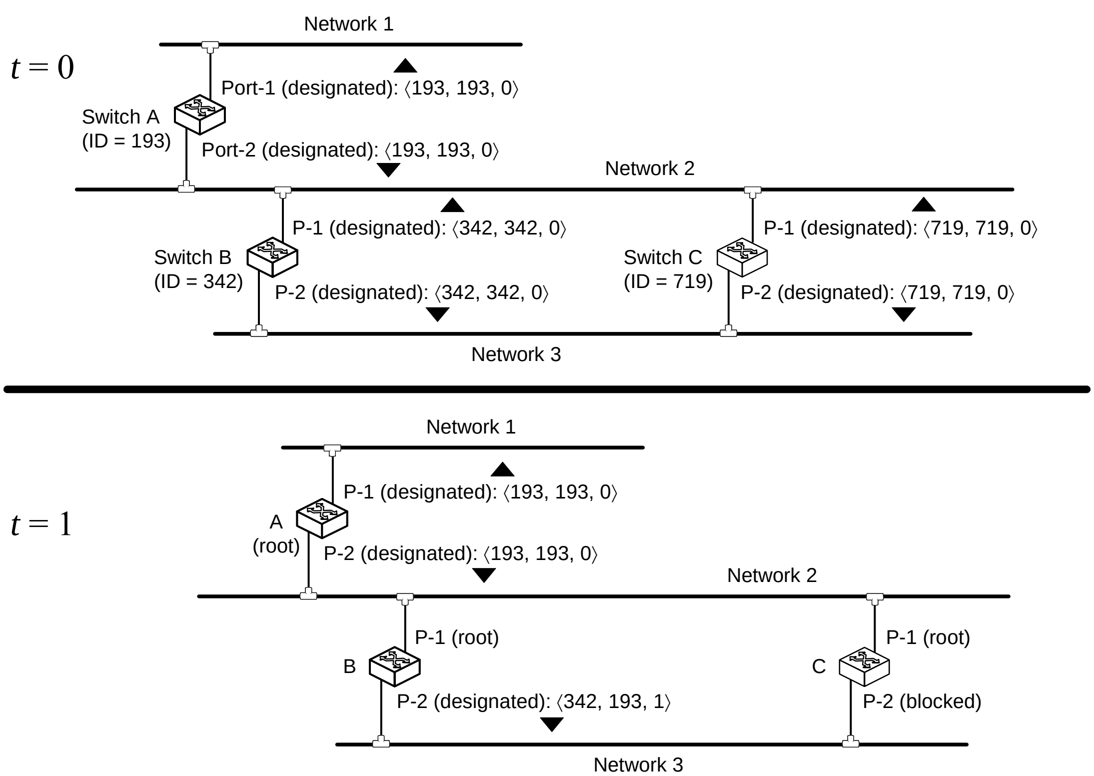

At $t = 0$:

Initially each switch selects itself as the "root switch" and for all the attached networks as the "designated switch."

Messages (each switch on both ports): Switch *A*: ⟨193, 193, 0⟩; Switch *B*: ⟨342, 342, 0⟩; Switch *C*: ⟨719, 719, 0⟩.

At $t = 1$:

Switch *A* selects itself as the "designated switch" for Network segments 1 and 2 because it has the lowest ID on both segments. Both of its ports become "designated."

Switch *B* selects *A* as the "designated switch" for Network 2 and selects itself as the "designated switch" for Network 3, because these are the lowest ID on their respective network segments. Its Port-1 becomes "root port" and Port-2 becomes "designated port."

Switch *C* selects *A* as the "designated switch" for Network 2 and selects *B* as the "designated switch" for Network 3, because they have the lowest ID on their respective network segments. Its Port-1 becomes "root port" and its Port-2 becomes blocked.

Messages: Switch *A*: ⟨193, 193, 0⟩ on both ports; Switch *B*: ⟨342, 193, 1⟩ on Port-2; Switch *C*: none.

At $t = 2$:

Messages: Switch *A*: ⟨193, 193, 0⟩ on both ports; Switch *B*: ⟨342, 193, 1⟩ on Port-2; Switch *C*: none.

(c)
The network reached the stable state in two iterations. After this, only switch *A* (root) will generate configuration messages and switch *B* will forward these messages only over its Port-2 for which it is the designated switch.

(d)
After the network stabilizes, a frame sent by station *X* to station *Y* will traverse the path shown in the figure below. Notice that, unlike routing, in LAN switches the frame is not first transmitted to the "next hop" and then relayed by that hop. That is, switch *A* does not "relay" a packet from host *X* to host *Y*, although *A* is elected as the root of the spanning tree. This is because the spanning tree protocol (STP) operates transparently to the backward learning algorithm, which learns the switching tables based on the network topology configured by STP unknown to the backward learning algorithm. In our example, only switch *B* will "relay" a packet from host *X* to host *Y*.

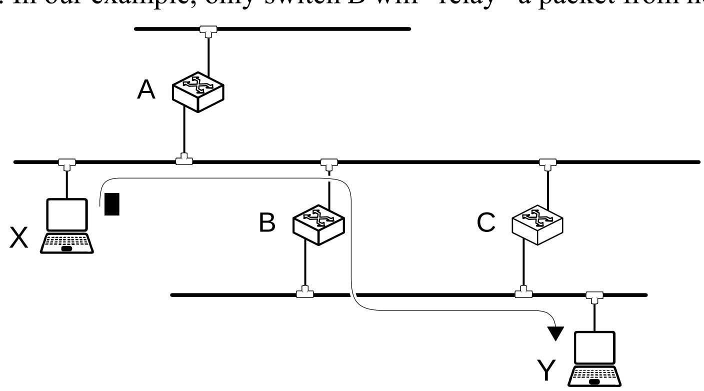

Problem 1.43 — Solution

Problem 1.44 — Solution

The solution for three stations using the 802.11 protocol is similar to that of Problem 1.18 and is shown in the figure below. Again, the third station has the smallest initial backoff and transmits the first. The other two stations will freeze their countdown when they sense the carrier as busy. After the first frame STA3 randomly chooses the backoff value equal to 4, and the other two stations resume their previous countdown. STA1 reaches zero first and transmits, while STA2 and STA3 freeze their countdown waiting for the carrier to become idle. Next, STA3 transmits its second frame and randomly chooses the backoff value equal to 3, and the other two stations again resume their previous countdown. STA2 finally transmits its first frame but STA3 simultaneously transmits its third frame and there is a collision.

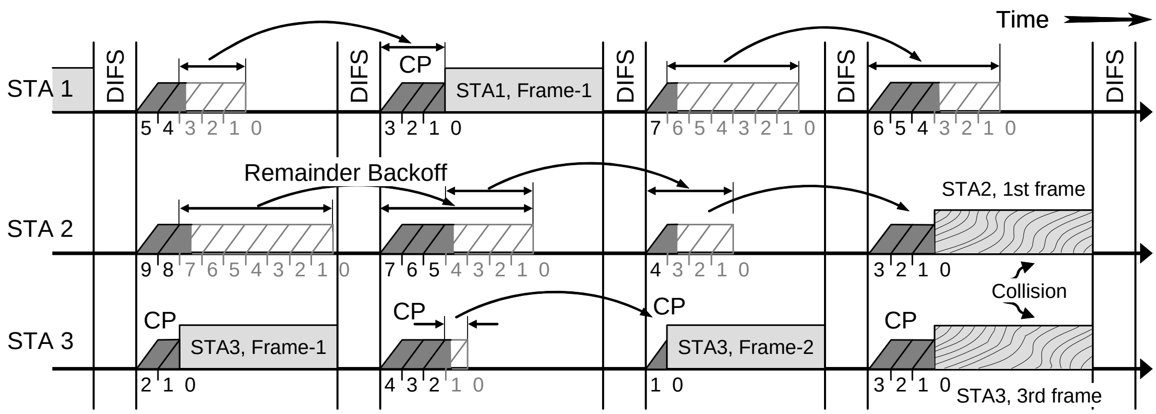

Problem 1.45 — Solution

There is no access point, as this is an independent BSS (IBSS). Remember that 802.11 stations cannot transmit and receive simultaneously, so once a station starts receiving a packet, it cannot contend for transmission until the next transmission round.

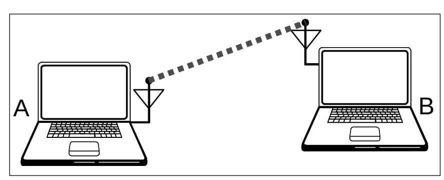

The solution is shown in the figure below (check also Figure 1-75 in Example 1.6). We make arbitrary assumption that packet arrives first at *B* and likewise that after the first transmission *A* claims the channel before *B*. Notice that the durations of EIFS and ACK_Timeout are the same (Table 1-6). Station *B* selects the backoff=6 and station *A* transmits the first, while station *B* carries over its frozen counter and resumes the countdown from backoff=2.

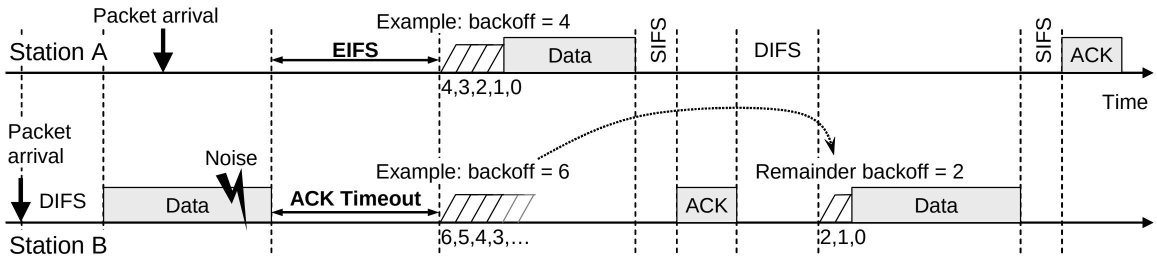

Problem 1.46 — Solution

Problem 2.1 — Solution

(a)

Scenario 1: the initial value of `TimeoutInterval` is picked as 3 seconds.

At time 0, the first segment is transmitted and the initial value of `TimeoutInterval` is set as 3 s. The timer will expire before the ACK arrives, so the segment is retransmitted and this time the RTO timer is set to twice the previous size, which is 6 s. The ACK for the initial transmission arrives at 5 s, but the sender cannot distinguish whether this is for the first transmission or for the retransmission. This does not matter, because the sender simply accepts the ACK, doubles the congestion window size (it is in slow start) and sends the second and third segments. The RTO timer is set at the time when the second segment is transmitted and the value is 6 s (unchanged). At 8 s a duplicate ACK will arrive for the first segment and it will be ignored, with no action taken. At 10 s, the ACKs will arrive for the second and third segments, the congestion window doubles and the sender sends the next four segments. `SampleRTT` is measured for both the second and third segments, but we assume that the sender sends the fourth segment immediately upon receiving the ACK for the second. This is the first `SampleRTT` measurement so

$$\texttt{EstimatedRTT} = \texttt{SampleRTT} = 5 \text{ s}$$

$$\texttt{DevRTT} = \texttt{SampleRTT} / 2 = 2.5 \text{ s}$$

and the RTO timer is set to

$$\texttt{TimeoutInterval} = \texttt{EstimatedRTT} + 4 \cdot \texttt{DevRTT} = 15 \text{ s}$$

After the second `SampleRTT` measurement (ACK for the third segment), the sender will have

$$\texttt{EstimatedRTT} = (1-\alpha) \cdot \texttt{EstimatedRTT} + \alpha \cdot \texttt{SampleRTT}$$

$$= 0.875 \times 5 + 0.125 \times 5 = 5 \text{ s}$$

$$\texttt{DevRTT} = (1-\beta) \cdot \texttt{DevRTT} + \beta \cdot |\, \texttt{SampleRTT} - \texttt{EstimatedRTT} \,|$$

$$= 0.75 \times 2.5 + 0.25 \times 0 = 1.875 \text{ s}$$

but the RTO timer is already set to 15 s and remains so while the fifth, sixth, and seventh segments are transmitted. The following table summarizes the values that `TimeoutInterval` is set to for the segments sent during the first 11 seconds:

Times when the RTO timer is set	**RTO timer values**
$t = 0$ s (first segment is transmitted)	`TimeoutInterval` = 3 s (initial guess)
$t = 3$ s (first segment is retransmitted)	`TimeoutInterval` = 6 s (RTO doubling)
$t = 10$ s (fourth and subsequent segments)	`TimeoutInterval` = 15 s (estimated value)

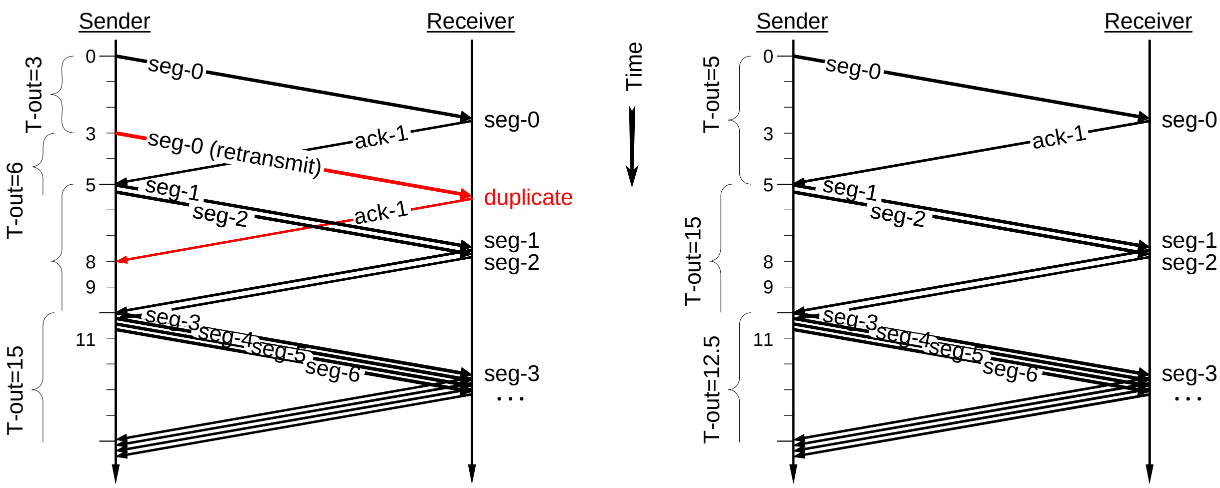

Scenario 1: initial T-out = 3 s   Scenario 2: initial T-out = 5 s

(b)

As shown in the above figure, the sender will transmit seven segments during the first 11 seconds and there will be a single (unnecessary) retransmission.

(c)

Scenario 2: the initial value of `TimeoutInterval` is picked as 5 seconds.

This time the sender correctly guessed the actual RTT interval. Therefore, the ACK for the first segment will arrive before the RTO timer expires. This is the first `SampleRTT` measurement and, as above, `EstimatedRTT` = 5 s, `DevRTT` = 2.5 s. When the second segment is transmitted, the RTO timer is set to `TimeoutInterval` = `EstimatedRTT` + 4 · `DevRTT` = 15 s.

After the second `SampleRTT` measurement (ACK for the *second* segment), the sender will have, as above, `EstimatedRTT` = 5 s, `DevRTT` = 1.875 s.

When the fourth segment is transmitted, the RTO timer is set to

$$\texttt{TimeoutInterval} = \texttt{EstimatedRTT} + 4 \cdot \texttt{DevRTT} = 12.5 \text{ s}$$

After the third `SampleRTT` measurement (ACK for the *third* segment), the sender will have

$$\texttt{EstimatedRTT} = 0.875 \times 5 + 0.125 \times 5 = 5 \text{ s}$$

$$\texttt{DevRTT} = 0.75 \times 1.875 + 0.25 \times 0 = 1.40625 \text{ s}$$

but the RTO timer is already set to 12.5 s and remains so while the fifth, sixth, and seventh segments are transmitted. The following table summarizes the values that `TimeoutInterval` is set to for the segments sent during the first 11 seconds:

Times when the RTO timer is set	RTO timer values
t = 0 s (first segment is transmitted)	`TimeoutInterval` = 3 s (initial guess)
t = 5 s (second segment is transmitted)	`TimeoutInterval` = 15 s (estimated value)
t = 10 s (fourth and subsequent segments)	`TimeoutInterval` = 12.5 s (estimated val.)

Problem 2.2 — Solution

The congestion window diagram is shown in the figure below.

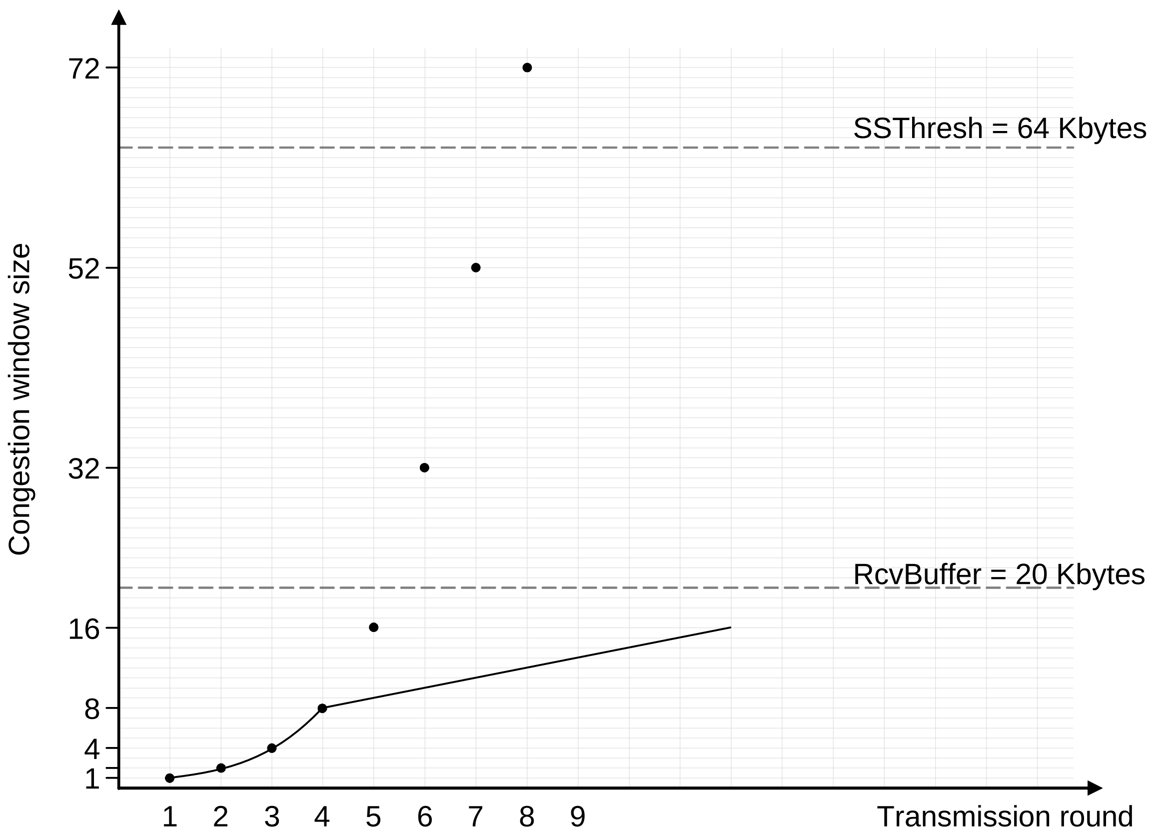

First, notice that because both hosts are fast and there is no packet loss, the receiver will never buffer the received packets, so sender will always get notified that `RcvWindow` = 20 Kbytes, which is the receive buffer size.

The congestion window at first grows exponentially. However, in transmission round #6 the congestion window of 32 × `MSS` = 32 Kbytes exceeds `RcvWindow` = 20 Kbytes. At this point the sender will send only min{`CongWin`, `RcvWindow`} = 20 segments and when these get acknowledged, the congestion window grows to 52 × `MSS`, instead of 64 × `MSS` under the exponential growth. Thereafter, the sender will keep sending only 20 segments and the congestion window will keep growing by 20 × `MSS`.

It is very important to notice that the growth is *not exponential* after the congestion window becomes $32 \times \texttt{MSS}$.

In transmission round #8 the congestion window grows to $72 \times \texttt{MSS}$, at which point it exceeds the slow start threshold (initially set to 64 Kbytes), and the sender enters the *congestion avoidance* state.

This diagram has the same shape under different network speeds, the only difference being that a transmission round lasts longer, depending on the network speed.

Problem 2.3 — Solution

Problem 2.4 — Solution

Notice that sender *A* keeps a single RTO retransmission timer for all outstanding packets. Every time a regular, non-duplicate ACK is received, the timer is reset if there remain outstanding packets. Thus, although a timer is set for segments sent in round 3×RTT, including segment #7, the timer is **reset** at time 4×RTT because packet #7 is unacknowledged. This is why the figure below shows the start of the timer for segment #7 at time 4×RTT, rather than at 3×RTT.

At time 5×RTT, sender *A* has not yet detected the loss of #7 (neither the timer expired, nor three dupACKs were received), so $\texttt{CongWin} = 7\times\texttt{MSS}$ (remains constant). There are two segments in flight (segments #7 and #8), so at time = 5×RTT sender *A* could send up to

$$\texttt{EffctWin} = \min\{\texttt{CongWin}, \texttt{RcvWindow}\} - \texttt{FlightSize} = \min\{7, 64\} - 2 = 5\times\texttt{MSS}$$

but it has nothing left to send, *A* sends a 1-byte segment to keep the connection alive. Recall that TCP guarantees *reliable transmission*, so although sender sent all data it cannot close the connection until it receives acknowledgement that all segments successfully reached the receiver. Ditto for sender *B* at time = 7×RTT.

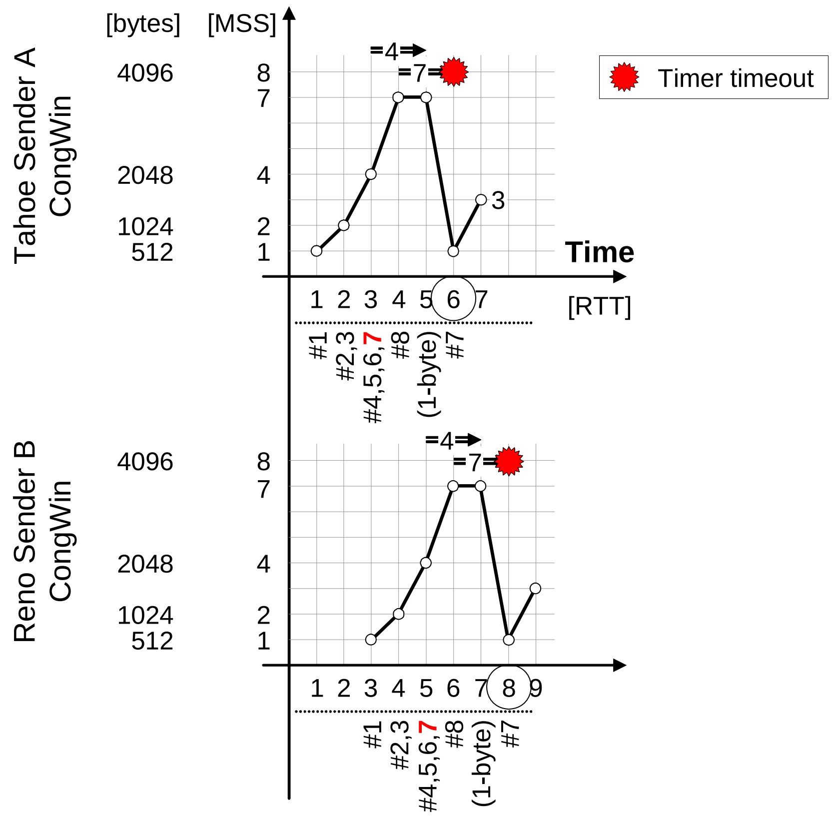

The *A*'s timer times out at time = 6×RTT (before three dupACKs are received), and *A* re-sends segment #7 and enters slow start. At 7×RTT the cumulative ACK for both segments #7 and #8 is received by *A* and it, therefore, increases `CongWin` = 1 + 2 = 3×`MSS`, but there is nothing left to send.

The Reno sender, *B*, behaves in the same way as the Tahoe sender because the loss is detected by the expired RTO timer. Both types of senders enter slow start after timer expiration. (Recall that the difference between the two types of senders is only in Reno implementing *fast recovery*, which takes place after three dupACKs are received.)

Problem 2.5 — Solution

The range of congestion window sizes is [1, 16]. Because the loss is detected when `CongWindow` = 16×`MSS`, `SSThresh` is set to 8×`MSS`. Thus, the congestion window sizes in consecutive transmission rounds are: 1, 2, 4, 8, 9, 10, 11, 12, 13, 14, 15, and 16 `MSS` (see the figure below). This averages to 9.58×`MSS` per second (recall, a transmission round is RTT = 1 sec), and a mean utilization of $\frac{9.58 \times 8}{128}$ [Kbps/Kbps] = 0.59875, or about 60%.

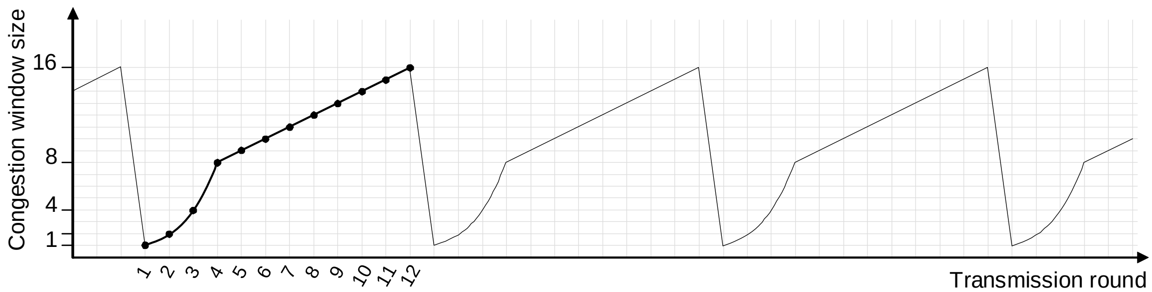

Problem 2.6 — Solution

The solution of Problem2.5 is an idealization that cannot occur in reality. A better approximation is as follows. The event sequence develops as follows:

packet ***loss happens at a router*** (last transmitted segment), current `CongWin` = 16×`MSS`.

the sender receives 16-1=15 ACKs which is not enough to grow `CongWin` to 17

but it still sends 16 new segments, last one will be lost

the sender receives 15 dupACKs, ***loss detected at the sender***

retransmit the oldest outstanding packet, `CongWin` ← 1

the sender receives cumulative ACK for 16 recent segments, except for the last one

`CongWin` ← 2, `FlightSize` = 1×`MSS`, send one new segment

the sender receives 2 dupACKs, `FlightSize` = 3×`MSS`, `EffctWin` = 0, sends one 1-byte segment

the sender receives 3rd dupACK, retransmits the oldest outstanding packet, CW ← 1

the sender receives cumulative ACK for 4 recent segments (one of them was 1-byte), `FlightSize` ← 0

`CongWin` ← 2, the sender resumes slow start

Problem 2.7 — Solution

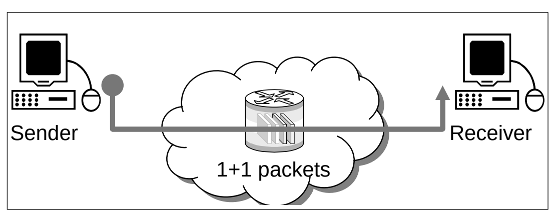

`MSS` = 512 bytes

`SSThresh` = 3×`MSS`

`RcvBuffer` = 2 KB = 4×`MSS`

`TimeoutInterval` = 3×RTT

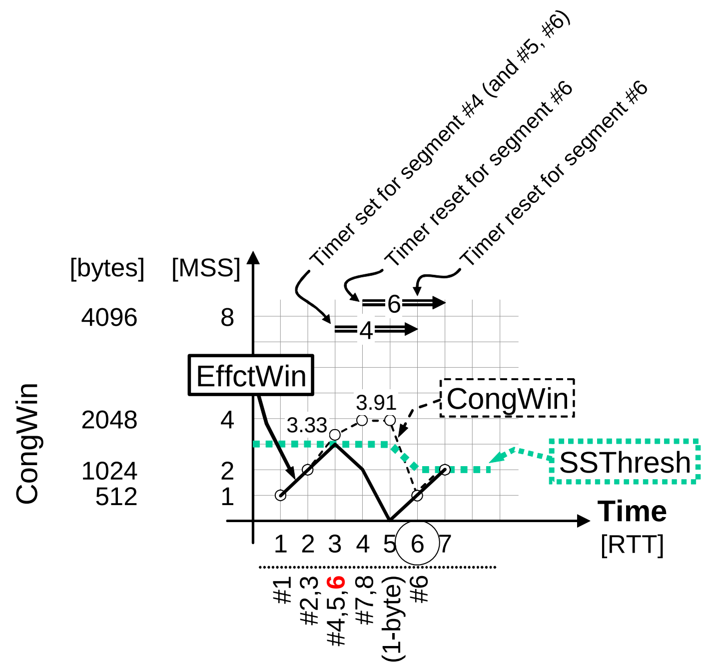

At time = 3×RTT, after receiving the acknowledgement for the 2nd segment, the sender's congestion window size reaches `SSThresh` and the sender enters additive increase mode. Therefore, the ACK for the 3rd segment is worth $\text{MSS} \times \dfrac{\text{MSS}}{\text{CongWindow}_{t\text{-}1}} = 1 \times \dfrac{1}{3} =$ 0.33 MSS. $\texttt{CongWindow}_{3\times\text{RTT}}$ = 3.33×`MSS`. Therefore, the sender sends 3 segments: 4th, 5th, and 6th. The router receiver 3 segments but can hold only 1+1=2. It will start transmitting the 4th segment, store the 5th segment in its buffer, and discard the 6th segment due to the lack of buffer space (there is a waiting room for one packet only). The acknowledgements for the 4th and 5th segments add $1 \times \dfrac{1}{3.33} = 0.3 \times \texttt{MSS}$ and $1 \times \dfrac{1}{3.63} = 0.28 \times \texttt{MSS}$, respectively, so $\texttt{CongWindow}_{4\times\text{RTT}}$ = 3.91×`MSS`. The effective window is smaller by one segment because the 6th segment is outstanding:

$$\texttt{EffctWin} = \lfloor \min\{\texttt{CongWin}, \texttt{RcvWindow}\} - \texttt{FlightSize} \rfloor = \lfloor \min\{3.91, 4\} - 1 \rfloor = 2\times\text{MSS}$$

At time = 4×RTT the sender sends two segments, both of which successfully reach the receiver. Acknowledgements for 7th and 8th segments are duplicate ACKs, but this makes only two duplicates so far, so the loss of #6 is still not detected. Notice that at this time the receive buffer stores two segments (`RcvBuffer` = 2 KB = 4 segments), so the receiver starts advertising `RcvWindow` = 1 Kbytes = 2 segments.

The sender computes

$$\texttt{EffctWin} = \lfloor \min\{\texttt{CongWin}, \texttt{RcvWindow}\} - \texttt{FlightSize} \rfloor = \lfloor \min\{3.91, 2\} - 3 \rfloor = 0$$

so it sends a 1-byte segment at time = 5×RTT.

At time = 6×RTT, the loss of the 6th segment is detected *via three duplicate ACKs*. Recall that the sender in fast retransmit does not use the above formula to determine the current `EffctWin`—it simply retransmits the segment that is suspected lost. That is why the above figure shows

`EffctWin` = 2 at time = 6×RTT. The last `EffctWin`, at time = 7×RTT, equals 2×`MSS` but there are no more data left to send.

Therefore, the answers are:

(a)

The first loss (segment #6 is lost in the router) happens at 3×RTT, so

$\texttt{CongWindow}_{3\times\text{RTT}} = 3.33\times\texttt{MSS}$.

(b)

The loss of the 6th segment is detected via three duplicate ACKs at time = 6×RTT.

At this time, not-yet-acknowledged segments are: 6th, 7th, and 8th, a total of three.

Problem 2.8 — Solution

In solving the problem, we should keep in mind that the receive buffer size is set relatively small to 2Kbytes = 8×`MSS`.

In the transmission round *i*, the sender sent segments *k*, *k*+1, …, *k*+7, of which the segment *k*+3 is lost. The receiver receives the four segments *k*+4, …, *k*+7, as out-of-order and buffers them and sends back four duplicate acknowledgements. In addition, the receiver notifies the sender that the new `RcvWindow` = 1 Kbytes = 4×`MSS`.

At *i*+1, the sender first receives three regular (non-duplicate!) acknowledgements for the first three successfully transferred segments, so `CongWin` = 11×`MSS`. Then, four duplicate acknowledgements will arrive while `FlightSize` = 5. After receiving the first three dupACKs, Reno sender reduces the congestion window size by half, $\texttt{CongWin} = \lfloor 11 / 2 \rfloor = 5\times\texttt{MSS}$.

The new value of $\texttt{SSThresh} = \lfloor \texttt{CongWin} / 2 \rfloor + 3\times\texttt{MSS} = 8\times\texttt{MSS}$.

Because Reno sender enters *fast recovery*, each dupACK received after the first three increment the congestion window by one `MSS`. Therefore, `CongWin` = **6**×`MSS`. The effective window is:

$$\texttt{EffctWin} = \min\{\texttt{CongWin}, \texttt{RcvWindow}\} - \texttt{FlightSize} = \min\{6, 4\} - 5 = -1 \quad (\#)$$

Thus, the sender is allowed to send nothing but the oldest unacknowledged segment, *k*+3, which is suspected lost.

There is an interesting observation to make here, as follows. Knowing that the receive buffer holds the four out-of-order segments and it has four more slots free, it may seem inappropriate to use the formula (#) above to determine the effective window size. After all, there are four *free* slots in the receive buffer, so that should not be the limiting parameter! The sender's current knowledge of the network tells it that the congestion window size is 6×`MSS` so this should allow sending more!? Read on.

The reason that the formula (#) is correct is that you and I know what receiver holds and where the unaccounted segments are currently residing. But the sender does not know this! It only knows that currently `RcvWindow` = 4×`MSS` and there are five segments *somewhere* in the network. As far as the sender knows, they still may show up at the receiver. So, it must not send anything else.

At i+2, the sender receives ACK asking for segment k+8, which means that all five outstanding segments are acknowledged at once. Because the congestion window size is still below the `SSThresh`, the sender increases `CongWin` by 5 to obtain `CongWin` = 11×`MSS`. Notice that by now the receiver notifies the sender that the new `RcvWindow` = 2 Kbytes = 8×`MSS`, because all the receive buffer space freed up.

The new effective window is:

$$\texttt{EffctWin} = \min\{\texttt{CongWin}, \texttt{RcvWindow}\} - \texttt{FlightSize} = \min\{11, 8\} - 0 = 8\times\texttt{MSS}$$

so the sender sends the next eight segments, k+8, …, k+15.

Next time the sender receives ACKs, it's already in *congestion avoidance* state, so it increments `CongWin` by 1 in every transmission round (per one RTT).

Notice that, although `CongWin` keeps increasing, the sender will keep sending only **eight** segments per transmission round because of the receive buffer's space limitation.

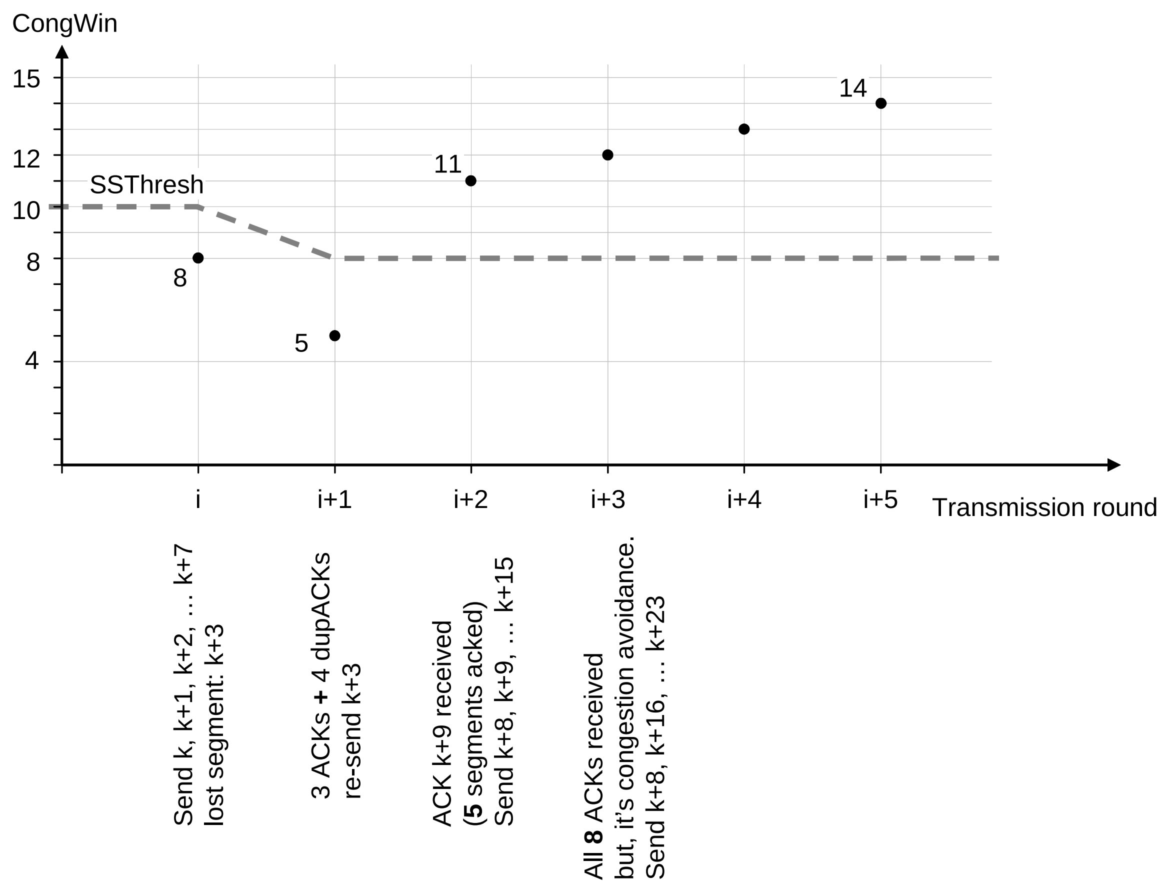

Some networking books give a simplified formula for computing the slow-start threshold size after a loss is detected as `SSThresh` = `CongWin` / 2 = 5.5×`MSS`. Rounding `CongWin` down to the next integer multiple of `MSS` is often not mentioned and neither is the property of fast recovery to increment `CongWin` by one `MSS` for each dupACK received after the first three that triggered the retransmission of segment *k*+3.

In this case, the sender in our example would immediately enter congestion avoidance, and the corresponding diagram is as shown in the figure below.

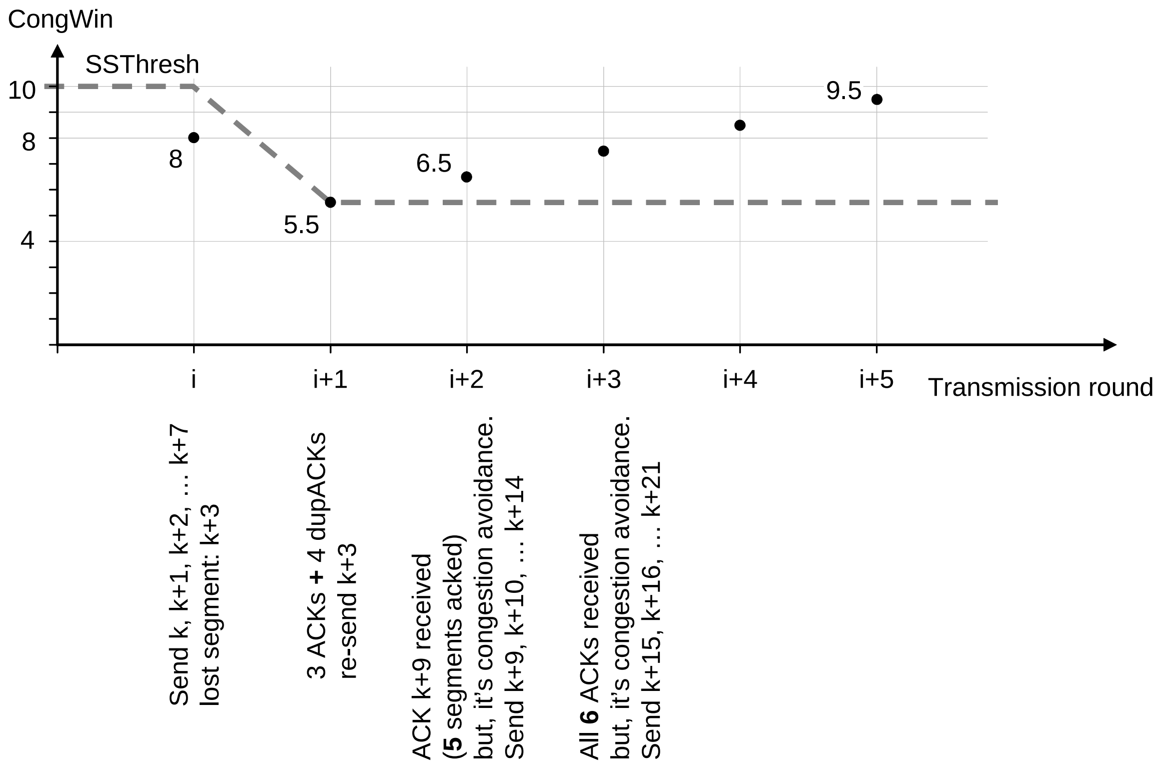

Problem 2.9 — Solution

We can ignore the propagation times because they are negligible relative to the packet transmission times (mainly due to short the distance between the transmitter and the receiver). Also, the transmission times for the acknowledgements can be ignored. Because the transmission just started, the sender is in the slow start state. Assuming that the receiver sends only *cumulative acknowledgements*, the total time to transmit the first 15 segments of data is (see the figure):

$$4 \times 0.8 + 15 \times 8 = 123.2 \text{ ms.}$$

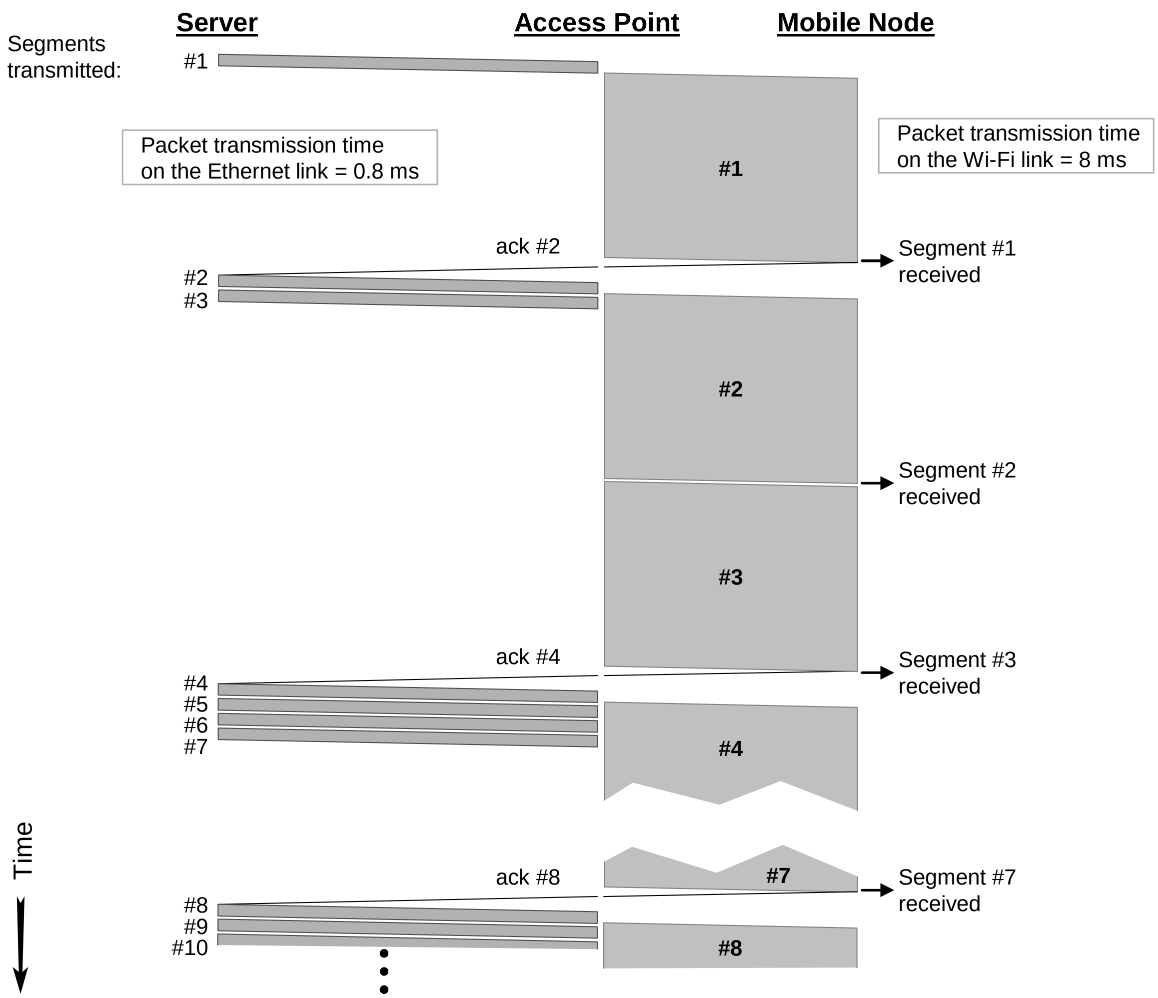

The timing diagram is as shown in the figure.

Problem 2.10 — Solution

(a)

During the slow start phase, an incoming acknowledgment will allow the sender to send two segments. The acknowledgements to these two segments will be separated by 1 ms, as shown in the figure below.

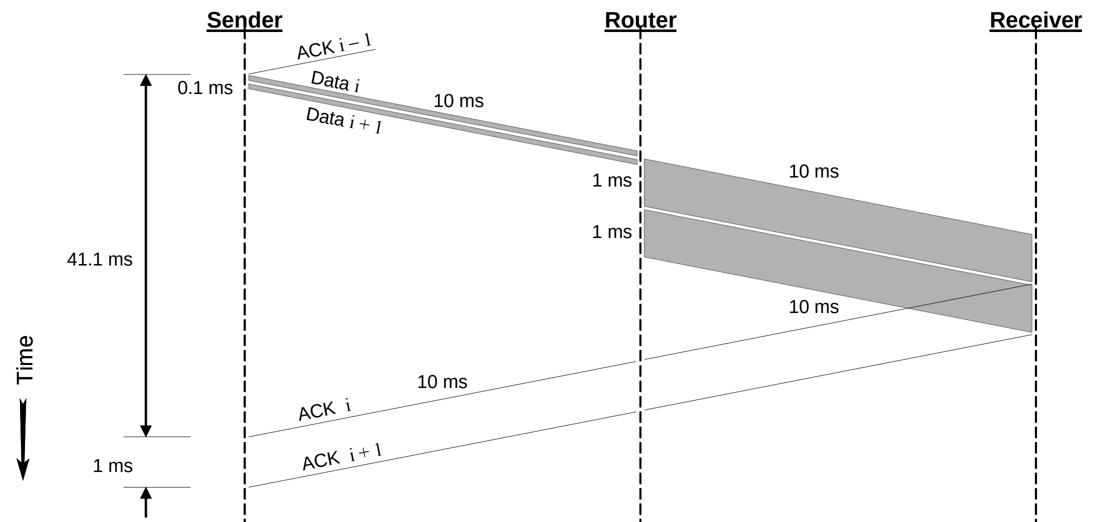

(b)

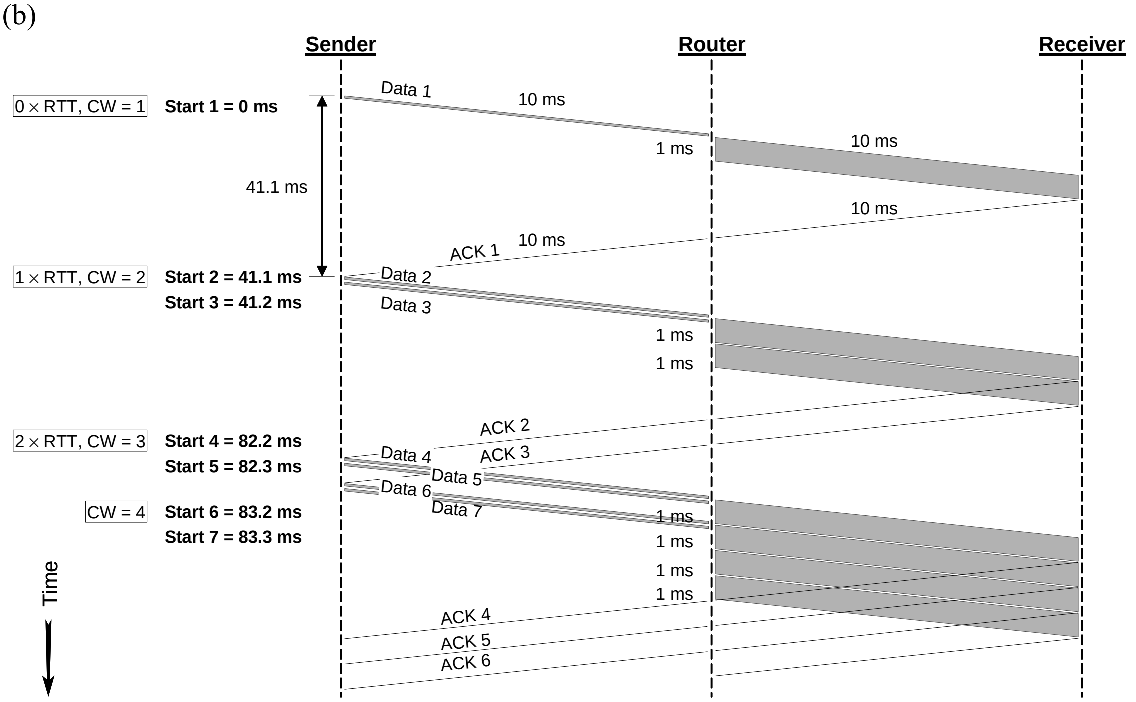

The start times for the transmissions of the first seven segments are as shown in the figure:

(c)

We can make the following observations:

o1. Except for the first round, packets always arrive in pairs at the router. The reason for this was explained under item (a); that is, during the slow start phase, each acknowledgement will allow the sender to send two segments back-to-back.

o2. The time gap between the packet pairs is 1 ms, because the time gap between the acknowledgements (during slow start) is 1 ms. This gives enough time to the router to transmit one packet before the arrival of the next pair.

o3. We can think conceptually that from each pair, one packet is relayed by the router and the other remains in the router's buffer. (This is *not* true, because the router transmits packets in the order of their arrival (FIFO), so it will first transmit any packets remaining from previous pairs, but we can think this way conceptually.)

o4. Because router can hold 9+1=10 packets, the first loss will happen when ≥ 20 packets is sent on one round. In 5th round, time = 5×RTT, the `CongWin` = 32, so this is when the first loss will happen

o5. The packets sent in this round will find the following number of packets already at the router (packet pairs are separated by ||):

0, 1 || 1, 2 || 2, 3 || 3, 4 || 4, 5 || 5, 6 || 6, 7 || 7, 8 || 8, 9 || 9, 10 || 9, 10 || 9, 10 || 9, 10 || 9, 10 || 9, 10 || 9, 10 .

o6. Therefore, the 20th packet of this round will find 10 packets already at the router and this packet will be lost. This is the 41st packet from the start of sending at time = 0.

o7. By the time the next pair arrives, the router will have transmitted one packet, so 21st packet finds 9 packets already at the router, but its companion in the pair, 22nd packet finds 10 packets and is lost

o8. This pattern repeats until the last, which is 32nd packet of this round.

o9. A total of 7 packets will be lost, starting with 20th, 22nd, 24th, …, and 32nd.

o10. At time, 6×RTT, the congestion window will grow up to 32 + 19 = 51

o11. After this ≥ 3 × dupACKs will be received and the sender will go into the multiplicative decrease phase

Therefore, the congestion window sizes for the first 6×RTT are: 1, 2, 4, 8, 16, 31, 51.

(d)

As shown in (c), the first packet will be lost in the 5th round, and it is the 20th packet of this round. Its ordinal number is #51, determined as follows:
1 + 2 + 4 + 8 + 16 = 31 (from previous four rounds) + 20 (from the 5th round) = 51

(e)

…

Problem 2.11 — Solution

Transmission delay for all three scenarios is:

$$t_x = \frac{\text{packet length}}{\text{bandwidth}} = \frac{8192 \text{ bits}}{1000000 \text{ bits per second}} = 8.192 \text{ ms}$$

In the first scenario (RTT_1 = 0.01 sec), the round-trip time is about the same as transmission delay. The sender can send up to one segment and start with a second one before it receives the acknowledgement for the first segment. Conversely, in the third scenario (RTT_3 = 1 sec), the sender can send a burst of up to 122 segments before receiving the acknowledgement for the first segment of this burst.

We ignore the initial slow-start phase and consider that the sender settled into a congestion-avoidance state. Because the network exhibits periodic behavior where every tenth packet is lost, we need to determine the sender's cyclic behavior. The figure illustrates the solution and the description follows below the figure.

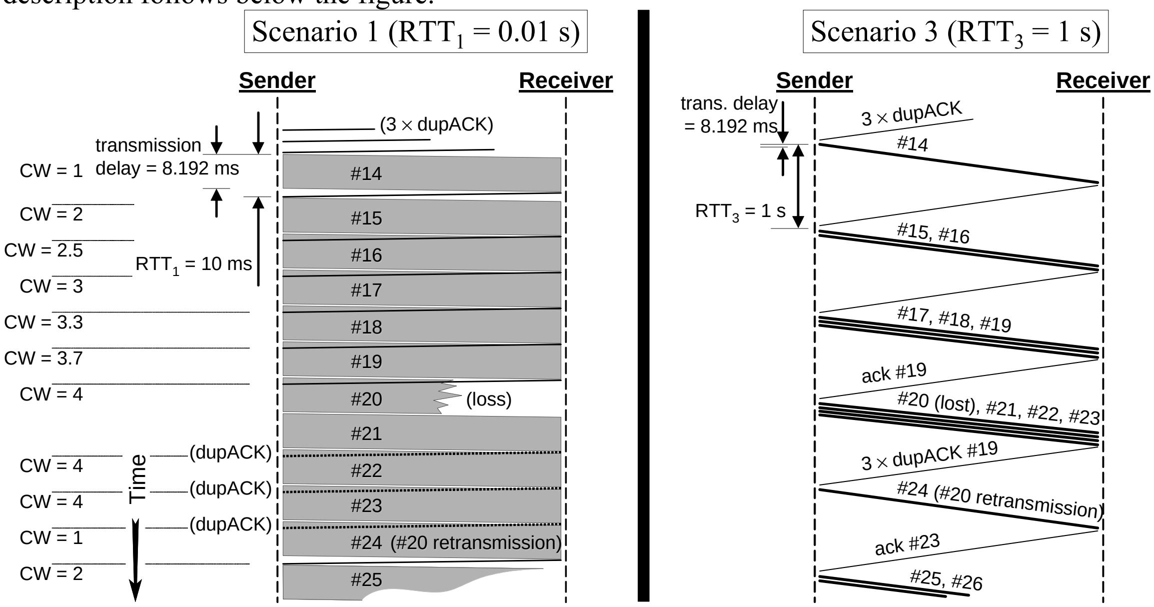

To detect a lost 10^{th} segment, the sender must have sent at least three more to receive three duplicate acknowledgements. Let us assume the sender detects the loss (via 3 dupACKs) after it has sent 13 segments. The 14^{th} transmission will be retransmission of the 10^{th} segment (the lost one). The receiver fills the gap and cumulatively acknowledges up to the segment #13, requesting the sender to send the next segment (#14). This table shows how the first scenario sender sends segments and receives acknowledgements:

packet number	10^{th}	11^{th}	12^{th}	13^{th}	14^{th}	15^{th}	16^{th}	17^{th}	18^{th}	19^{th}	20^{th}	21^{st}	22^{nd}	23^{rd}	24^{th}	25^{th}
segment seq. num.	10 (×)	11	12	13	10	14	15	16	17	18	19 (×)	20	21	22	19	23
ACK	_	**10**	**10**	**10**	14	15	16	17	18	19	_	**19**	**19**	**19**	23	24
CongWin				**1**	2	2.5	3	3.3	3.7	3.7		**3.7**	**3.7**	**1**	2	2.5

This cycle will repeat. The sender is sending segments back-to-back. The only overhead it experiences is that every tenth packet is lost, which implies that during 10 transmission delays it successfully sends 9 segments. Hence, the average data rate the sender achieves is: (9/10) × 1 Mbps = 900 Kbps.

Let us now consider the third scenario with RTT_3 = 1 sec. Because the round-trip time is much longer than transmission time and every 10^{th} packet is lost, we can consider that burst transmissions are clocked to RTT intervals.

RTT		n–1			n	n+1		n+2			n+3				n+4
packet number	10^{th}	11^{th}	12^{th}	13^{th}	14^{th}	15^{th}	16^{th}	17^{th}	18^{th}	19^{th}	20^{th}	21^{st}	22^{nd}	23^{rd}	24^{th}
segment seq. num.	10 (×)	11	12	13	10	14	15	16	17	18	19 (×)	20	21	22	19
ACK	_	**10**	**10**	**10**	14	15	16	17	18	19	_	**19**	**19**	**19**	23
CongWin	_			1	2		2.5, 3			3.3, 3.7, 4				1	2

In this scenario, the cycle lasts 4×RTT and the sender successfully sends 9 segments of which 6 are acknowledged and 3 are buffered at the receiver and will be acknowledged in the next cycle. Hence, the average data rate the sender achieves is:

$$\frac{9 \times t_x}{4 \times \text{RTT}} \times 1\,\text{Mbps} = \frac{0.073728}{4} \times 1000000 = 18432\text{ bps} \approx 18.4\text{ Kbps}$$

Obviously, the sender in the third scenario achieves much lower rate than the one in the first scenario. The reason for this is that the first-scenario sender receives feedback quickly and reacts quickly. Conversely, the third-scenario sender receives feedback very slowly and accordingly reacts slowly—it is simply unable to reach the potentially achievable rate.

Problem 2.12 — Solution

Object size,	O = 1 MB = 2^{20} bytes = 1048576 bytes = 8388608 bits
Segment size `MSS` = 1 KB,	S = `MSS` × 8 bits = 8192 bits
Round-trip time	RTT = 100 ms
Transmission rate,	R = as given for each individual case (see below)

There are a total of $L = \dfrac{1\,MB}{1\,KB} = \dfrac{2^{20}}{2^{10}} = 2^{10} = 1024$ segments (packets) to transmit.

(a)

Bottleneck bandwidth, R = 1.5 Mbps, data sent continuously:

$$\frac{O}{R} = \frac{2^{20} \times 8}{1.5 \times 10^6} = \frac{1048576 \times 8}{1500000} = \frac{8388608}{1500000} = 5.59\ \text{sec}$$

latency = 2 × RTT + O / R = 200 × 10^{-3} + 5.59 sec = 5.79 sec

(b)

R = 1.5 Mbps, Stop-and-wait

latency = $2 \times \text{RTT} + \left(\frac{S}{R} + \text{RTT}\right) \times L = 0.2 + \left(\frac{8192}{1500000} + 0.1\right) \times 1024$ = 108.19 sec

(c)

R = ∞, Go-back-20

Because transmission time is assumed equal to zero, all 20 packets will be transmitted instantaneously and then the sender waits for the ACK for all twenty. Thus, data will be sent in chunks of 20 packets:

latency = $2 \times \text{RTT} + \text{RTT} \times \frac{L}{20}$ = 0.2 + 5.12 = 5.22 sec

(d)

$R = \infty$, TCP Tahoe

Because the transmission is error-free and the bandwidth is infinite, there will be no loss, so the congestion window will grow exponentially (*slow start*) until it reaches the slow start threshold `SSThresh` = 65535 bytes = 64 × `MSS`, which is the default value. From there on it will grow linearly (*additive increase*). Therefore, the sequence of congestion window sizes is as follows:

Congestion window sizes:

$$\underbrace{1,\ 2,\ 4,\ 8,\ 16,\ 32,\ 64,}_{\text{Slow start -- total 7 bursts}}\quad \underbrace{65,\ 66,\ 67,\ 68,\ 69,\ 70,\ \ldots}_{\text{Additive increase -- total 13 bursts}}$$

Then, slow start phase consists of 7 bursts, which will transfer the first 127 packets. The additive increase for the remaining 1024 − 127 = 897 packets consists of at most 897/64 ≈ 14 bursts. Quick calculation gives the following answer:

Assume there will be thirteen bursts during the additive increase.

With a constant window of 64 × `MSS`, this gives 13 × 64 = 832.

On the other hand, additive increase adds 1 for each burst, so starting from 1 this gives $1+2+3+ \ldots + 13 = \frac{13\times(13+1)}{2} = 91$ packets.

Therefore, starting with the congestion window size of 64 × `MSS`, the sender can in 13 bursts send up to a total of 832 + 91 = 923 packets, which is more than 832.

Finally, sender needs total of 7 + 13 = 20 bursts:

latency = 2 × RTT + 20 × RTT = 2.2 sec.

Problem 2.13 — Solution

It is easier to solve this problem if we represent the relevant parameters in a tabular form as in the table below (see also the figure below).

First, we know that this is a TCP Reno sender, currently in the slow start phase, with `CongWin`(t_i) = 400 bytes. Given that `MSS` = 200 bytes, the sender sends a "burst" of two segments back-to-back. According to Listing 2-1 (summarized on page 147), when sending a segment, the sender needs to set the RTO timer if it is not already running. The RTO timer is not running at t_i because all previous segments were successfully acknowledged. So, using equation (2.2), the sender obtains:

`TimeoutInterval`(t_i) = `EstimatedRTT`(t_i) + 4 · `DevRTT`(t_i) = 100.8 + 4×9 = 136.8 ms

where `EstimatedRTT`(t_i) and `DevRTT`(t_i) are determined as follows (using default values for α = 0.125 and β = 0.25):

`EstimatedRTT`(t_i) = $(1-\alpha)$ · `EstimatedRTT`(t_{i-1}) + α · `SampleRTT`(t_i) = $0.875 \times 100 + 0.125 \times 106 = 100.8$ ms

`DevRTT`(t_i) = $(1-\beta)$ · `DevRTT`(t_{i-1}) + β · |`SampleRTT`(t_i) − `EstimatedRTT`(t_{i-1})| = $0.75 \times 10 + 0.25 \times |106 - 100| = 9$ ms

The sender does not modify the RTO timer when sending the second segment, because the timer is already running.

Notice that the top row of the table shows the *times when each segment is* ***transmitted***. Although the problem statement does not require determining these, we can easily determine them as follows. We assume for simplicity that $t_i = 0$ ms. Given that $t_x << t_p$ and the propagation times are in the range from 50 ms to 100 ms, we can assume for simplicity that the segment transmission time equals $t_x = 1$ ms. However, assuming any other small number for t_x would not make a difference for the problem solution.

The second segment is sent right after the first one, so $t_{i+1} = t_i + 1 = 1$ ms. Notice that at t_{i+1} there is no acknowledgement that arrived, so the sender does not calculate `EstimatedRTT`(t_{i+1}) and `DevRTT`(t_{i+1}) and these entries are left empty in the table.

Time t [ms]	t_i (=0 ms)	t_{i+1} (=1 ms)	t_{i+2} (=105)	t_{i+3} (=106)	t_{i+4} (=107)	t_{i+5} (=108)	t_{i+6} (=237.5)	t_{i+7} (=288)	$t_{i+7}+t_x$ (=289)
CongWin(t)	2×`MSS`		4×`MSS`				1×`MSS`	3×`MSS`	
SSThresh(t)	64 KBytes						2×`MSS`		
RTT(t)	105	93	179	182	165	193	154	171	
SampRTT(t)	106		105						
EstimRTT(t)	100.8		101.3						
DevRTT(t)	9		7.8						
TimeoutInt	136.8 ms		132.5				265		
Transmiss'n type	Burst #1		Burst #2				Re-transmit	Burst #3	

The sender then waits idle until an acknowledgement arrives for the two segments that it sent.

The ACK for the first segment (shown in the figure as ACK #3) will arrive at the time $t_i +$ `RTT`$(t_i) = 0 + 105 = 105$ ms.

The ACK for the second segment (shown in the figure as ACK #1 - duplicate) will arrive at time t_{i+1} + `RTT`$(t_{i+1}) = 1 + 93 = 94$ ms. This is earlier than the ACK for the first segment by 11 ms and we assume that the second segment also arrived before the first one to the receiver. Because segment #2 is out of order, the receiver will immediately send a duplicate ACK, asking again for segment #1. The sender will do nothing when it receives the duplicate ACK. (Recall that the TCP sender retransmits only when it receives ≥3×dupACKs, and it does not sample RTT for duplicate ACKs.)

When the first segment arrives, the receiver will acknowledge both outstanding segments by sending a *cumulative* ACK #3. When the sender receives ACK #3 at 105 ms, it will increment its congestion window by 2×`MSS` and send four more segments at times t_{i+2}, $t_{i+3}=t_{i+2}+t_x$, $t_{i+4}=t_{i+3}+t_x$, and $t_{i+5}=t_{i+4}+t_x$.

At the time of arrival of ACK #3, the sender measures the sample RTT. This time corresponds to t_{i+2} and `SampleRTT`(t_{i+2}) = `RTT`(t_i) = 105 ms. The sender calculates `EstimatedRTT`(t_{i+2}) and `DevRTT`(t_{i+2}) as shown in the table.

Because ACK #3 acknowledged all the outstanding segments and the RTO timer was reset, the sender will calculate the new value `TimeoutInterval`(t_{i+2}) = 132.5 ms and set the RTO timer when it sends the third segment at t_{i+2} = 105 ms. Three more segments will be sent in this burst at times t_{i+3}=106 ms, t_{i+4}=107 ms, and t_{i+5}=108 ms.

Given that `RTT`(t_{i+4})=165 ms is the shortest for all the segments from the second burst, we will again assume that segment #5 arrived out of order and the receiver will immediately send a duplicate ACK (shown in the figure as ACK #3 - duplicate). The acknowledgements for the four segments of the second burst will arrive at these times (see also the figure):

For segment #5 (ACK #3 - duplicate) arrives at: t_{i+4} + `RTT`(t_{i+4}) = 107 + 165 = 272 ms.

For segment #3 (ACK #4) arrives at: t_{i+2} + `RTT`(t_{i+2}) = 105 + 179 = 284 ms.

For segments #4 and #5 (ACK #6) arrives at: t_{i+3} + `RTT`(t_{i+3}) = 106 + 182 = 288 ms.

For segment #6 (ACK #7) arrives at: t_{i+5} + `RTT`(t_{i+5}) = 108 + 193 = 301 ms.

All of the above times are too late for the RTO timer, which will expire at time t_{i+6} = t_{i+2} + `TimeoutInterval`(t_{i+2}) = 105 + 132.5 = 237.5 ms. Notice that both TCP Tahoe and TCP Reno behave the same on RTO timer expiration. They both enter slow start, so there will be no difference in behavior and `CongWin`(t_{i+6}) = 1×`MSS` = 200 bytes. Also, according to equation (2.1), the timeout interval is set to twice the previous value, so `TimeoutInterval`(t_{i+6}) = 2 × `TimeoutInterval`(t_{i+2}) = 265 ms.

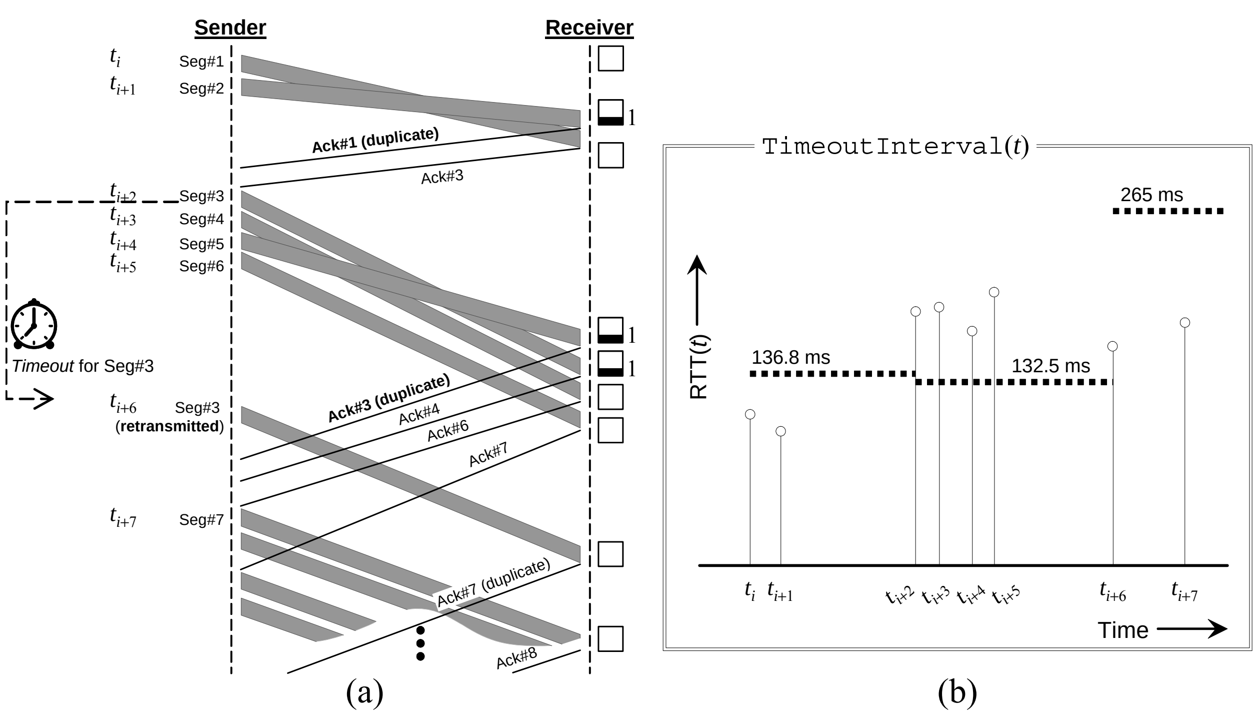

Because of the expired RTO timer, the sender will retransmit the oldest outstanding segment at time t_{i+6} = 237.5 ms. This segment is a copy of the third segment, but we count it as the seventh transmission. Given that RTT(t_{i+6}) = 154 ms, the corresponding acknowledgement will arrive at the time t_{i+6} + RTT(t_{i+6}) = 237.5 + 154 = 391.5 ms.

Meanwhile, the duplicate ACK #3 will arrive and the sender will ignore it. Next, the ACK for segment #3 (shown in the figure as ACK #4) will arrive at time equal 284 ms. Notice that this is the ACK for the original segment #3, not the retransmitted one. Although at this time segment #5 is already at the receiver, it is still out of order: there is a gap because segment #4 still did not arrive, so the receiver is asking for #4. The sender thinks that the receiver is acknowledging the retransmitted segment, so it increments `CongWin` to 2×`MSS`. There are still three outstanding segments (#4, #5 and #6), so `EffectiveWindow` = `CongWindow` − `FlightSize` = 0 and the sender remains quiet. Notice also that now `CongWindow` reached `SSThresh`, so the sender enters *congestion avoidance*.

At the time equal 288 ms, ACK #6 arrives acknowledging segments #4 and #5. Because the sender is in congestion avoidance and `CongWin` = 2×`MSS`, the two-segment ACK will count as one by equation (2.5), so `CongWin` becomes 3×`MSS`. There is one more segment unaccounted (#6), so `FlightSize` = 1 and `EffectiveWindow` = 2 and the sender transmits two segments at times t_{i+7} and $t_{i+7} + t_x$.

(a)
The table below shows the congestion window sizes `CongWin`(t) and the sequence numbers of the segments transmitted from the sender. (Recall that the sender's sequence number for the segment that will be transmitted at t_i equals 30 and `MSS` = 200 bytes.)

Time t [ms]	t_i (=0 ms)	t_{i+1} (=1 ms)	t_{i+2} (=105)	t_{i+3} (=106)	t_{i+4} (=107)	t_{i+5} (=108)	t_{i+6} (=237.5)	t_{i+7} (=288)	$t_{i+7}+t_x$ (=289)
CongWin(t)	2×`MSS`		4×`MSS`				1×`MSS`	3×`MSS`	
SeqNum(t)	30	230	430	630	830	1030	**430**	1230	1430

(b)
To determine the times when the ACKs will arrive, we simply add the corresponding RTT to the time the segment departed, as was already done above for most of the segments. The table below shows the times of ACK arrivals, the sequence numbers of the ACKs, and the values of `RcvWindow` as carried in each ACK segment (recall that the receiver's buffer size `RcvWindow`(t_i) = 1000 bytes and also see how the figure above indicates the receive buffer occupancy).

Time t [ms]	t_i (=0 ms)	t_{i+1} (=1 ms)	t_{i+2} (=105)	t_{i+3} (=106)	t_{i+4} (=107)	t_{i+5} (=108)	t_{i+6} (=237.5)	t_{i+7} (=288)	$t_{i+7}+t_x$ (=289)
ACK arrives	105 ms	94 ms	284	288	278	301	391.5	459	
ACKSeqNo	430	30/**dup**	630	1030	430/**dup**	1230	1230/**dup**	1430	
RcvWindow	1000	800	800	1000	800	1000	1000	1000	

(c)
The values of `EstimatedRTT`(t) and `DevRTT`(t) are shown in the first table above. Recall that the TCP retransmission-timer management algorithm measures `SampleRTT` for segments that have been transmitted *once* and *not* for segments that have been retransmitted. It also ignores dupACKs. An extract from the first table shows that `EstimatedRTT`(t) and `DevRTT`(t) are calculated only at t_i and t_{i+2}:

Time t [ms]	t_i (=0 ms)	t_{i+1} (=1 ms)	t_{i+2} (=105)	t_{i+3} (=106)	t_{i+4} (=107)	t_{i+5} (=108)	t_{i+6} (=237.5)	t_{i+7} (=288)	$t_{i+7}+t_x$ (=289)
RTT(t)	105	93	179	182	165	193	154	171	
SampRTT(t)	106		105						
EstimRTT(t)	100.8		101.3						
DevRTT(t)	9		7.8						

(d)
The TCP sender will set its retransmission timer three times during the considered interval, and the values of `TimeoutInterval`(t) are as follows:

Time t [ms]	t_i (=0 ms)	t_{i+1} (=1 ms)	t_{i+2} (=105)	t_{i+3} (=106)	t_{i+4} (=107)	t_{i+5} (=108)	t_{i+6} (=237.5)	t_{i+7} (=288)	$t_{i+7}+t_x$ (=289)
TimeoutInt	136.8 ms		132.5				265		

Problem 2.14 — Solution

Problem 2.15 — Solution

During the specified period of time, the sender will receive only 5 acknowledgements for new data. This will happen at times: at 89 for segment #33 (which was transmitted earlier, at time t = 82); at 96 for the retransmitted segment #35; at 104 for segment #37; at 113 for segment #39; and at 120 for segment #41. Check the pseudocode at the end of Section 2.1.2 to see that all other ACKs are ignored because they are acknowledging already acknowledged data (they are called "duplicate ACKs").

By examining Figure 2-19, we can see that the retransmitted segments #33, #35, and #37 arrive at an idle router, so they will be transmitted first and their measured RTT values equal 6. However, segment #39 will find segment #73 in front of it at the router, and segment #41 will find segments #80 and #81 in front of it at the router. The measured values of `SampleRTT`(t) as read from Figure 2-19 are:

`SampleRTT`(89) = 6; `SampleRTT`(96) = 6; `SampleRTT`(104) = 6; `SampleRTT`(113) = 7; `SampleRTT`(120) = 8;

The calculation of `TimeoutInterval`(t) is straightforward using Eq. (2.2).

Problem 3.1 — Solution

Problem 3.2 — Solution

Notice that the first packet is sent at 20 ms, so its playout time is 20 + 210 = 230 ms. The playout of all subsequent packets are spaced apart by 20 ms (unless a packet arrives too late and is discarded).

Notice also that the packets are labeled by sequence numbers. Therefore, although packet #6 arrives before packet #5, it can be scheduled for playout in its correct order.

Packet sequence number	Arrival time r_i [ms]	Playout time p_i [ms]
#1	195	230
#2	245	250
#3	270	270
#4	295	discarded (>290)
#6	300	330
#5	310	310
#7	340	350
#8	380	discarded (>370)
#9	385	390
#10	405	410

The playout schedule is also illustrated in this figure:

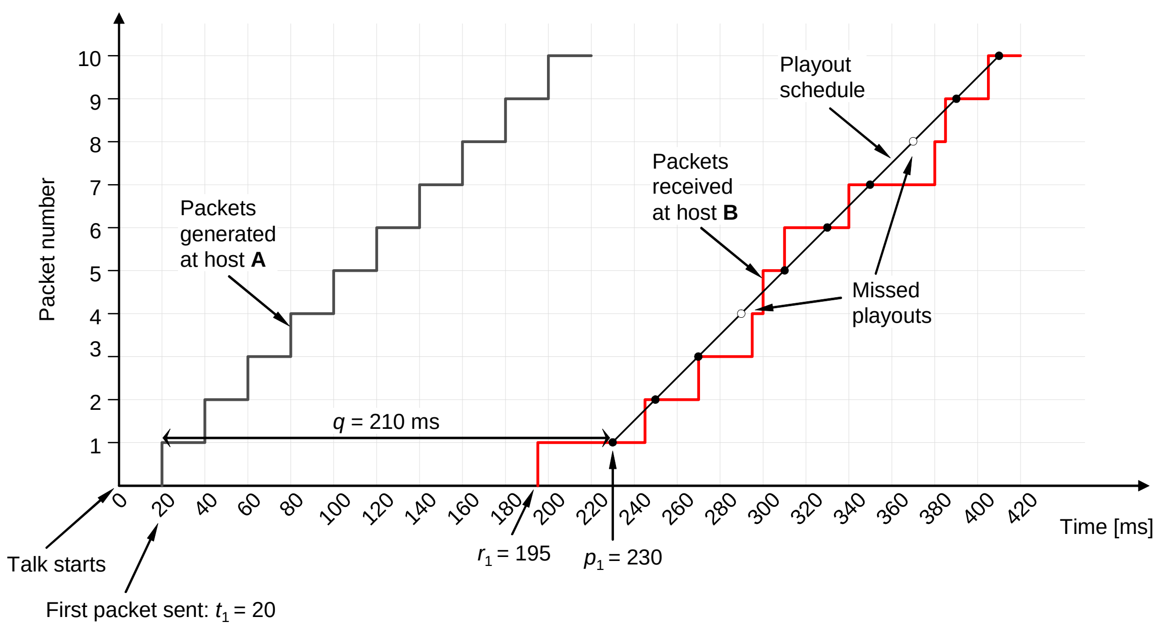

Problem 3.3 — Solution

(a)

Packet sequence number	Arrival time r_i [ms]	Playout time p_i [ms]

#1	95	170
#2	145	190
#3	170	210
#4	135	230
#6	160	250
#5	275	discarded (>270)
#7	280	290
#8	220	310
#9	285	330
#10	305	350

(b)

The minimum propagation delay given in the problem statement is 50 ms. Hence, the maximum a packet can be delay for playout is 100 ms. Because the source generates a packet every 20 ms, the maximum number of packets that can arrive during this period is 5. Therefore, the required size of memory buffer at the destination is 6 × 160 bytes = 960 bytes. (The buffer should be able to hold 6 packets, rather than 5, because I assume that the last arriving packet is first buffered and then the earliest one is removed from the buffer and played out.)

Problem 3.4 — Solution

The length of time from when the first packet in this talk spurt is generated until it is played out is:

$$q_k = \hat{\delta}_k + K \cdot \hat{\upsilon}_k = 90 + 4 \times 15 = 150 \text{ ms}$$

The playout times for the packets including k+9th are obtained by adding this amount to their timestamp, *because they all belong to the same talk spurt*. Notice that the k+5th packet is lost, but this is not interpreted as the beginning of a new talk spurt. Also, when calculating $\hat{\delta}_{k+6}$ we are missing $\hat{\delta}_{k+5}$, but we just use $\hat{\delta}_{k+4}$ in its stead.

The new talk spurt starts at k+10, because there is no gap in sequence numbers, but the difference between the timestamps of subsequent packets is $t_{k+10} - t_{k+9} = 40$ ms > 20 ms, which indicates the beginning of a new talk spurt. The length of time from when the first packet in this new talk spurt is generated until it is played out is:

$$q_{k+10} = \hat{\delta}_{k+10} + K \cdot \hat{\upsilon}_{k+10} = 92.051 + 4 \times 15.9777 = 155.9618 \text{ ms} \approx 156 \text{ ms}$$

and this is reflected on the playout times of packets k+10 and k+11.

Packet seq. #	Timestamp t_i [ms]	Arrival time r_i [ms]	Playout time p_i [ms]	Average delay $\hat{\delta}_i$ [ms]	Average deviation $\hat{\upsilon}_i$
k	400	480	550	90	15
k+1	420	510	570	90	14.85
k+2	440	570	590	90.4	15.0975
k+3	460	600	610	90.896	15.4376
k+4	480	605	630	91.237	15.6209

k+7	540	645	690	91.375	15.6009
k+6	520	650	670	91.761	15.8273
k+8	560	680	710	92.043	15.9486
k+9	580	690	730	92.223	15.9669
k+10	620	695	776	92.051	15.9777
k+11	640	705	796	91.78	16.0857

Problem 3.5 — Solution

Given the playout delay, we can determine the constant K using Eq. (3.1) to find out which value of K gives the playout delay of 300 ms. Then, we use the chart shown in Figure A-7 in Appendix A to guess approximately what is the percentage of packets that will arrive to late to be played out.

Problem 3.6 — Solution

Problem 3.7 — Solution

The solutions for (a) and (b) are shown in the figure below.

(c)

If the RPF uses pruning and routers *E* and *F* do not have attached hosts that are members of the multicast group, then there will be 6 packets less forwarded in the entire network per every packet sent by *A*, compared to the case (b).

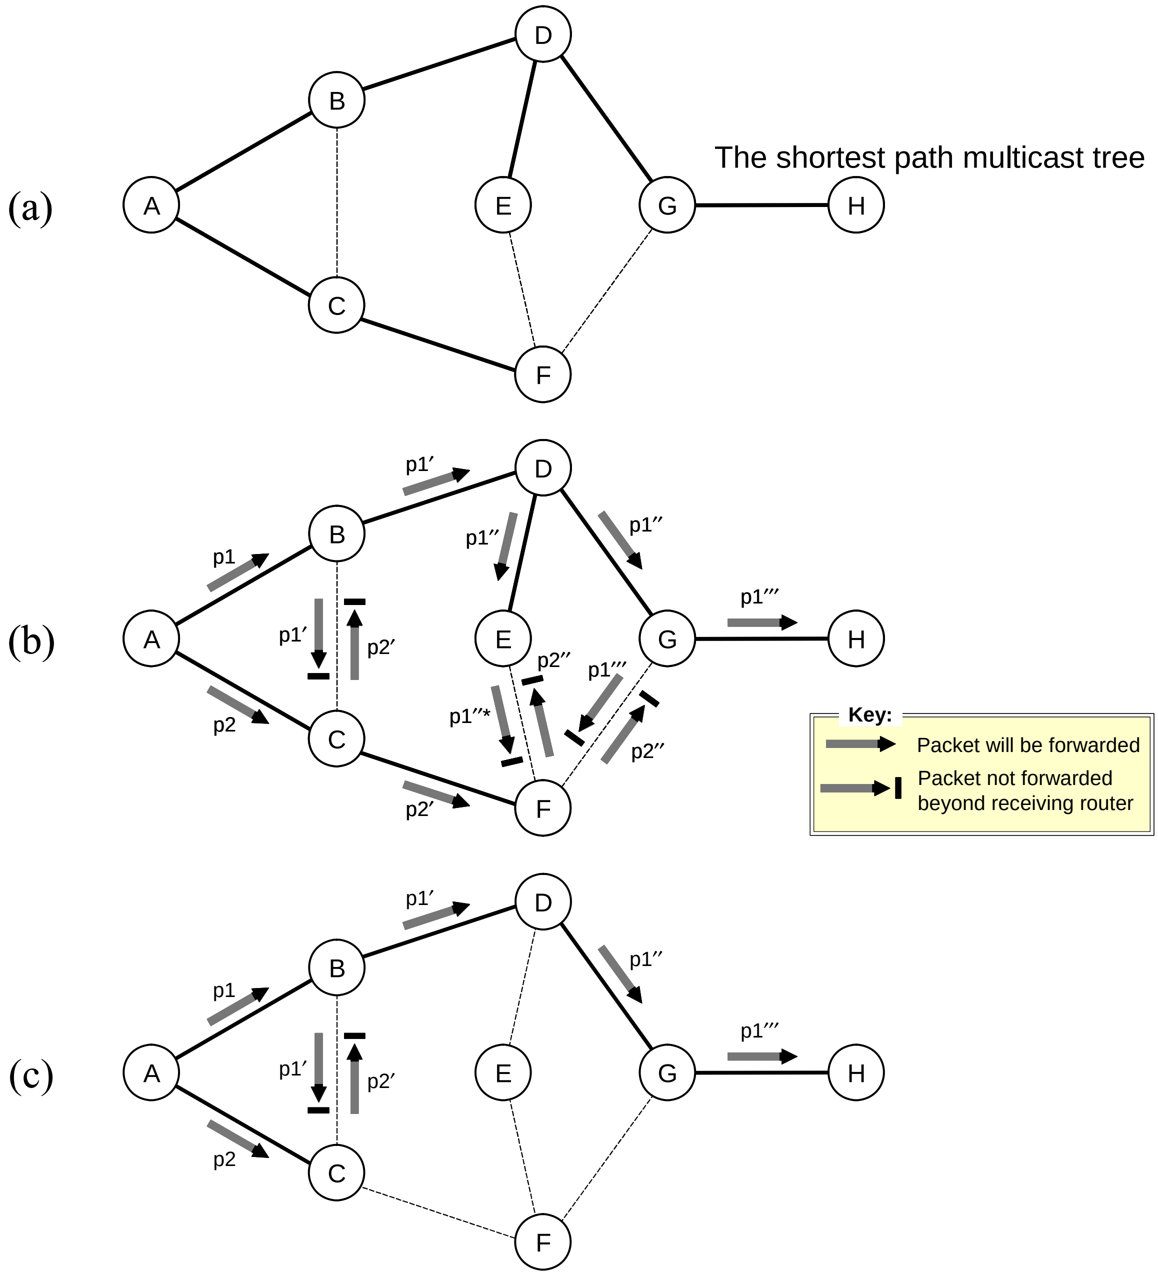

Problem 4.1 — Solution

Problem 4.2 — Solution

Problem 4.3 — Solution

The architecture of a first-generation router is shown in Figure 4-5(a). As shown in Figure 4-6, in this architecture every packet must cross the system bus two times on its way from the input port to the output port. We assume that processing times in line cards and the CPU are negligible.

(a)

If packets arrive simultaneously on all four ports, it will take $T_4 = 4 \times 2 \times \frac{L}{R}$ time units to move the packets from their input ports to their output ports. During this time, two packets can arrive on each input port. If packets continue arriving at the link data rate R, which is the *peak rate*, then during each cycle two packets will arrive on an input port and only one will be moved to the output port. The remaining packets will accumulate and the delay can grow arbitrarily high, depending on how long the period of peak-rate arrivals lasts.

The figure below illustrates an example of peak rate behavior. We assume that packets will arrive on all input ports simultaneously. However, to move the packet to the CPU, the input ports will access the system bus randomly, without priority access. Similarly, CPU will randomly access the system bus to move the packets to their output ports. As already stated, it takes T_4 time units move four incoming packets to their output ports, but at the same time up to 2 new packets can arrive at each input port. As illustrated in the figure, there will be a queue buildup on each input port. Notice also that there may be queue buildup on output ports, for example, if packets from different input ports are heading to the same output port. Because of random access to the system bus, even packets from the same port may need to queue, as illustrated for output port 3 in the figure below at time t_3.

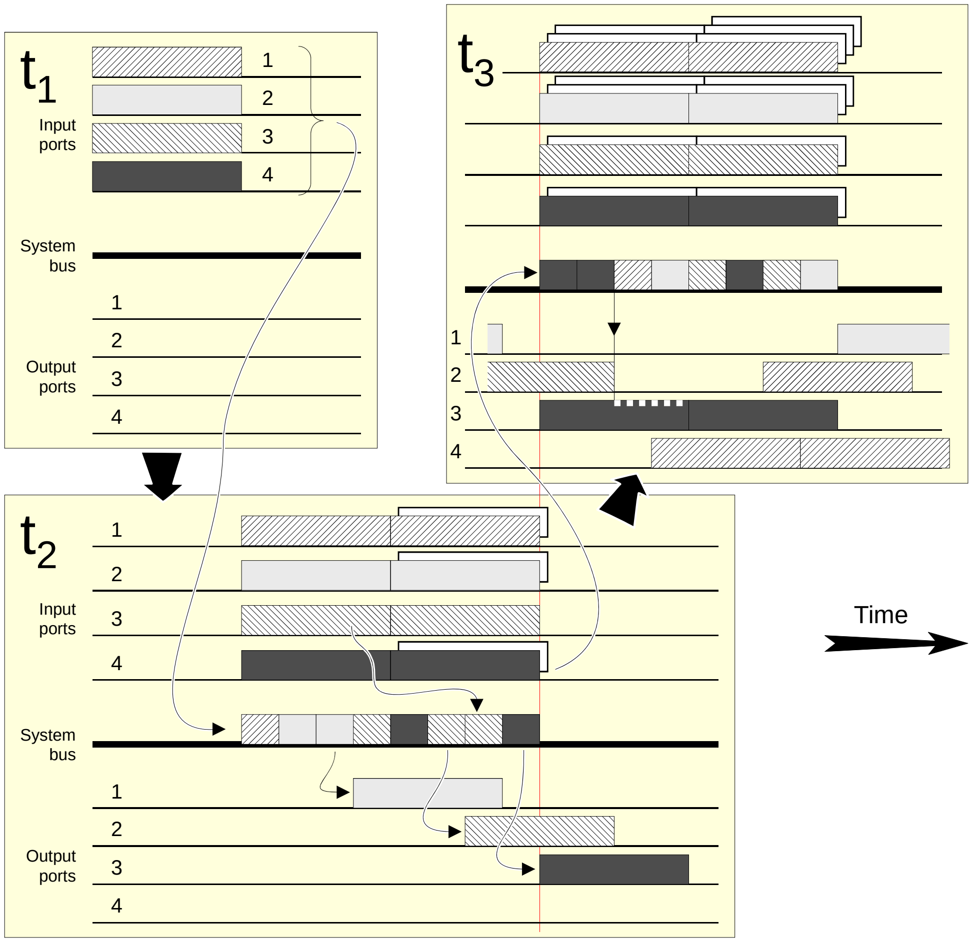

75

(b)

Technically speaking, there is *neither head-of-line blocking nor output blocking* in this system. Head-of-line blocking happens when one packet may be heading to an idle port, but in front of it may be another packet headed for an output port that is currently busy, and the former packet must wait until the one in front of it departs. This assumes that packets are not moved to their output port until their output port becomes idle. However, in the first generation routers, both input and output ports must be able to queue packets. (An input port must be able to queue a newly arriving packet while some previously arrived packets are waiting for access to the system bus, to be moved to the CPU for forwarding table lookup. The input port does not know where the incoming packet is heading, so all incoming packets are queued in FCFS manner to wait for access to the system bus and transfer to CPU.

CPU will examine the packet's header and decide to which output port it goes. Once CPU decides the output port, the packet is lined up in a FCFS queue, where the same queue is for *all* output ports. However, the head-of-line packet is never waiting for its output port to become idle; rather, it is waiting for the system bus access to get moved to its output port (regardless of whether this output port is currently idle or busy).

Output blocking occurs if two or more packets (all from different input ports) are headed to the same output port and the switching fabric is unable to move the simultaneously. The first generation routers are based on system bus as their switching fabric, so packets are always moved sequentially, one-by-one, and *never* simultaneously. Therefore, any delays that packet experience at their input ports neither qualify neither as head-of-line blocking nor as output blocking.

This is not to say that there are no queuing delays in this system. Quite opposite, there will be a major buildup of packets both at input and output queues during peak-arrival-rate periods, as explained above in part (a).

Problem 4.4 — Solution

Problem 4.5 — Solution

Check Problem 1.32 — Solution to see that the next hop router for the arrived packets is as follows:

Packet arrived from	Packet destination IP address	Next hop	Output port
B	63.67.145.18	G	6 (tag: 110)
C	223.123.59.47	D	3 (tag: 011)
G	223.125.49.47	C	2 (tag: 010)

The packets will traverse the Banyan switch as shown in the figure below. The top part of the figure shows how the router's network ports are wired to the switching fabric. Every network port is bidirectional and connected by a full-duplex link to another router.

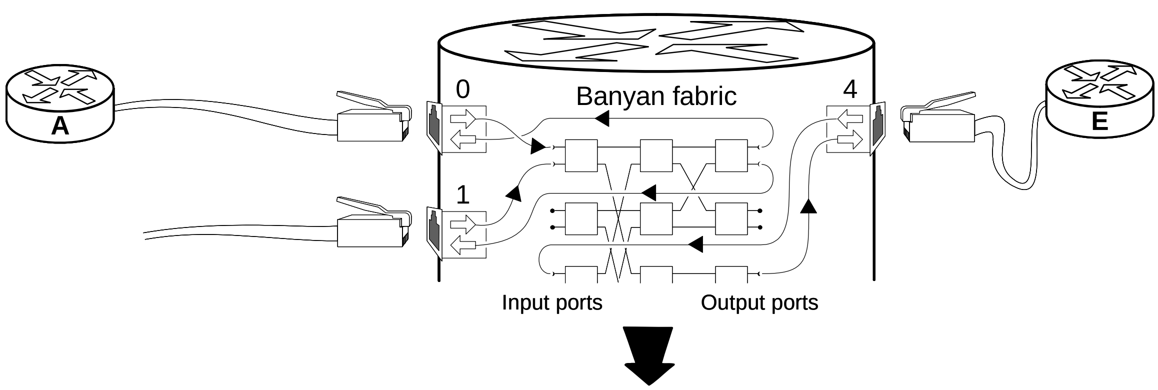

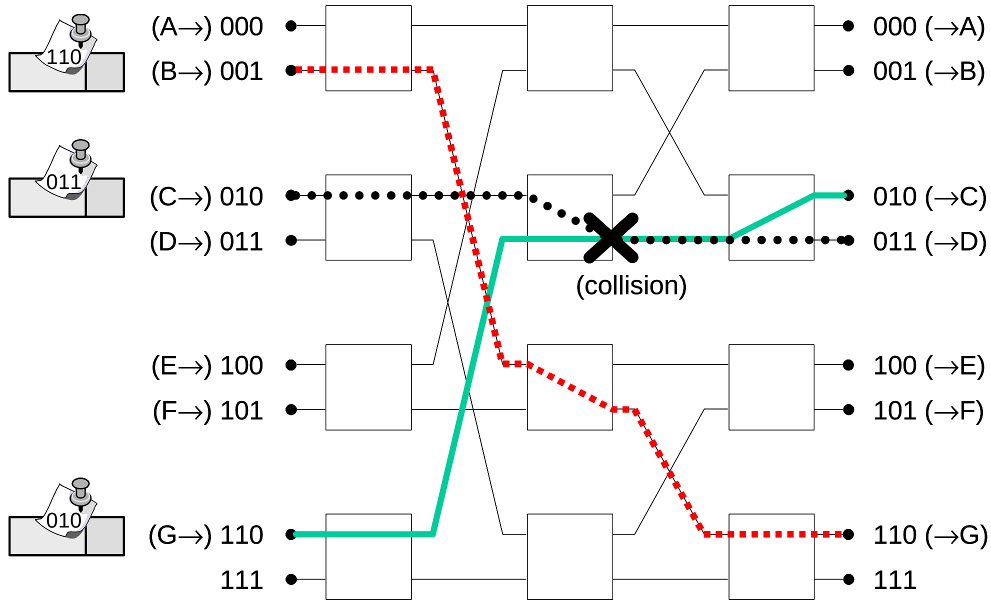

Notice that although all packets go to different output ports, there will be a collision in the second stage of the Banyan fabric because packets from ports 2 and 6 are trying to go to the same output of the second-stage 2×2 switching element.

Problem 4.6 — Solution

Problem 4.7 — Solution

Two alternative solutions for the Batcher-banyan fabric are shown:

(a)

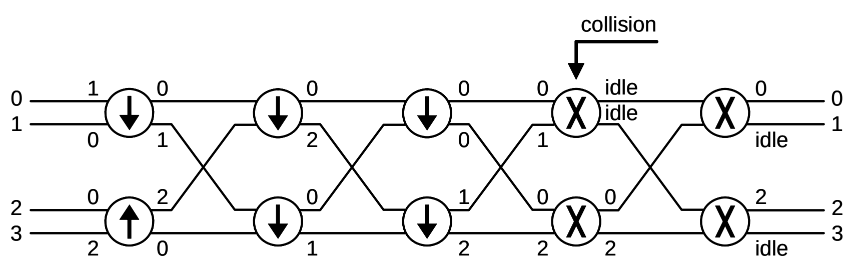

(b)

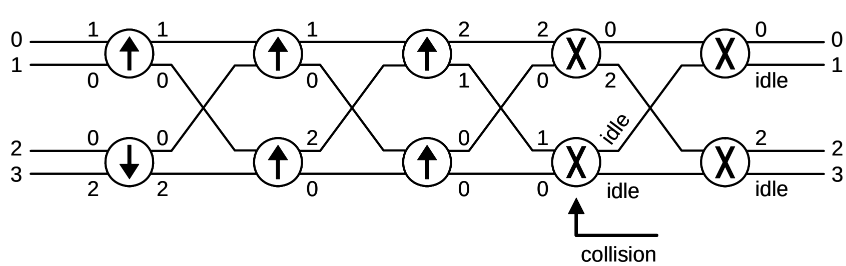

Result: the packets at inputs I0 and (I1 or I2) are lost.

The reader may notice that the above 4×4 Batcher network looks different from that in Figure 4-10(b). This just means that the same sorting problem can be solved in different ways, with different Batcher networks.

Problem 4.8 — Solution

Problem 4.9 — Solution

The figure below shows the components of the router datapath delay (compare to Figure 4-13). We ignore the forwarding decision delay. Transmission delays as well as reception delays are the same on all communication lines. Crossbar traversal delay equals $\frac{1}{2} t_x$.

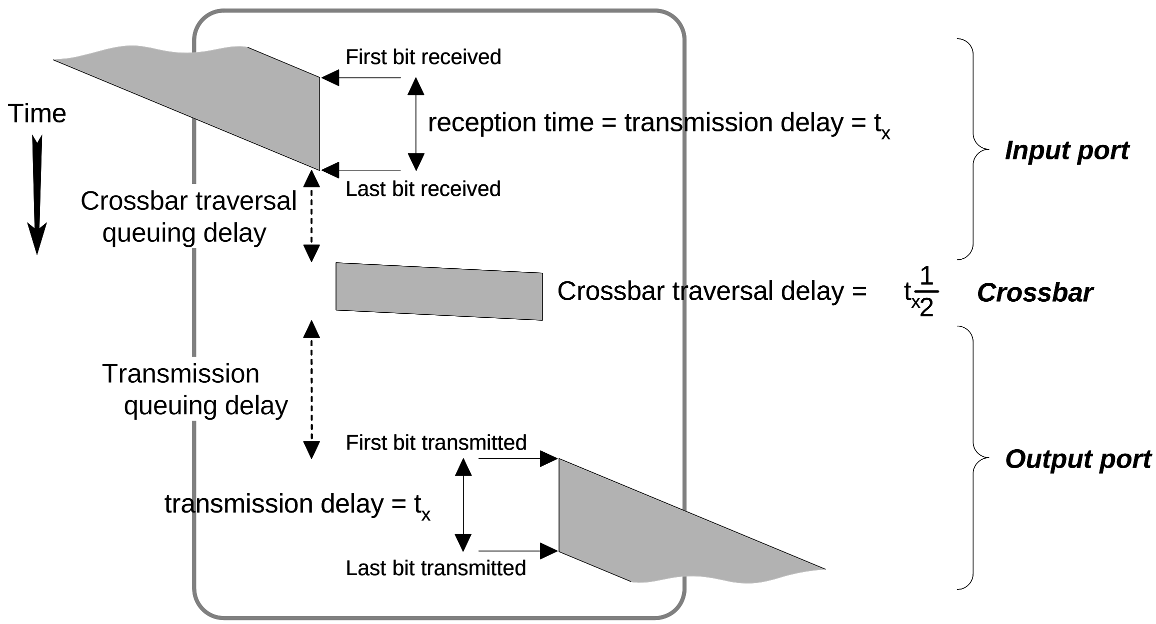

The timing diagram of packet handling in the router is shown in the figure below.

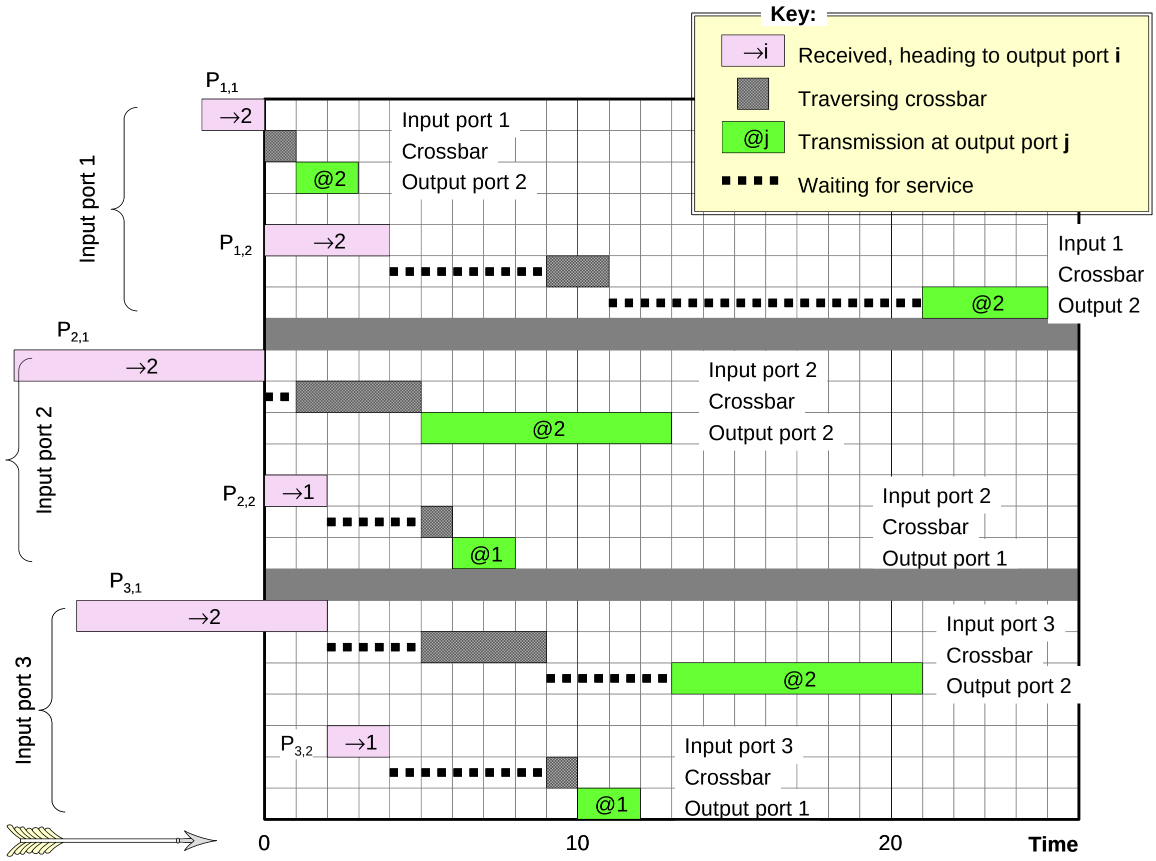

Only the first packet on input port 1 ($P_{1,1}$) will experience no waiting at all. All other packets will experience some form of blocking. The second packet on input port 1 ($P_{1,2}$) must wait before traversing the crossbar because the first packet on input port 2 ($P_{2,1}$) is currently traversing the crossbar and then it must wait for $P_{3,1}$ (although $P_{3,1}$ is on a higher-index port, it arrived before $P_{1,2}$)—so $P_{1,2}$ is experiencing *output blocking*. $P_{2,1}$ is also experiencing output blocking because it has to wait for $P_{1,1}$ to traverse the crossbar. Finally, packet $P_{3,1}$ is also experiencing output blocking because it has to wait for $P_{2,1}$ to traverse the crossbar.

Packets $P_{2,2}$ and $P_{3,2}$ are experiencing *head-of-line blocking*, because they could traverse the crossbar if it were not for packets $P_{2,1}$ and $P_{3,1}$, respectively, which are if front of them in their respective queue and are blocking their access to the crossbar.

Output blocking and head-of-line blocking both prevent crossbar traversal and therefore cause queuing delay before crossbar traversal.

Notice also that packets $P_{1,2}$ and $P_{3,1}$ must wait at the output port for their turn for transmission. This is indicated as transmission queuing delay in the first figure above.

Problem 4.10 — Solution

Problem 4.11 — Solution

Problem 4.12 — Solution

Problem 4.13 — Solution

The problem statement gives arrival times A_i and service times X_i. From these we need to determine total delays T_i in the system for each customer. Notice that there is a single server, so if a new customer arrives while one customer is served, the new customer should join the waiting line. Figure 4-17 illustrates that the total delay for a customer consists of waiting plus service time, i.e., $T_i = W_i + X_i$.

(a)

The following figure illustrates the arrivals $A(t)$, delays T_i and departures $B(t)$. The first customer arrives at an idle server, so immediately goes into service, i.e., $W_1 = 0$ and $T_1 = X_1$. The second customer arrives at time $t = 2$ and finds the first customer in service, so has to wait one time unit, $W_2 = 1$ and $T_2 = W_2 + X_2$. The length of the observed interval is 36 time units, until the last customer that arrived during the observed interval (customer #10) departs.

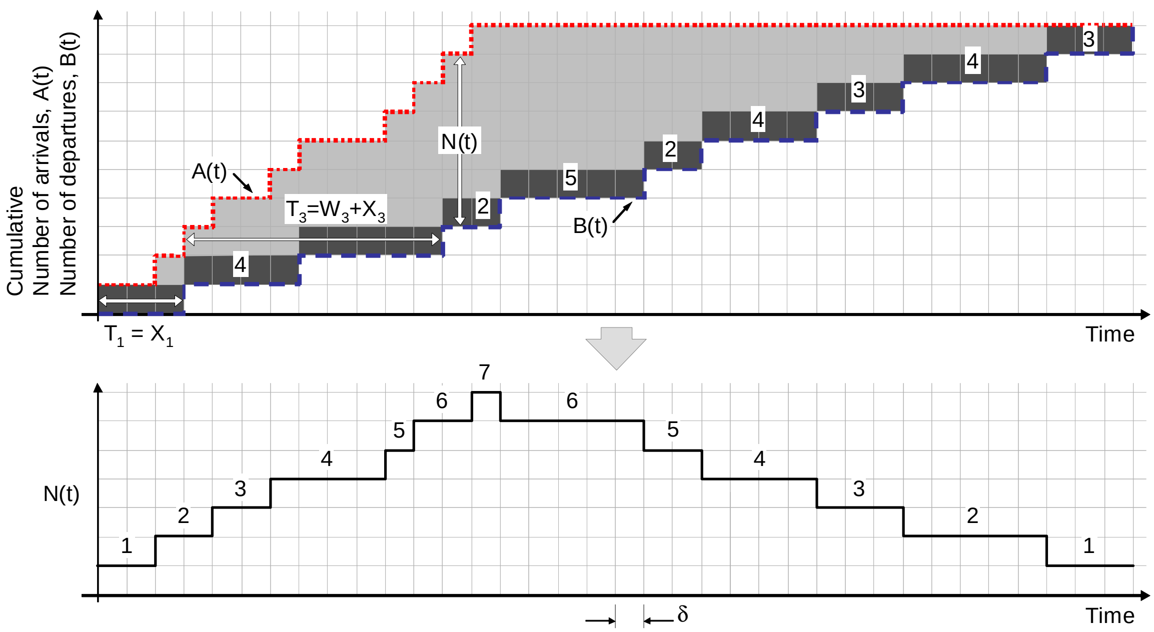

(b)

The lower chart indicates the current number of customers in the system, $N(t)$. We can calculate the average number N over the observed interval as follows

$$N = \frac{(5\times 1)+(7\times 2)+(5\times 3)+(8\times 4)+(3\times 5)+(7\times 6)+(1\times 7)}{36} = \frac{130}{36} = 3.6\dot{1}\ \text{customers}$$

The average delay T per customer over the observed interval is

$$T = \frac{3+(1+4)+(4+5)+(8+2)+(8+5)+(12+2)+(11+4)+(14+3)+(16+4)+(20+3)}{10} = \frac{129}{10} = 12.9$$

Assuming that the arrival rate is $\lambda = 1$ customer/unit-of-time, the Little's Law should yield $N =$

$\lambda \cdot T$. However, over the observed interval we have 3.61 ≠ 1×12.9. Therefore, the system does *not* satisfy the Little's Law over the observed interval.

Problem 4.14 — Solution

Problem 4.15 — Solution

Problem 4.16 — Solution

(a)

This is an $M/M/1$ queue with the arrival rate λ = 950,000 packets/sec and service rate μ = 1,000,000 packets/sec. The expected queue waiting time is:

$$W = \frac{\lambda}{\mu \cdot (\mu - \lambda)} = \frac{950000}{1000000 \times (1000000 - 950000)} = 19 \times 10^{-6} \text{ sec}$$

(b)

The time that an average packet would spend in the router if no other packets arrive during this time equals its service time, which is $\frac{1}{\mu} = \frac{1}{1000000} = 1 \times 10^{-6}$ sec

(c)

By Little's Law, the expected number of packets in the router is

$$N = \lambda \cdot T = \lambda \cdot \left(W + \frac{1}{\mu} \right) = 950000 \times 20 \times 10^{-6} = 19 \text{ packets}$$

Problem 4.17 — Solution

Problem 4.18 — Solution

Given

Data rate is 9600 bps ⇒ the average service time is $\frac{1}{\mu} = \frac{\text{average packet length}}{\text{link data rate}} = \frac{1000 \times 8}{9600} = 0.83$

∴ μ = 1.2

Link is 70% utilized ⇒ the utilization rate is ρ = 0.7

For exponential message lengths: $M/M/1$ queue with μ = 1.2, ρ = 0.7, the average waiting time is $W = \frac{\rho}{\mu \cdot (1 - \rho)} = \frac{0.7}{1.2 \times 0.3} = 1.94 \text{ sec}$.

For constant-length messages we have $M/D/1$ queue and the average waiting time is derived in the solution of Problem 4.22(b) below as: $W = \dfrac{\rho}{2 \cdot \mu \cdot (1-\rho)} = \dfrac{0.7}{2 \times 1.2 \times 0.3} = 0.97 \text{ sec}$.

It is interesting to notice that constant-length messages have 50 % shorter expected queue waiting time than the exponentially distributed length messages.

Problem 4.19 — Solution

The single repairperson is the server in this system and the customers are the machines. Define the system state to be the number of operational machines. This gives a Markov chain, which is the same as in an $M/M/1/m$ queue with arrival rate μ and service rate λ. The required probability is simply p_m for such a queue. Because the sum of state probabilities is $\sum_{i=0}^{m} p_i = 1$, the fraction of time the system spends in state m equals p_m. From Eq. (4.8), we have the steady-state proportion of time where there is no operational machine as $p_m = \dfrac{\rho^m \cdot (1-\rho)}{1-\rho^{m+1}}$.

Problem 4.20 — Solution

This can be modeled as an $M/M/1/m$ system, because the there are a total of K users, and there can be up to K tasks in the system if their file requests coincide. The average service time is $\dfrac{1}{\mu} = \dfrac{\text{average packet length}}{\text{throughput rate}} = \dfrac{A \times R}{R} = A$ and the service rate is μ = $1/A$. The user places the request, but may need to wait if there are already pending requests of other users. Let W denote the waiting time once the request is placed but before the actual transmission starts, which is unknown. Every user comes back, on average, after $A+B+W$ seconds. Hence, the arrival rate is λ = $\dfrac{K}{A+B+W}$.

From Little's Law, given the average number N of customers in the system, the average waiting delay per customer is $W = T - A = \dfrac{N}{\lambda} - A$. The time T is from the moment the user places the request until the file transfer is completed, which includes waiting after the users who placed their request earlier but are not yet served, plus the time it takes to transfer the file (service time), which on average equals A seconds. (Only one customer at a time can be served in this system.)

Then, $\lambda = \dfrac{N}{W+A} = \dfrac{K}{A+B+W}$ and from here: $W = \dfrac{N \cdot (A+B) - K \cdot A}{K-N}$

For an $M/M/1/m$ system, the average number N of users requesting the files is:

$$N = \frac{\rho}{1-\rho} - \frac{(K+1) \cdot \rho^{K+1}}{1-\rho^{K+1}}$$

where $\rho = \lambda / \mu$ is the utilization rate. Finally, the average time it takes a user to get a file since completion of his previous file transfer is $\boldsymbol{A} + \boldsymbol{B} + \boldsymbol{W}$.

Problem 4.22 — Solution

This is an *M*/*D*/1 queue with deterministic service times. Recall that *M*/*D*/1 is a sub-case of *M*/*G*/1. Given: Service rate, μ = 1/4 = 0.25 items/sec; arrival rate, λ = 0.2 items/sec.

(a)

Mean service time $\overline{X}$ = 4 sec.

$$\overline{X}^2 = \frac{1}{\mu^2} \text{ and } N_Q = \frac{\lambda^2 \cdot \overline{X}^2}{2 \cdot (1-\rho)} = \frac{\lambda^2}{2 \cdot \mu^2 \cdot (1-\lambda/\mu)} = 16$$

The second moment of service time for the deterministic case is obtained as

$$0 = E\{x-\mu\}^2 = E\{x^2\} - E^2\{x\} \text{ and from here, we have } E\{x^2\} = E^2\{x\} = \overline{X}^2 = \frac{1}{\mu^2}$$

(b)

The total time spent by a customer in the system, T, is $T = W + \overline{X}$, where W is the waiting time in the queue $W = \dfrac{\rho}{2 \cdot \mu \cdot (1-\rho)} = 8 \text{ sec}$ so the total time T = 12 sec.

Problem 4.23 — Solution

Problem 4.24 — Solution

Problem 5.1 — Solution

Problem 5.2 — Solution

Problem 5.3 — Solution

Problem 5.4 — Solution

Recall that packet-by-packet FQ is *non-preemptive*, so the packet that is already in transmission will be let to finish regardless of its finish number. Therefore, the packet of class 3 currently in transmission can be ignored from further consideration. It is interesting to notice that the first packet from flow 1 has a smaller finish number, so we can infer that it must have arrived *after* the packet in flow 3 was already put in service.

The start round number for servicing the currently arrived packet equals the current round number, because its own queue is empty. Hence, $F_{2,1} = R(t) + L_{2,1} = 85000 + 1024\times8 = 93192$.

Therefore, the order of transmissions under FQ is: $pkt_{2,1} < pkt_{1,1} < pkt_{1,2}$; that is, the newly arrived packet goes first (after the one currently in service is finished).

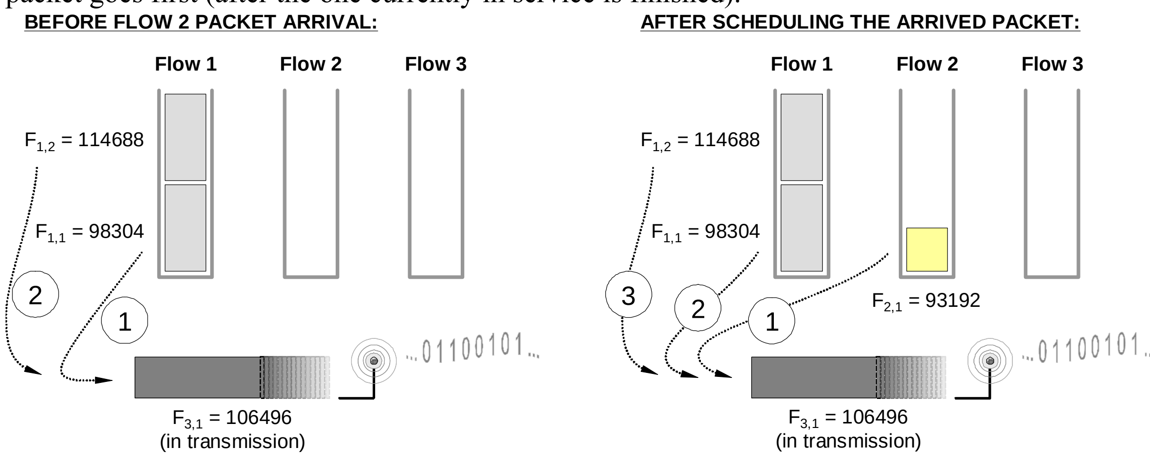

Problem 5.5 — Solution

Problem 5.6 — Solution

$F_{1,1} = L_{1,1}/w_1 = 200\times8/3 = 533.3$

$F_{2,1} = L_{2,1}/w_2 = 50\times8/1 = 400$

$F_{3,1} = L_{3,1}/w_3 = 1000\times8/1.5 = 5333.3$

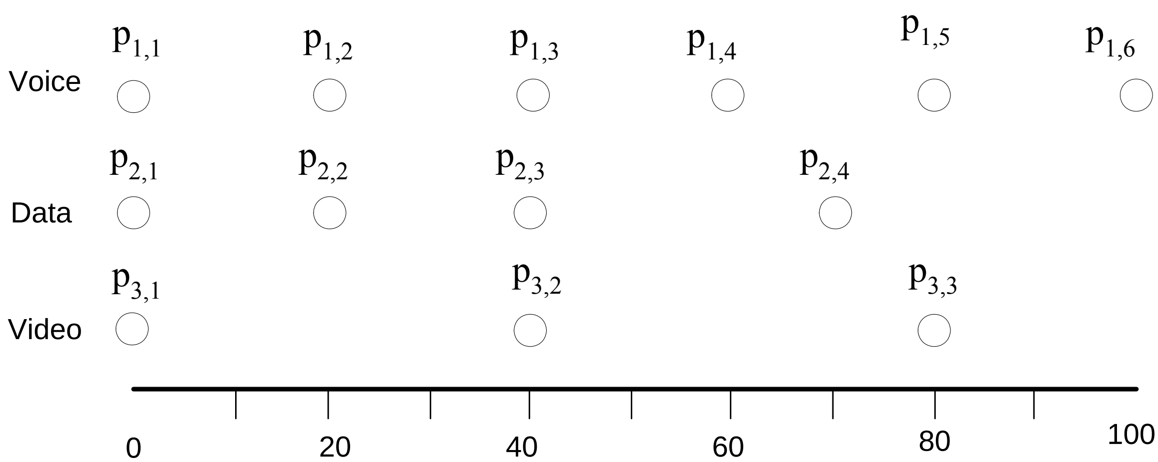

1. At time t = 0, $P_{2,1}$ is served first; transmission time $t_x = 400/1M = 0.4ms$;
2. At time 0.4ms, $P_{1,1}$ is served; $t_x = 1600/1M = 1.6ms$;
3. At time 2ms, $P_{3,1}$ is served; $t_x = 8000/1M = 8ms$;
4. At time 10ms, no packets in any queues
5. At time t = 20ms

   $R(t) = t\times C/N = 20000/2=10000$

   $F_{1,2} = \max(F_{1,1}, R(t)) + L_{1,2}/w_1 =10000 + 533.3 = 10533.3$

   $F_{2,2} = 10000 + 500\times 8/1 = 14000$
6. At time 20ms, $P_{1,2}$ is served; $t_x = 1.6ms$
7. At time 21.6ms, $P_{2,2}$ is served; $t_x = 4000/1M = 4ms$
8. At time t = 40ms

   $R(t) = 40000/3 = 13333.3$

   $F_{1,3} = \max(F_{1,2}, R(t)) + L_{1,3}/w_1 = 13333.3 + 533.3 = 13866.6$

   $F_{2,3} = \max(F_{2,2}, R(t)) + L_{2,3}/w_2 = 14000 + 8000/1= 22000$

   $F_{3,2} = \max(F_{3,1}, R(t)) + L_{3,2}/w_2 = 5333.3 + 5333.3 = 10666.6$
9. At time 40ms, $P_{3,2}$ is served; $t_x = 8000/1M = 8ms$
10. At time 48ms, $P_{1,3}$ is served; $t_x = 1.6ms$
11. At time 49.6, $P_{2,3}$ is served; $t_x = 8ms$
12. At time t = 60, $P_{1,4}$ is served because there are no other queued packets, $t_x = 1.6ms$

    $R(t) = 60000$

    $F_{1,4} = \max(F_{1,3}, R(t)) + L_{1,4}/w_1 = 60000 + 533.3 = 60533.3$
13. At time 70, $P_{2,4}$ is served; $t_x = 0.4ms$
14. At time t = 80

R(t) = 80000/2 = 40000

F1,5 = max($F_{1,4}$, R(t)) + $L_{1,5}$/w1 = 60533.3 + 533.3 = 61066.6

F3,3 = max($F_{3,2}$, R(t)) + $L_{3,3}$/w3 = 40000 + 5333.3 = 45333.3

15. At time 80, $P_{3,3}$ is served

16. $P_{1,5}$

17 At last, $P_{1,6}$

Problem 5.7 — Solution

(a) Packet-by-packet FQ

The following figure helps to determine the round numbers, based on bit-by-bit GPS. The packets are grouped in two groups, as follows. Regardless of the scheduling discipline, all of the packets that arrived by 300 s will be transmitted by 640 s. It is easy to check this by using a simple FIFO scheduling. Therefore, the round number $R(t)$ can be considered independently for the packets that arrive up until 300 s vs. those that arrive thereafter. This is shown as Part A and B in the above figure. (Resetting the round number is optional, only for the sake of simplicity.)

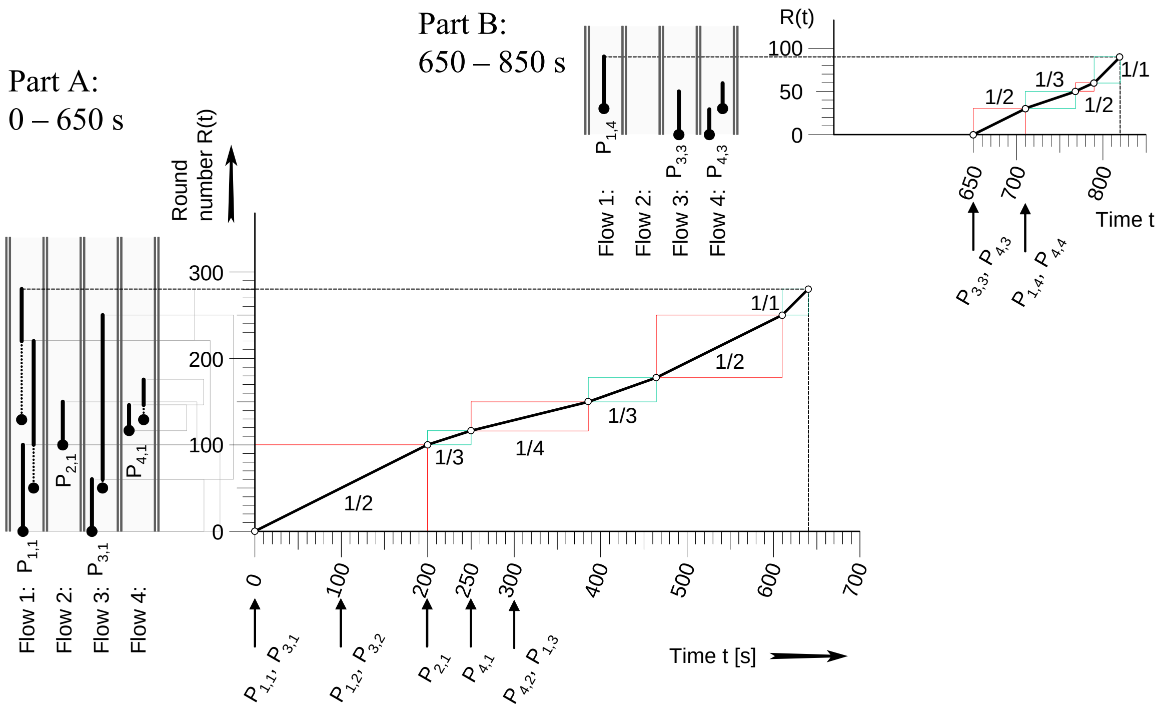

The packet arrivals on different flows are illustrated on the left hand side of the figure, in the round number units. Thus, e.g., packet $P_{2,1}$ arrives at time $t_{2,1}$ = 200 s or round number $R(t_{2,1})$ = 100. The following table summarizes all the relevant computations for packet-by-packet FQ.

Arrival times & state	Parameters	Values under packet-by-packet FQ
$t = 0$: $\{P_{1,1}, P_{3,1}\}$ arrive server idle, q's empty	Finish numbers	$R(0) = 0$; $F_{1,1} = L_{1,1} = 100$; $F_{3,1} = L_{3,1} = 60$
	Transmit periods	Start/end($P_{3,1}$): 0→60 sec; Start/end($P_{1,1}$): 60→160 s
$t = 100$: $\{P_{1,2}, P_{3,2}\}$ $P_{1,1}$ in transmission	Finish numbers	$R(t) = t{\cdot}C/N = 100{\times}1/2 = 50$ $F_{1,2} = \max\{F_{1,1}, R(t)\} + L_{1,2} = 100 + 120 = 220$;

All queues empty		$F_{3,2} = \max\{0, R(t)\} + L_{3,2} = 50 + 190 = 240$
	Transmit periods	Start/end($P_{1,2}$): 160→280 s; Queued packets: $P_{3,2}$
$t = 200$: $\{P_{2,1}\}$ arrives $P_{1,2}$ in transmission $P_{3,2}$ in queue	Finish numbers	$R(t) = t{\cdot}C/N = 200{\times}1/2 = 100$ $F_{2,1} = \max\{0, R(t)\} + L_{2,1} = 100 + 50 = 150$ $F_{3,2} = 240$ (unchanged);
	Transmit periods	$P_{1,2}$ ongoing; Queued packets: $P_{2,1} < P_{3,2}$
$t = 250$: $\{P_{4,1}\}$ arrives $P_{1,2}$ in transmission $\{P_{2,1}, P_{3,2}\}$ in queues	Finish numbers	$R(t) = (t{-}t'){\cdot}C/N + R(t') = 50{\times}1/3 + 100 = 116\ 2/3$ $F_{2,1} = 150$ (unchanged); $F_{3,2} = 240$ (unchanged); $F_{4,1} = \max\{0, R(t)\} + L_{4,1} = 116.6\dot{7} + 30 = 146.6\dot{7}$
	Transmit periods	Start/end($P_{4,1}$): 280→310 s; Queued pkts: $P_{2,1} < P_{3,2}$
$t = 300$: $\{P_{4,2}, P_{1,3}\}$ $P_{4,1}$ in transmission $\{P_{2,1}, P_{3,2}\}$ in queues	Finish numbers	$R(t) = (t{-}t'){\cdot}C/N + R(t') = 50{\times}1/4 + 116.6\dot{7} = 129\ 1/6$ $F_{1,3} = \max\{0, R(t)\}{+}L_{1,3} = 129.1\dot{6} + 60 = 189.1\dot{6}$; $F_{2,1} = 150$ (unchanged); $F_{3,2} = 240$ (unchanged); $F_{4,2} = \max\{0, R(t)\} + L_{4,2} = 129.1\dot{6} + 30 = 159.1\dot{6}$
	Transmit periods	$P_{4,1}$ ongoing; Queued packets: $P_{2,1} < P_{4,2} < P_{1,3} < P_{3,2}$ Start/end($P_{2,1}$): 310→360 s; s/e($P_{4,2}$): 360→390 s; Start/end($P_{1,3}$): 390→450 s; s/e($P_{3,2}$): 450→640 s.
At $t = 640$ s, round number reset, $R(t) = 0$, because the system becomes idle.		
$t = 650$: $\{P_{3,3}, P_{4,3}\}$ server idle, q's empty	Finish numbers	$R(0) = 0$; $F_{3,3} = L_{3,3} = 50$; $F_{4,3} = L_{4,3} = 30$
	Transmit periods	Start/end($P_{4,3}$): 650→680 sec; s/e($P_{3,3}$): 680→730 s.
$t = 710$: $\{P_{1,4}, P_{4,4}\}$ $P_{3,3}$ in transmission All queues empty	Finish numbers	$R(t) = (t{-}t'){\cdot}C/N + R(t') = 110{\times}1/2 + 0 = 55$ $F_{1,4} = \max\{0, R(t)\} + L_{1,4} = 55 + 60 = 115$; $F_{4,4} = \max\{30, R(t)\} + L_{4,4} = 55 + 30 = 85$
	Transmit periods	$P_{3,3}$ ongoing; Queued packets: $P_{4,4} < P_{1,4}$ Start/end($P_{4,4}$): 730→760 s; s/e($P_{1,4}$): 760→820 s.

(b) Packet-by-packet WFQ; Weights for flows 1-2-3-4 are 4:2:1:2

The round number computation, based on bit-by-bit GPS, remains the same as in the figure above. The only difference is in the computation of finish numbers under packet-by-packet WFQ, see Eq. (5.3), as summarized in the following table.

Packets $P_{1,4}$ and $P_{4,4}$ end up having the same finish number (70); the tie is broken by a random drawing so that $P_{1,4}$ is decided to be serviced first, ahead of $P_{4,4}$.

Arrival times & state	**Parameters**	**Values under packet-by-packet WFQ**
$t = 0$: $\{P_{1,1}, P_{3,1}\}$ arrive server idle, q's empty	Finish numbers	$R(0) = 0$; $F_{1,1} = L_{1,1}/w_1 = 100/4 = 25$; $F_{3,1} = 60$
	Transmit periods	Start/end($P_{1,1}$): 0→100 s; Start/end($P_{3,1}$): 100→160 s
$t = 100$: $\{P_{1,2}, P_{3,2}\}$ $P_{3,1}$ in transmission All queues empty	Finish numbers	$R(t) = t{\cdot}C/N = 100{\times}1/2 = 50$ $F_{1,2} = \max\{F_{1,1}, R(t)\} + L_{1,2}/w_1 = 100 + 120/4 = 130$; $F_{3,2} = \max\{0, R(t)\} + L_{3,2}/w_3 = 50 + 190/1 = 240$
	Transmit periods	Start/end($P_{1,2}$): 160→280 s; Queued packets: $P_{3,2}$
$t = 200$: $\{P_{2,1}\}$ arrives $P_{1,2}$ in transmission $P_{3,2}$ in queue	Finish numbers	$R(t) = t{\cdot}C/N = 200{\times}1/2 = 100$ $F_{2,1} = \max\{0, R(t)\} + L_{2,1}/w_2 = 100 + 50/2 = 125$ $F_{3,2} = 240$ (unchanged);

	Transmit periods	$P_{1,2}$ ongoing; Queued packets: $P_{2,1} < P_{3,2}$
$t = 250$: $\{P_{4,1}\}$ arrives $P_{1,2}$ in transmission $\{P_{2,1}, P_{3,2}\}$ in queues	Finish numbers	$R(t) = (t-t')\cdot C/N + R(t') = 50\times 1/3 + 100 = 116\ 2/3$ $F_{2,1} = 125$ (unchanged); $F_{3,2} = 240$ (unchanged); $F_{4,1} = \max\{0, R(t)\} + L_{4,1}/w_4 = 116.6\dot{7} + 30/2 = 131.6\dot{7}$
	Transmit periods	Start/end($P_{2,1}$): 280→330 s; Queued pkts: $P_{4,1} < P_{3,2}$
$t = 300$: $\{P_{4,2}, P_{1,3}\}$ $P_{2,1}$ in transmission $\{P_{3,2}, P_{4,1}\}$ in queues	Finish numbers	$R(t) = (t-t')\cdot C/N + R(t') = 50\times 1/4 + 116.6\dot{7} = 129\ 1/6$ $F_{1,3} = \max\{0, R(t)\} + L_{1,3}/w_1 = 129.1\dot{6} + 60/4 = 144.1\dot{6}$; $F_{3,2} = 240$ (unchanged); $F_{4,1} = 131.6\dot{7}$ (unchanged); $F_{4,2} = \max\{131.6\dot{7}, R(t)\} + L_{4,2}/w_4 = 146.6\dot{7}$
	Transmit periods	$P_{2,1}$ ongoing; Queued packets: $P_{4,1} < P_{1,3} < P_{4,2} < P_{3,2}$ Start/end($P_{4,1}$): 330→360 s; s/e($P_{1,3}$): 360→420 s; Start/end($P_{4,2}$): 420→450 s; s/e($P_{3,2}$): 450→640 s.
At $t = 640$ s, round number reset, $R(t) = 0$, because the system becomes idle.		
$t = 650$: $\{P_{3,3}, P_{4,3}\}$ server idle, q's empty	Finish numbers	$R(0) = 0$; $F_{3,3} = L_{3,3}/w_3 = 50$; $F_{4,3} = L_{4,3}/w_4 = 15$
	Transmit periods	Start/end($P_{4,3}$): 650→680 sec; s/e($P_{3,3}$): 680→730 s.
$t = 710$: $\{P_{1,4}, P_{4,4}\}$ $P_{3,3}$ in transmission All queues empty	Finish numbers	$R(t) = (t-t')\cdot C/N + R(t') = 110\times 1/2 + 0 = 55$ $F_{1,4} = \max\{0, R(t)\} + L_{1,4}/w_1 = 55 + 60/4 = 70$; $F_{4,4} = \max\{30, R(t)\} + L_{4,4}/w_4 = 55 + 30/2 = 70$
	Transmit periods	$P_{3,3}$ ongoing; Queued pkts: $P_{1,4} = P_{4,4}$ (tie ⇒ random) Start/end($P_{1,4}$): 730→790 s; s/e($P_{4,4}$): 790→820 s.

Finally, the following table summarizes the order/time of departure:

Packet #	Arrival time [sec]	Packet size [bytes]	Flow ID	Departure order/ time under FQ	Departure order/ time under WFQ
1	0	100	1	#2 / 60 s	#1 / 0 s
2	0	60	3	#1 / 0 s	#2 / 100 s
3	100	120	1	#3 / 160 s	#3 / 160 s
4	100	190	3	#8 / 450 s	#8 / 450 s
5	200	50	2	#5 / 310 s	#4 / 280 s
6	250	30	4	#4 / 280 s	#5 / 330 s
7	300	30	4	#6 / 360 s	#7 / 420 s
8	300	60	1	#7 / 390 s	#6 / 360 s
9	650	50	3	#10 / 680 s	#10 / 680 s
10	650	30	4	#9 / 650 s	#9 / 650 s
11	710	60	1	#12 / 760 s	#11 / 730 s (tie)
12	710	30	4	#11 / 730 s	#11 / 790 s (tie)

Problem 5.8 — Solution

Let us assign the high-priority queue index 1, and the other two queues have indices 2 and 3. We modify Eq. (5.2) as follows. For a high priority packet $P_{1,j}$, the finish number is:

$$F_{1,j} = \max\{F_{1,j-1}, R(t_a)\} + L_{1,j} \tag{5.2$'$}$$

If the priority packet is at the head of its own queue and no other packet is currently serviced, then $F_{1,j} = R(t_a) + L_{1,j}$. If a non-priority packet $P_{i,k}$, $i \neq 1$, is currently in service, then we use its finish number $F_{i,k}$ to compute $F_{1,j} = F_{i,k} + L_{1,j}$, because a currently serviced packet is not preempted.

For a non-priority packet $P_{i,j}$, $i \neq 1$:

$$F_{i,j} = \max\{F_{1,\text{last}}, F_{i,j-1}, R(t_a)\} + L_{i,j}; \quad i \neq 1 \tag{5.2″}$$

where $F_{1,\text{last}}$ symbolizes the finish number of the *last* packet currently in the priority queue.

Any time new packet arrives (whether to a priority or non-priority queue), we must recompute the finish numbers and sort the packets in the ascending order of their finish numbers.

The order and departure times are as follows:

$P_{2,1}$ at $t = 0$ (not preempted) | $P_{1,1}$ at $t = 6$ | $P_{1,2}$ at $t = 8$ | $P_{3,1}$ at $t = 10$ | $P_{3,2}$ at $t = 12$ | $P_{2,2}$ at $t = 13$ (the tie between $P_{2,2}$ and $P_{3,2}$ is resolved by flipping a coin).

Problem 5.9 — Solution

(a)
The solution is shown in the figure blow. Notice that LSR *C* can decide that packets for destination `17.1.1/24` and `17.3.1/24` belong to the same FEC (forwarding equivalence class). The reason for this is that they traverse the same path across the MPLS domain (from LSR *F* to LSR *C*) and neither one is assigned preferential treatment. Therefore, in this case there is no reason to build more than a single LSP tunnel for the traffic from *F* to *C*.

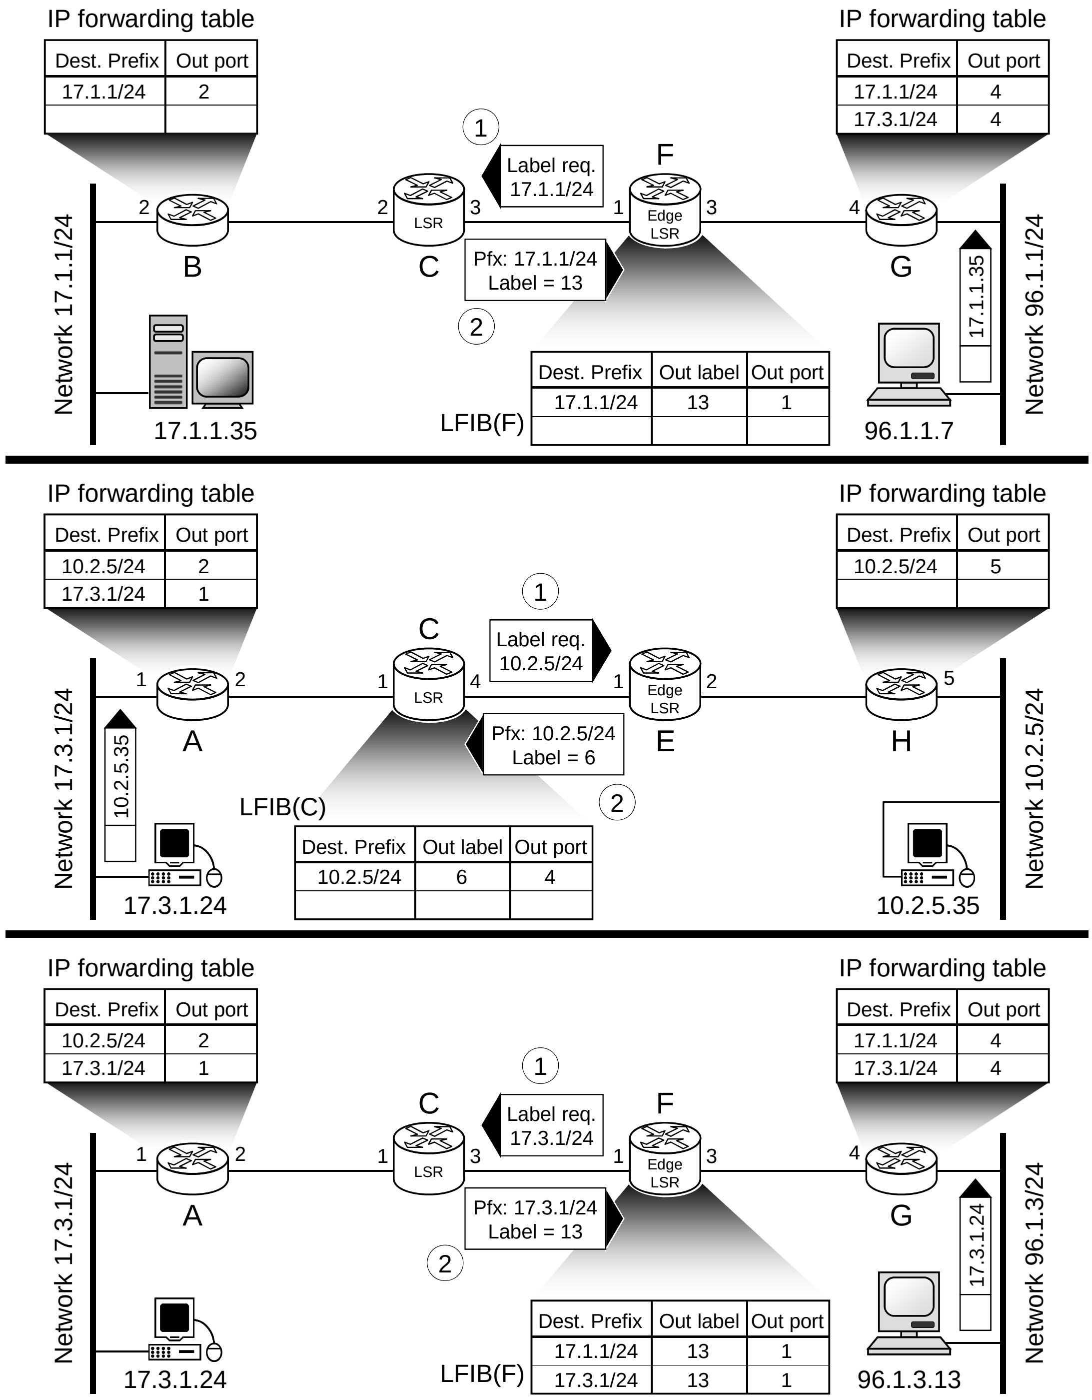

(b)

Forwarding will work as illustrated in the figure below. LSR *C* uses its conventional IP forwarding table to forward packets to destinations `17.1.1.35` (first packet) and `17.3.1.24` (3[rd] packet). Similarly, LSR *E* uses its IP forwarding table to forward packets to destination `10.2.5.35` (2[nd] packet).

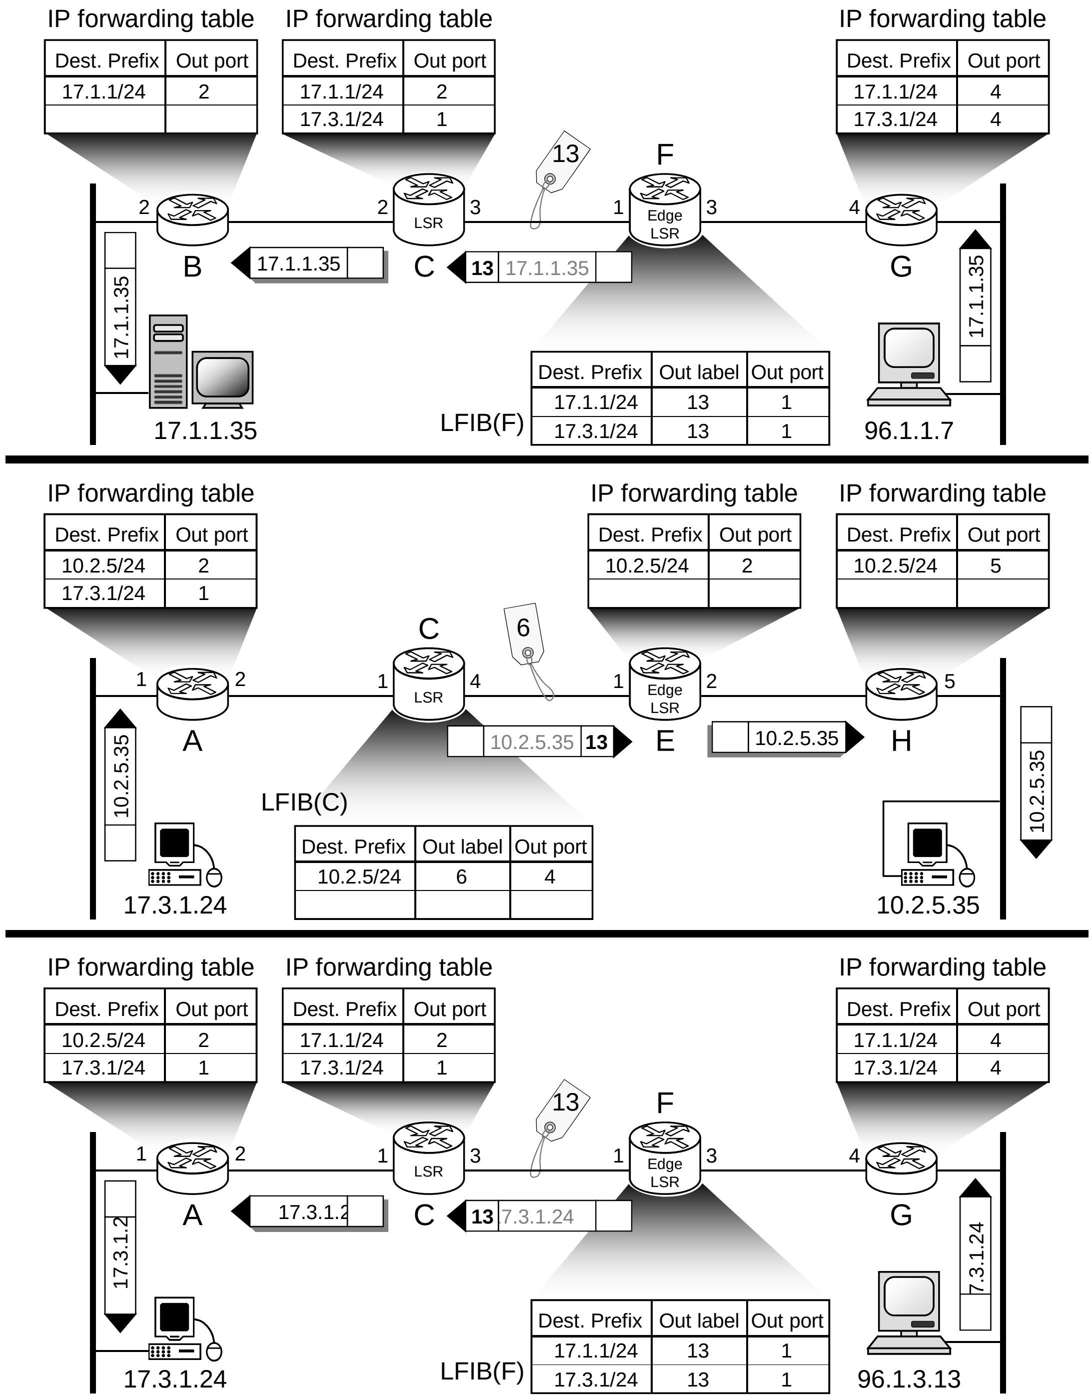

(c)

The minimum number of FECs is 2 and the minimum number of LSP tunnels that need to be set up is accordingly 2: *F*→*C* and *C*→*E*.

Appendix A: Probability Refresher

Random Events and Their Probabilities

An **experiment** (or, **observation**) is a procedure that yields one of a given set of possible outcomes. **Outcome** is a specific result of an experiment. The **sample space** *S* of the experiment is the set of possible outcomes that can occur in an experiment. An **event** *E* is a collection of outcomes, which is a subset of the sample space. For example, tossing a coin results in one of two possible outcomes: heads (H) or tails (T). Tossing a coin twice results in one of four possible outcomes: HH (two heads), HT (a head followed by a tail), TH, or TT (see Figure A-1(a)). Similarly, tossing a coin three times results in one of eight possible outcomes: HHH, HHT, HTH, HTT, THH, THT, TTH, or TTT. Consider the experiment of tossing a coin twice, where we can define a number of possible events:

Event *A*: The event "two heads" consists of the single outcome A = {HH}. (This event is equivalent to the event "no tails.")

Event *B*: The event "exactly one head" consists of two outcomes B = {HT, TH}. (This event is equivalent to the event "exactly one tail.")

Event *C*: The event "at least one head" consists of three outcomes C = {HH, HT, TH}. (This event is equivalent to the events "at most one tail" and "not two tails.")

We also define a null event ("nothing happened"), which is symbolized by ∅. Given a sample space *S*, the total number of events that can be defined is the set that contains all of the subsets of *S*, including both ∅ and *S* itself. This set is called the *power set* of set *S* and is denoted $\mathbb{P}(S)$, or $\mathbb{P}S$, or 2^S. In the example of tossing a coin twice, the power set of *S* contains $2^4 = 16$ events, including the null event ∅. The events consisting of a single outcome are: {HH}, {HT}, {TH}, {TT}. The events consisting of pairs of outcomes are:
{HH, HT}, {HH, TH}, {HH, TT}, {HT, TH}, {HT, TT}, {TH, TT}.

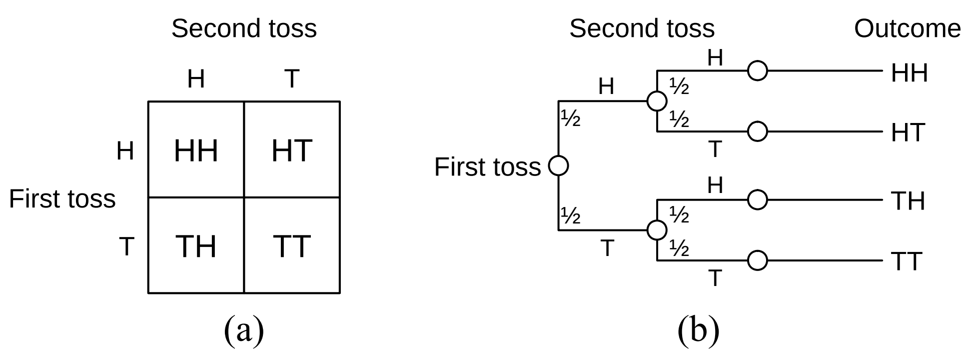

Figure A-1. (a) Possible outcomes of two coin tosses. (b) "Tree diagram" of possible outcomes of two coin tosses.

The events consisting of triples of outcomes are:
{HH, HT, TH}, {HH, HT, TT}, {HH, TH, TT}, {HT, TH, TT}.

Finally, the event consisting of all four outcomes is: {HH, HT, TH, TT}.

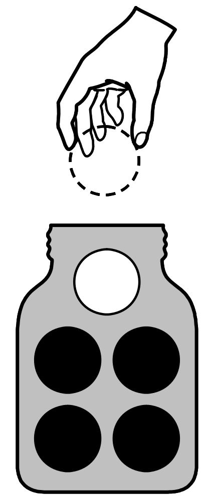

Figure A-2. Jar with black and white balls.

We say that an event is *random* when the result of an experiment is not known before the experiment is performed. For example, a coin toss is random because we do not know if it will land heads or tails. If we knew how it would land, then the event would not be random because its outcome would be predetermined. For a random event, the best we can do is *estimate* the probability of the event.

One way to define the probability of a random event is as the relative frequency of occurrence of an experiment's outcome, when repeating the experiment indefinitely. Consider a jar with five balls: four black and one white (Figure A-2). Imagine that you reach into the jar and retrieve a ball, examine its color, and put it back. If you repeat this experiment many times, then, on average, four out of five times you will retrieve a black ball and one out of five times you will retrieve the white ball. Therefore, the probability of an outcome could be defined as the frequency of the outcome.

We would like to know not only the probability of individual outcomes, but also the probability of events. Let us first consider the case when all outcomes are equally likely, and later we will consider the general case. In the case of equally likely outcomes, the probability of an event is equal to the number of outcomes in an event (cardinality of the event) divided by the number of possible outcomes (cardinality of the sample space): $p(E) = \frac{|E|}{|S|}$. For example, tossing a fair coin has two equally likely outcomes: heads and tails are equally likely on each toss and there is no relationship between the outcomes of two successive tosses. When a coin is tossed twice, the number of outcomes is four (Figure A-1(a)). The probabilities for the three events defined above are:

Event A: The probability of the event "two heads," A = {HH}, equals $p(A) = 1/4$.

Event B: The probability of the event "exactly one head," B = {HT, TH}, equals $p(B) = (1 + 1)/4 = 1/2$.

Event C: The probability of the event "at least one head," C = {HH, HT, TH}, equals $p(C) = (1 + 1 + 1)/4 = 3/4$.

The Principles of Probability Theory

Given a sample space S with finite number of elements (or, outcomes), we assume that a **probability** value $p(x)$ is attached to each element x in S, with the following properties

1. $0 \le p(x) \le 1$, for all $x \in S$
2. $\sum_{x \in S} p(x) = 1$

Given an event E, that is, a subset of S, the probability of E is defined as the sum of the probabilities of the outcomes in the event E

$$p(E) = \sum_{x \in E} p(x)$$

Therefore, the probability of the entire sample space is 1, and the probability of the null event is 0.

Events can also be combined. Given two events A and B, their **intersection** or **conjunction** is the event that consists of all outcomes common to both events, and is denoted as $\{A \text{ and } B\}$ or $\{A \cap B\}$. For example, the event C defined above ("at least one head") consists of three outcomes $C = \{\text{HH, HT, TH}\}$. We can define another event, "at least one tail," which also consists of three outcomes $D = \{\text{HT, TH, TT}\}$. The conjunction of C and D ("at least one head and at least one tail") consists of the outcomes HT and TH. (Note that this event is equivalent to the event "one head and one tail.") Such an event is called a **joint event** (or, compound event) and the probability of such and event is called a **joint probability** (or, compound probability).

Another type of event combination involves outcomes of either of two events. The **union** or **disjunction** of two events consists of all the outcomes in either event, denoted as $\{A \text{ or } B\}$ or $\{A \cup B\}$. For example, consider again the event "at least one head," $C = \{\text{HH, HT, TH}\}$, and the event "at least one tail," $D = \{\text{HT, TH, TT}\}$. The event "at least one head or at least one tail" consists of $\{A \text{ or } B\} = \{\text{HH, HT, TH, TT}\}$. This disjunction equals the entire sample space S, because there must be at least one head or at least one tail in any coin-toss experiment.

An important property of random events is *independence* of one another. When two events are independent, the occurrence of one of the events gives no information about the probability that the other event occurs. One event does not influence another, or stated formally, the events E and F are **independent** if and only if $p(E \text{ and } F) = p(E) \cdot p(F)$.

To illustrate the above concepts, consider the following scenario of two containers with black and white balls (Figure A-3). First you decide randomly from which container to draw a ball from, and then you draw a ball from the selected vessel. To decide the vessel to draw from, you roll a die. If the die comes up 1 or 2, you draw a ball from the jar; otherwise, you draw from the urn (i.e., when the die comes up 3, 4, 5, or 6). Let us define the following events: $VJ = \{\text{die roll outcomes: 1, 2}\}$, $VU = \{\text{die roll outcomes: 3, 4, 4, 6}\}$, $BB = \{\text{black ball taken}\}$, $BW = \{\text{white ball taken}\}$.

A joint event can be defined as $JB = \{\text{black ball taken from Jar}\}$, which is a conjunction of $VJ \cap BB$. Another joint event is $UW = \{\text{white ball taken from Urn}\}$. One can notice that VJ and BB are not independent, because the occurrence of one of the events gives useful information about the probability that the other event occurs. That is, we know that the probability of taking a black ball is high die comes up 1 or 2, because the fraction of black balls is greater in the jar than in the urn.

Many any problems are concerned with a numerical value associated with the outcome of an experiment. For example, we may be interested in the total number of packets that end up with errors when 100 packets are transmitted. To study problems of this type we introduce the concept of a random variable. A **random variable** is a function from the sample space of an experiment to the set of real numbers. That is, a random variable assigns a real number to each possible

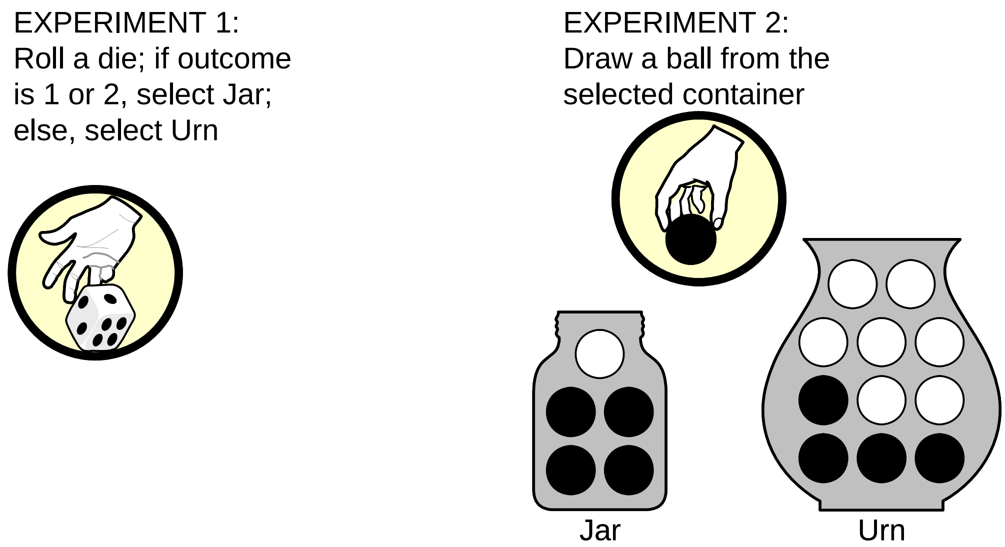

Figure A-3. Two experiments using a die, and a jar and urn with black and white balls.

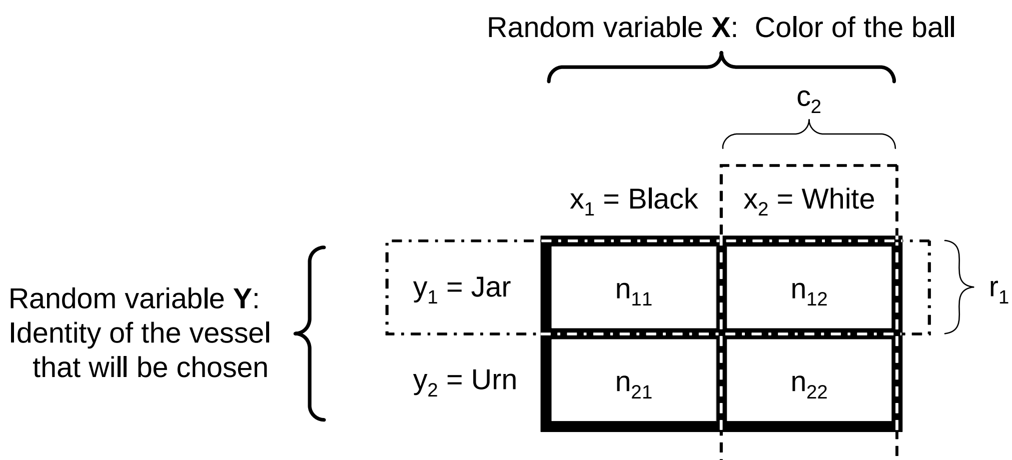

Figure A-4. Matrix of probabilities of random variables from Figure A-3.

outcome. (It is worth repeating that a random variable is a function; it is *not* a variable, and it is *not* random!)

In the example of Figure A-3, the identity of the vessel that will be chosen is a random variable, which we shall denote by Y. This random variable can take one of two possible values, namely *jar* or *urn*. Similarly, the color of the ball that will be drawn from the vessel is also a random variable and will be denoted by X. It can take either of the values *black* or *white*.

Consider the matrix representation in Figure A-4. Suppose that we run N times the experiments in Figure A-3 and record the outcomes in the matrix. Each cell of the matrix records the fraction of the total number of samples that turned out in the specific way. For example, n_{11} records the fraction of samples (out of the N total) where the selected vessel was *jar* and the color of the ball taken from the vessel was *black*. The **joint probability** of the random variables X and Y is the probability that X will take the value x_i and at the same time Y will take the value y_j, which is written as $p(X = x_i, Y = y_j)$. For example, in Figure A-4, $X = black$ and $Y = jar$. In a general case where variable X can take M different values, $X = x_i$, $i = 1, \ldots, M$, and variable Y can take L different values, $Y = y_j$, $j = 1, \ldots, L$, we can write the joint probability as

$$p(X = x_i, Y = y_j) = \frac{n_{ij}}{N} \tag{A.1}$$

(Of course, for this to hold, we are assuming that $N \to \infty$.) Similarly, the probability that X takes the value x_i (e.g., the probability that the ball color is *black*) regardless of the value of Y is written as $p(X = x_i)$ and is given by the fraction of the total number of points that fall in column i, so that

$$p(X = x_i) = \frac{c_i}{N} \tag{A.2}$$

The number of instances in column i in Figure A-4 is just the sum of the number of instances in each cell of that column. Therefore, we have $c_i = \sum_{j=1}^{L} n_{ij}$. If we plug this to equation (A.2) and then use (A.1), we will obtain

$$p(X = x_i) = \frac{\sum_{j=1}^{L} n_{ij}}{N} = \sum_{j=1}^{L} \frac{n_{ij}}{N} = \sum_{j=1}^{L} p(X = x_i, Y = y_j) \tag{A.3}$$

This is known as the **sum rule of probability**. The probability $p(X = x_i)$ is sometimes called the **marginal probability**, because it is obtained by marginalizing, or summing out, the other variables (in this case Y).

Let us fix the value of the random variable X so that $X = x_i$, and consider the fraction of such instances for which $Y = y_j$. In the example of Figure A-4, we could assume that X = *white* and consider the fraction of instances for which Y = *jar*. This is written as $p(Y = y_j \mid X = x_i)$ and is called **conditional probability** of $Y = y_j$ given $X = x_i$. It is obtained by finding the fraction of those points in column i that fall in cell i,j and is given as

$$p(Y = y_j \mid X = x_i) = \frac{n_{ij}}{c_i} \tag{A.4}$$

Starting with equation (A.1) and using equations (A.2) and (A.4), we can derive the following

$$\begin{aligned} p(X = x_i, Y = y_j) &= \frac{n_{ij}}{N} = \frac{n_{ij}}{c_i} \cdot \frac{c_i}{N} \\ &= p(Y = y_j \mid X = x_i) \cdot p(X = x_i) \end{aligned} \tag{A.5}$$

This relationship is known as the **product rule of probability**.

Statistics of Random Variables

If X is a *discrete random variable*, define $S_X = \{x_1, x_2, \ldots, x_N\}$ as the *range* of X. That is, the value of X belongs to S_X.

Probability mass function (PMF): $P_X(x) = P[X = x]$

Properties of X with $P_X(x)$ and S_X:

a) $P_X(x) \geq 0 \quad \forall x$

b) $\sum_{x \in S_X} P_X(x) = 1$

c) Given $B \subset S_X$, $P[B] = \sum_{x \in B} P_X(x)$

Define a and b as upper and lower bounds of X if X is a *continuous random variable*.

Cumulative distribution function (CDF): $F_X(x) = P[X \leq x]$

Probability density function (PDF): $f_X(x) = \dfrac{dF(x)}{dx}$

Properties of X with PDF $f_X(x)$:

a) $f_X(x) \geq 0 \quad \forall\, x$

b) $F_X(x) = \int_{-\infty}^{x} f_X(u) \cdot du$

c) $\int_{-\infty}^{\infty} f_X(x) \cdot dx = 1$

Expected value:

The *mean* or first moment.

Continuous RV case: $E[X] = \mu_X = \int_a^b x \cdot f_X(x) \cdot dx$

Discrete RV case: $E[X] = \mu_X = \sum_{k=1}^{N} x_k \cdot P_X(x_k)$

Variance:

Second moment minus first moment-squared: $Var[X] = E\left[(X - \mu_X)^2\right] = E\left[X^2\right] - \mu_X^2$

Continuous RV case: $E\left[X^2\right] = \int_a^b x^2 \cdot f_X(x) \cdot dx$

Discrete RV case: $E\left[X^2\right] = \sum_{k=1}^{N} x_k^2 \cdot P_X(x_k)$

Standard Deviation: $\sigma_X = \sqrt{Var[X]}$

Bayes' Theorem

Consider the following scenario of two containers with black and white balls (Figure A-5). First you decide randomly from which container to draw a ball, then you draw a ball from the selected container, and finally you report the ball color to your friend. To decide the container to draw from, you roll a die. If the die comes up 1 or 2, you draw a ball from the jar; otherwise, you draw

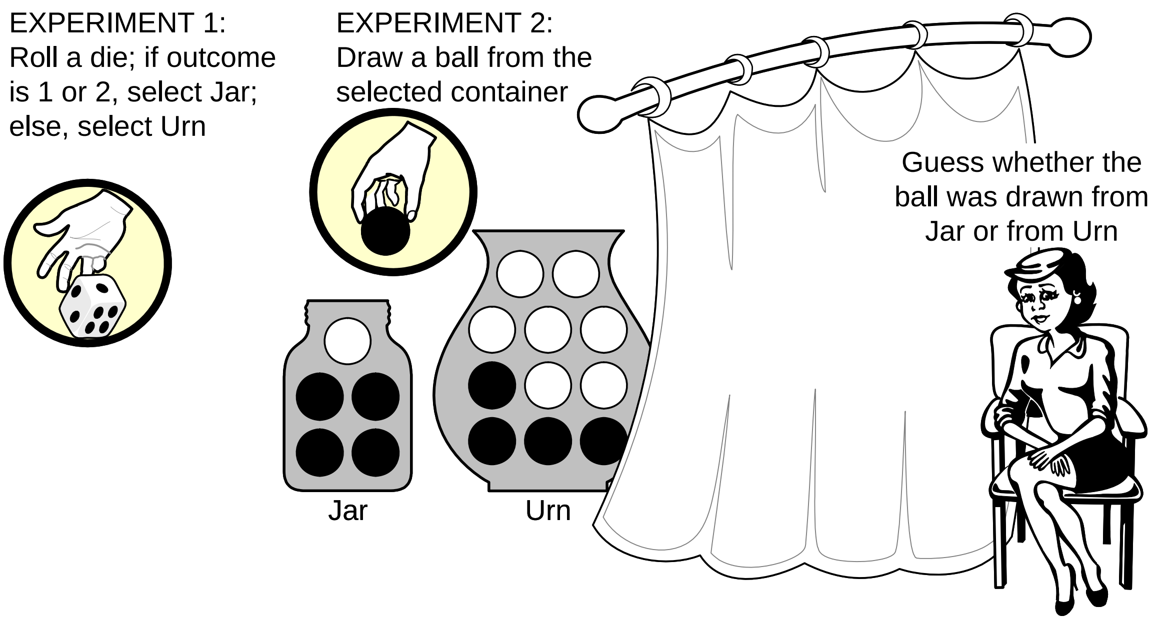

Figure A-5. Two experiments use a die, and a jar and urn with black and white balls. The person behind a curtain is trying to guess from which vessel a ball is drawn, given the color of the ball.

from the urn (i.e., when the die comes up 3, 4, 5, or 6). Your friend is sitting behind a curtain and cannot observe which container the ball was drawn from. Your friend needs to answer a question such as: "given that a black ball was drawn, is it more likely that the ball was drawn from the jar or urn?". In other words, given that a black ball was drawn, what is the probability that it came from the jar?

On one hand, we know that the fraction of black balls is greater in the jar than in the urn. On the other hand, we know that as the result of the roll of the die, it is twice as likely that you have drawn the ball from the urn. How can we combine the evidence from our sample (black drawing outcome) with our prior belief based on the roll of the die? To help answer this question, consider again the representation in Figure A-4.

$$p(Y \mid X) = \frac{p(X \mid Y) \cdot p(Y)}{p(X)} \tag{A.7}$$

Suppose that an experiment can have only two possible outcomes. For example, when a coin is flipped, the possible outcomes are heads and tails. Each performance of an experiment with two possible outcomes is called a **Bernoulli trial**, after the Swiss mathematician James Bernoulli (1654-1705). In general, a possible outcome of a Bernoulli trial is called a *success* or a *failure*. If p is the probability of a success and q is the probability of a failure, it follows that $p + q = 1$.

Many problems can be solved by determining the probability of k successes when an experiment consists of n mutually independent Bernoulli trials. (Bernoulli trials are **mutually independent** if the conditional probability of success on any given trial is p, given any information whatsoever about the outcomes of other trials.)

The probability of exactly k successes in n independent Bernoulli trials, with probability of success p and probability of failure $q = 1 - p$, is $b(k; n, p) = C(n, k) \cdot p^k \cdot q^{n-k}$. When $b(k; n, p)$ is considered as a function of k, we call this function the **binomial distribution**.

Random Processes

A *process* is a naturally occurring or designed sequence of operations or events, possibly taking up time, space, expertise or other resource, which produces some outcome. A process may be identified by the changes it creates in the properties of one or more objects under its influence.

A function may be thought of as a computer program or mechanical device that takes the characteristics of its input and produces output with its own characteristics. Every process may be defined functionally and every process may be defined as one or more functions.

An example random process that will appear later in the text is Poisson process. It is usually employed to model arrivals of people or physical events as occurring at random points in time. *Poisson process* is a counting process for which the times between successive events are *independent and identically distributed* (IID) exponential random variables. For a Poisson process, the number of arrivals in any interval of length τ is Poisson distributed with a parameter $\lambda \cdot \tau$. That is, for all t, $\tau > 0$,

$$P\{A(t+\tau) - A(t) = n\} = e^{-\lambda\tau} \frac{(\lambda\tau)^n}{n!}, \qquad n = 0,1,... \tag{A.8}$$

The average number of arrivals within an interval of length τ is $\lambda\tau$ (based on the mean of the Poisson distribution). This implies that we can view the parameter λ as an arrival rate (average number of arrivals per unit time). If X represents the time between two arrivals, then $P(X > x)$, that is, the probability that the interarrival time is longer than x, is given by $e^{-x/\lambda}$. An interesting property of this process is that it is *memoryless*: the fact that a certain time has elapsed since the last arrival gives us no indication about how much longer we must wait before the next event arrives. An example of the Poisson distribution is shown in Figure A-6.

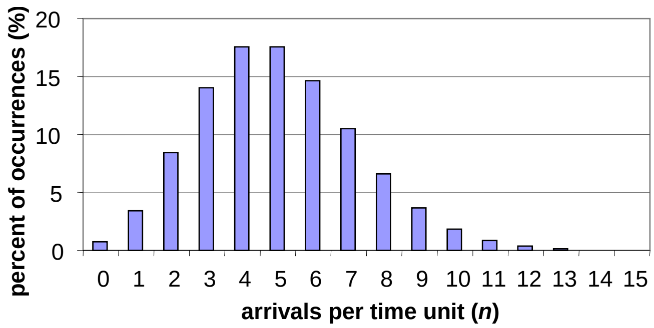

Figure A-6. The histogram of the number of arrivals per unit of time (τ = 1) for a Poisson process with average arrival rate λ = 5.

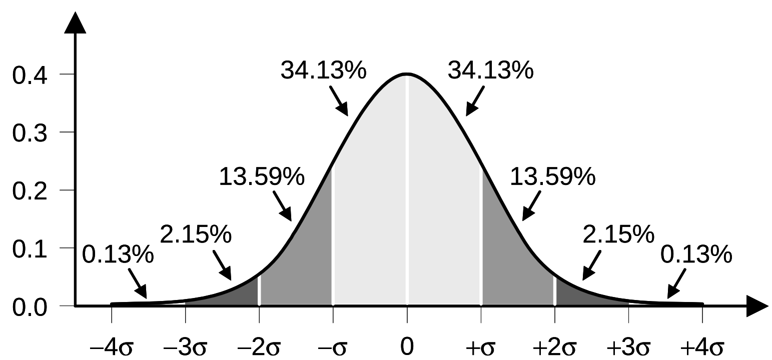

Figure A-7. Areas between selected points under the normal curve.

This model is not entirely realistic for many types of sessions and there is a great amount of literature which shows that it fails particularly at modeling the LAN traffic. However, such simple models provide insight into major tradeoffs involved in network design, and these tradeoffs are often obscured in more realistic and complex models.

Markov process is a random process with the property that the probabilities of occurrence of the various possible outputs depend upon one or more of the preceding outputs.

Statistics Review

Proportions of Area Under the Normal Curve

Figure A-7 shows a coarse partition of areas under the normal curve $N(\mu,\sigma)$. Statistical tables are often used to obtain finer partitioning, as shown in Table A-1. To use this table, it is necessary to convert the raw magnitude to a so-called *z*-score. The *z*-score is a standard deviate that allows for using the standard normal distribution $N(0,1)$, the one which has a mean μ = 0.0, a standard deviation σ = 1.0, and a total area under the curve equal to 1.0.

The values in Table A-1 represent the proportion of area under the standard normal curve. The table contains *z*-scores between 0.0 and 4.00 (i.e., four standard deviations, or $4\times\sigma$), with 0.01 increments. Because the normal distribution is symmetrical, the table represents *z*-scores ranging between −4.00 and 4.00.

Figure A-8 illustrates how to read Table A-1. Suppose you want to know how big is the area under the normal curve from the mean up to $1.5\times\sigma$, i.e., $z = 1.50$ and how much remains beyond. First, we look up Column A and find $z = 1.50$. Second, we read the associated values in Columns B and C, which represent the area between mean and *z*, and the area beyond *z*, respectively.

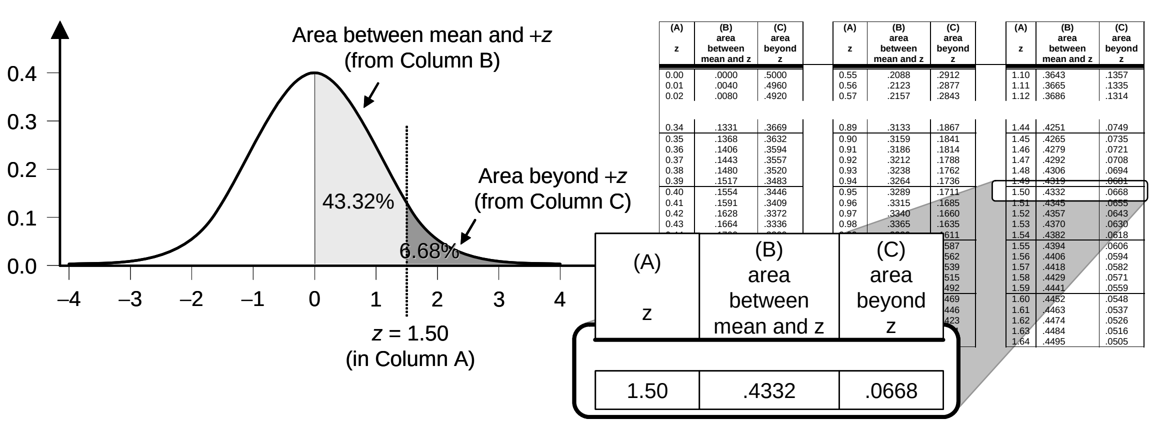

Figure A-8. Illustration of how to read Table A-1 on the next page.

Table A-1: Proportions of area under the normal curve. (*continued below*)

(A) z	(B) area between mean and z	(C) area beyond z	(A) z	(B) area between mean and z	(C) area beyond z	(A) z	(B) area between mean and z	(C) area beyond z
0.00	.0000	.5000	0.55	.2088	.2912	1.10	.3643	.1357
0.01	.0040	.4960	0.56	.2123	.2877	1.11	.3665	.1335
0.02	.0080	.4920	0.57	.2157	.2843	1.12	.3686	.1314
0.03	.0120	.4880	0.58	.2109	.2810	1.13	.3708	.1292
0.04	.0160	.4840	0.59	.2224	.2776	1.14	.3729	.1271
0.05	.0199	.4801	0.60	.2257	.2743	1.15	.3749	.1251
0.06	.0239	.4761	0.61	.2291	.2709	1.16	.3770	.1230
0.07	.0279	.4721	0.62	.2324	.2676	1.17	.3790	.1210
0.08	.0319	.4681	0.63	.2357	.2643	1.18	.3810	.1190
0.09	.0359	.4641	0.64	.2389	.2611	1.19	.3830	.1170
0.10	.0398	.4602	0.65	.2422	.2578	1.20	.3849	.1151
0.11	.0438	.4562	0.66	.2454	.2564	1.21	.3869	.1131
0.12	.0478	.4522	0.67	.2486	.2514	1.22	.3888	.1112
0.13	.0517	.4483	0.68	.2517	.2483	1.23	.3907	.1093
0.14	.0557	.4443	0.69	.2549	.2451	1.24	.3925	.1075
0.15	.0596	.4404	0.70	.2580	.2420	1.25	.3944	.1056
0.16	.0636	.4364	0.71	.2611	.2389	1.26	.3962	.1038
0.17	.0675	.4325	0.72	.2642	.2358	1.27	.3980	.1020
0.18	.0714	.4286	0.73	.2673	.2327	1.28	.3997	.1003
0.19	.0753	.4247	0.74	.2704	.2296	1.29	.4015	.0985
0.20	.0793	.4207	0.75	.2734	.2266	1.30	.4032	.0968
0.21	.0832	.4168	0.76	.2764	.2236	1.31	.4049	.0951
0.22	.0871	.4129	0.77	.2794	.2206	1.32	.4066	.0934
0.23	.0910	.4090	0.78	.2823	.2177	1.33	.4082	.0918
0.24	.0948	.4052	0.79	.2852	.2148	1.34	.4099	.0901
0.25	.0987	.4013	0.80	.2881	.2119	1.35	.4115	.0885
0.26	.1026	.3974	0.81	.2910	.2090	1.36	.4131	.0869
0.27	.1064	.3936	0.82	.2939	.2061	1.37	.4147	.0853
0.28	.1103	.3897	0.83	.2967	.2033	1.38	.4162	.0838
0.29	.1141	.3859	0.84	.2995	.2005	1.39	.4177	.0823
0.30	.1179	.3821	0.85	.3023	.1977	1.40	.4192	.0808
0.31	.1217	.3783	0.86	.3051	.1949	1.41	.4207	.0793
0.32	.1255	.3745	0.87	.3078	.1922	1.42	.4222	.0778
0.33	.1293	.3707	0.88	.3106	.1894	1.43	.4236	.0764
0.34	.1331	.3669	0.89	.3133	.1867	1.44	.4251	.0749
0.35	.1368	.3632	0.90	.3159	.1841	1.45	.4265	.0735
0.36	.1406	.3594	0.91	.3186	.1814	1.46	.4279	.0721
0.37	.1443	.3557	0.92	.3212	.1788	1.47	.4292	.0708
0.38	.1480	.3520	0.93	.3238	.1762	1.48	.4306	.0694
0.39	.1517	.3483	0.94	.3264	.1736	1.49	.4319	.0681
0.40	.1554	.3446	0.95	.3289	.1711	1.50	.4332	.0668
0.41	.1591	.3409	0.96	.3315	.1685	1.51	.4345	.0655
0.42	.1628	.3372	0.97	.3340	.1660	1.52	.4357	.0643
0.43	.1664	.3336	0.98	.3365	.1635	1.53	.4370	.0630
0.44	.1700	.3300	0.99	.3389	.1611	1.54	.4382	.0618
0.45	.1736	.3264	1.00	.3413	.1587	1.55	.4394	.0606
0.46	.1772	.3228	1.01	.3438	.1562	1.56	.4406	.0594
0.47	.1808	.3192	1.02	.3461	.1539	1.57	.4418	.0582
0.48	.1844	.3156	1.03	.3485	.1515	1.58	.4429	.0571
0.49	.1879	.3121	1.04	.3508	.1492	1.59	.4441	.0559
0.50	.1915	.3085	1.05	.3531	.1469	1.60	.4452	.0548
0.51	.1950	.3050	1.06	.3554	.1446	1.61	.4463	.0537
0.52	.1985	.3015	1.07	.3577	.1423	1.62	.4474	.0526
0.53	.2019	.2981	1.08	.3599	.1401	1.63	.4484	.0516
0.54	.2054	.2946	1.09	.3621	.1379	1.64	.4495	.0505

Table A-1 (*continued*)

(A) z	(B) area between mean and z	(C) area beyond z	(A) z	(B) area between mean and z	(C) area beyond z	(A) z	(B) area between mean and z	(C) area beyond z
1.65	.4505	.0495	2.22	.4868	.0132	2.79	.4974	.0026
1.66	.4515	.0485	2.23	.4871	.0129	2.80	.4974	.0026
1.67	.4525	.0475	2.24	.4875	.0125	2.81	.4975	.0025
1.68	.4535	.0465	2.25	.4878	.0122	2.82	.4976	.0024
1.69	.4545	.0455	2.26	.4881	.0119	2.83	.4977	.0023
1.70	.4554	.0446	2.27	.4884	.0116	2.84	.4977	.0023
1.71	.4564	.0436	2.28	.4887	.0113	2.85	.4978	.0022
1.72	.4573	.0427	2.29	.4890	.0110	2.86	.4979	.0021
1.73	.4582	.0418	2.30	.4893	.0107	2.87	.4979	.0021
1.74	.4591	.0409	2.31	.4896	.0104	2.88	.4980	.0020
1.75	.4599	.0401	2.32	.4898	.0102	2.89	.4981	.0019
1.76	.4608	.0392	2.33	.4901	.0099	2.90	.4981	.0019
1.77	.4616	.0384	2.34	.4904	.0096	2.91	.4982	.0018
1.78	.4625	.0375	2.35	.4906	.0094	2.92	.4982	.0018
1.79	.4633	.0367	2.36	.4909	.0091	2.93	.4983	.0017
1.80	.4641	.0359	2.37	.4911	.0089	2.94	.4984	.0016
1.81	.4649	.0351	2.38	.4913	.0087	2.95	.4984	.0016
1.82	.4656	.0344	2.39	.4916	.0084	2.96	.4985	.0015
1.83	.4664	.0336	2.40	.4918	.0082	2.97	.4985	.0015
1.84	.4671	.0329	2.41	.4920	.0080	2.98	.4986	.0014
1.85	.4678	.0322	2.42	.4922	.0078	2.99	.4986	.0014
1.86	.4686	.0314	2.43	.4925	.0075	3.00	.4987	.0013
1.87	.4693	.0307	2.44	.4927	.0073	3.01	.4987	.0013
1.88	.4699	.0301	2.45	.4929	.0071	3.02	.4987	.0013
1.89	.4706	.0294	2.46	.4931	.0069	3.03	.4988	.0012
1.90	.4713	.0287	2.47	.4932	.0068	3.04	.4988	.0012
1.91	.4719	.0281	2.48	.4934	.0066	3.05	.4989	.0011
1.92	.4726	.0274	2.49	.4936	.0064	3.06	.4989	.0011
1.93	.4732	.0268	2.50	.4938	.0062	3.07	.4989	.0011
1.94	.4738	.0262	2.51	.4940	.0060	3.08	.4990	.0010
1.95	.4744	.0256	2.52	.4941	.0059	3.09	.4990	.0010
1.96	.4750	.0250	2.53	.4943	.0057	3.10	.4990	.0010
1.97	.4756	.0244	2.54	.4945	.0055	3.11	.4991	.0009
1.98	.4761	.0239	2.55	.4946	.0054	3.12	.4991	.0009
1.99	.4767	.0233	2.56	.4948	.0052	3.13	.4991	.0009
2.00	.4772	.0228	2.57	.4949	.0051	3.14	.4992	.0008
2.01	.4778	.0222	2.58	.4951	.0049	3.15	.4992	.0008
2.02	.4783	.0217	2.59	.4952	.0048	3.16	.4992	.0008
2.03	.4788	.0212	2.60	.4953	.0047	3.17	.4992	.0008
2.04	.4793	.0207	2.61	.4955	.0045	3.18	.4993	.0007
2.05	.4798	.0202	2.62	.4956	.0044	3.19	.4993	.0007
2.06	.4803	.0197	2.63	.4957	.0043	3.20	.4993	.0007
2.07	.4808	.0192	2.64	.4959	.0041	3.21	.4993	.0007
2.08	.4812	.0188	2.65	.4960	.0040	3.22	.4994	.0006
2.09	.4817	.0183	2.66	.4961	.0039	3.23	.4994	.0006
2.10	.4821	.0179	2.67	.4962	.0038	3.24	.4994	.0006
2.11	.4826	.0174	2.68	.4963	.0037	3.25	.4994	.0006
2.12	.4830	.0170	2.69	.4964	.0036	3.30	.4995	.0005
2.13	.4834	.0166	2.70	.4965	.0035	3.35	.4996	.0004
2.14	.4838	.0162	2.71	.4966	.0034	3.40	.4997	.0003
2.15	.4842	.0158	2.72	.4967	.0033	3.45	.4997	.0003
2.16	.4846	.0154	2.73	.4968	.0032	3.50	.4998	.0002
2.17	.4850	.0150	2.74	.4969	.0031	3.60	.4998	.0002
2.18	.4854	.0146	2.75	.4970	.0030	3.70	.4999	.0001
2.19	.4857	.0143	2.76	.4971	.0029	3.80	.4999	.0001
2.20	.4861	.0139	2.77	.4972	.0028	3.90	.49995	.00005
2.21	.4864	.0136	2.78	.4973	.0027	4.00	.49997	.00003

References

1. I. Aad and C. Castelluccia, "Differentiation mechanisms for IEEE 802.11," *Proceedings of the IEEE Infocom 2001*, April 2001.
2. M. Allman, V. Paxson, and W. R. Stevens, "TCP congestion control," IETF Request for Comments RFC-2581, April 1999. Online at: http://www.apps.ietf.org/rfc/rfc2581.html
3. G. Almes, S. Kalidindi, and M. Zekauskas, "A one-way delay metric for IPPM," IETF Request for Comments RFC-2679, September 1999 (a). Online at: http://www.apps.ietf.org/rfc/rfc2679.html
4. G. Almes, S. Kalidindi, and M. Zekauskas, "A round-trip delay metric for IPPM," IETF Request for Comments RFC-2681, September 1999 (b). Online at: http://www.apps.ietf.org/rfc/rfc2681.html
5. G. Anastasi and L. Lenzini, "QoS provided by the IEEE 802.11 wireless LAN to advanced data applications: A simulation analysis," *ACM Wireless Networks*, vol. 6, no. 2, pp. 99-108, 2000.
6. D. Andersen, D. Bansal, D. Curtis, S. Seshan, and H. Balakrishnan, "System support for bandwidth management and content adaptation in Internet applications," *Proceedings of the USENIX OSDI Conference*, San Diego, CA, October 2000.
7. D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of data-handling system with multiple sources," *Bell System Technical Journal*, vol. 61, no. 8, pp. 1871-1894, 1982.
8. C. Aras, J. F. Kurose, D. Reeves, and H. Schulzrinne, "Real-time communication in packet-switched networks," *Proceedings of the IEEE*, vol. 82, no. 1, pp. 122-139, January 1994.
9. B. Badrinath, A. Fox, L. Kleinrock, G. Popek, P. Reiher, and M. Satyanarayanan, "A conceptual framework for network and client adaptation," *Mobile Networks and Applications (MONET)*, ACM / Kluwer Academic Publishers, vol. 5, pp. 221-231, 2000.
10. H. Balakrishnan, V. N. Padmanabhan, S. Seshan, and R. H. Katz, "A comparison of mechanisms for improving TCP performance over wireless links," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp.756-769, December 1997.
11. A. J. Ballardie, P. F. Francis, and J. Crowcroft, "Core based trees," *ACM SIGCOMM Computer Communication Review*, vol. 23, no. 4, pp. 85-95, August 1993.
12. P. Baran, "Introduction to Distributed Communications Network," RAND Corporation, Memorandum RM-3420-PR, August 1964. Online at: http://www.rand.org/publications/RM/baran.list.html
13. R. van der Berg, "How the 'Net works: An introduction to peering and transit," Social Science Electronic Publishing, Inc., September 2, 2008. Online at: http://ssrn.com/abstract=1443245
14. D. Bertsekas and R. Gallagher. *Data Networks*. 2nd edition, Prentice Hall, Upper Saddle River, NJ, 1992.
15. S. N. Bhatti and J. Crowcroft, "QoS-sensitive flows: Issues in IP packet handling," *IEEE Internet Computing*, vol. 4, no. 4, pp. 48-57, July 2000.
16. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services," IETF Request for Comments RFC-2475, December 1998. Online at: http://www.apps.ietf.org/rfc/rfc2475.html

17. M. S. Blumenthal and D. D. Clark, "Rethinking the design of the Internet: The end-to-end arguments vs. the brave new world," *ACM Transactions on Internet Technology*, vol. 1, no. 1, pp. 70-109, August 2001.

18. J. C. Bolot, "End-to-end packet delay and loss behavior in the Internet," *Proceedings of the ACM SIGCOMM '93 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 1993.

19. B. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang, "Recommendations on queue management and congestion avoidance in the Internet," IETF Request for Comments RFC-2309, April 1998. Online at: http://www.apps.ietf.org/rfc/rfc2309.html

20. L. S. Brakmo and L. L. Peterson, "TCP Vegas: End-to-end congestion avoidance on a global internet," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 8, pp. 1465-1480, October 1995.

21. L. Breslau, E. W. Knightly, S. Shenker, I. Stoica, and H. Zhang, "Endpoint Admission Control: Architectural Issues and Performance," *Proceedings of the ACM SIGCOMM 2000 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 2000.

22. M. Buettner and D. Wetherall, "An empirical study of UHF RFID performance," *Proceedings of the 14th ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2008)*, pp. 223-234, San Francisco, CA, September 2008.

23. Z. Cao and E. Zegura, "Utility max-min: An application-oriented bandwidth allocation scheme," *Proc. IEEE InfoCom*, vol. 2, pp. 793-801, 1999.

24. R. L. Carter and M. E. Crovella, "Measuring bottleneck link speed in packet-switched networks," Technical Report TR-96-006, Department of Computer Science, Boston University, March 1996. Online at: http://www.cs.bu.edu/faculty/crovella/papers.html

25. D. Chalmers and M. Sloman, "A survey of quality of service in mobile computing environments," *IEEE Communications Surveys*, vol. 2, no. 2, 1999.

26. H. S. Chhaya and S. Gupta, "Performance of asynchronous data transfer methods of IEEE 802.11 MAC protocol," *IEEE Personal Communications*, vol. 3, no. 5, October 1996.

27. M. Christiansen, K. Jeffay, D. Ott, and F.D. Smith, "Tuning RED for Web traffic," *IEEE/ACM Transactions on Networking*, vol. 9, no. 3, pp. 249-264, June 2001.

28. Cisco Systems Inc., "Interface Queue Management," Cisco white paper, posted August 1995. Online at: http://www.cisco.com/warp/public/614/16.html

29. Cisco Systems Inc., "Performance Measurements of Advanced Queuing Techniques in the Cisco IOS," Cisco white paper, posted July 1999. Online at: http://www.cisco.com/warp/public/614/15.html

30. Cisco Systems Inc., "Advanced QoS Services for the Intelligent Internet," Cisco white paper, posted June 2006. Online at: http://www.cisco.com/warp/public/cc/pd/iosw/ioft/ioqo/tech/qos_wp.htm

31. D. D. Clark, "Design philosophy of the DARPA Internet protocols," *Proceedings of the ACM SIGCOMM*, pp. 106-114, Stanford, CA, August 1988.

32. D. D. Clark, "The structuring of systems using upcalls," *Proceedings of the 10th ACM Symposium on Operating Systems*, pp. 171-180, December 1985.

33. D. D. Clark and W. Fang, "Explicit allocation of best-effort packet delivery service," *IEEE/ACM Transactions on Networking*, vol. 6, no. 4, pp. 362-373, August 1998.

34. D. D. Clark, S. Shenker, and L. Zhang, "Supporting real-time applications in an integrated services packet network: Architecture and mechanisms," *Proc. SIGCOMM '92*, Baltimore, MD, August 1992.

35. D. D. Clark and D. L. Tennenhouse, "Architectural considerations for a new generation of protocols," *ACM SIGCOMM Computer Communications Review*, vol. 20, no. 4, pp. 200-208, August 1990.

36. D. E. Comer, *Internetworking With TCP/IP, Volume I: Principles, Protocols, and Architecture*, 5th Edition, Pearson Prentice Hall, Upper Saddle River, NJ, 2006.

37. B. P. Crow, I. Widjaja, L. G. Kim, and P. T. Sakai, "IEEE 802.11 wireless local area networks," *IEEE Communications Magazine*, vol. 35, no. 9, pp. 116-126, September 1997.

38. W. Dapeng and R. Negi, "Effective capacity: A wireless link model for support of quality of service," *IEEE Transactions on Wireless Communications*, vol. 2, no. 4, pp. 630-643, July 2003.

39. B. S. Davie and Y. Rekhter, *MPLS: Technology and Applications*, Morgan Kaufmann Publishers (Academic Press), San Francisco, CA, 2000.

40. S. Deering and R. Hinden, "Internet Protocol, version 6 (IPv6) specification," IETF Request for Comments RFC-2460, December 1998. Online at: http://www.apps.ietf.org/rfc/rfc2460.html

41. L. De Ghein, *MPLS Fundamentals: A Comprehensive Introduction to MPLS Theory and Practice*, Cisco Press, Indianapolis, IN, 2007.

42. C. Demichelis and P. Chimento, "IP packet delay variation metric for IP performance metrics (IPPM)," IETF Request for Comments RFC-3393, November 2002. Online at: http://www.apps.ietf.org/rfc/rfc3393.html

43. D-J. Deng and R-S. Chang, "A priority scheme for IEEE 802.11 DCF access method," *IEICE Transactions on Communications*, E82-B-(1), January 1999.

44. J. van Duuren, "Fault-free digital radio communication and Hendrik C. A. van Duuren," *Proceedings of the IEEE*, vol. 89, no. 10, pp. 1540-1542, October 2001.

45. M. DuVal and T. Siep, "High Rate WPAN for Video," IEEE document: IEEE 802.15-00/029, Submission date: 6 March 2000. Online at: http://grouper.ieee.org/groups/802/15/pub/2000/Mar00/00029r0P802-15_CFA-Response-High-Rate-WPAN-for-Video-r2.ppt

46. I. Elhanany, M. Kahane, and D. Sadot, "Packet scheduling in next-generation multiterabit networks," *IEEE Computer*, vol. 34, no. 4, pp. 104-106, April 2001.

47. G. Fairhurst and L. Wood, "Advice to link designers on link Automatic Repeat reQuest (ARQ)," IETF Request for Comments RFC-3366, August 2002. Online at: http://www.apps.ietf.org/rfc/rfc3366.html

48. C. Ferran and S. Watts, "Videoconferencing in the field: A heuristic processing model," *Management Science*, vol. 54, no. 9, pp. 1565-1578, September 2008.

49. V. Firou, J. Le Boudec, D. Towsley, and Z. Zhang, "Theories and models for Internet quality of service," *Proceedings of the IEEE (Special Issue on Internet Technology)*, August 2002.

50. S. Floyd, "Congestion control principles," IETF Request for Comments RFC-2914, September 2000. Online at: http://www.apps.ietf.org/rfc/rfc2914.html

51. S. Floyd, T. Henderson, and A. Gurtov, "The NewReno modification to TCP's fast recovery algorithm," IETF Request for Comments RFC-3782, April 2004. Online at: http://www.rfc-editor.org/rfc/rfc3782.txt

52. S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397-413, August 1993.

53. Z. Fu, H. Luo, P. Zerfos, S. Lu, L. Zhang, and M. Gerla, "The impact of multihop wireless channel on TCP performance," *IEEE Transactions on Mobile Computing*, vol. 4, no. 2, pp. 209-221, March 2005.

54. F. C. Gärtner, "A Survey of Self-Stabilizing Spanning-Tree Construction Algorithms," Swiss Federal Institute of Technology (EPFL), Technical Report IC/2003/38, June 10, 2003. Online at: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.6.3874

55. P. Goyal, S. S. Lam, and H. M. Vin, "Determining end-to-end delay bounds in heterogeneous networks," *ACM Multimedia Systems*, vol. 5, no. 3, pp. 157-163, 1997.
[An earlier version of this paper appeared in *Proceedings of the Fifth International Workshop on*

Network and Operating System Support for Digital Audio and Video (NOSSDAV '95), Durham, NH, pp. 287-298, April 1995.]

56. J. P. Gray, "Line control procedures," *Proceedings of the IEEE*, vol. 60, no. 11, pp. 1301-1312, November 1972.

57. C. Greenhalgh, S. Benford, and G. Reynard, "A QoS architecture for collaborative virtual environments," *Proceedings of the ACM Multimedia Conference*, pp.121-130, 1999.

58. L. He and J. Walrand, "Pricing and revenue sharing strategies for Internet service providers," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 5, pp. 942-951, May 2006.

59. J. Hoe, "Improving the start-up behavior of a congestion control scheme for TCP," *Proceedings of the ACM SIGCOMM 1996 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, August 1996.

60. G. Holland and N. Vaidya, "Analysis of TCP performance over mobile ad hoc networks," *Proceedings of the 5th ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 1999)*, pp. 219-230, Seattle, WA, August 1999.

61. N. Hu and P. Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," *IEEE Journal on Selected Areas in Communications*, August 2003.

62. C. Huitema, *IPv6: The New Internet Protocol*, 2nd Edition, Prentice-Hall, Inc., Upper Saddle River, NJ, 1998.

63. N. C. Hutchinson and L. L. Peterson, "The *x*-kernel: An architecture for implementing network protocols," *IEEE Transactions on Software Engineering*, vol. 17, no. 1, pp. 64-76, January 1991.

64. D. S. Isenberg, "The dawn of the stupid network," *ACM netWorker*, vol. 2, no. 1, pp. 24-31, February/March 1998. Online at: http://www.isen.com/papers/Dawnstupid.html
An older version of this paper, "Rise of the stupid network," is available online here: http://www.rageboy.com/stupidnet.html

65. V. Jacobson, "Congestion avoidance and control," *ACM Computer Communication Review*, vol. 18, no. 4, pp. 314-329, August 1988. Online at: ftp://ftp.ee.lbl.gov/papers/congavoid.ps.Z

66. M. Jain and C. Dovrolis, "Pathload: A measurement tool for end-to-end available bandwidth," *Proceedings of the 3rd Passive and Active Measurements Workshop*, Fort Collins, CO, March 2002.

67. R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. John Wiley & Sons, New York, NY, 1991.

68. C. Jin, D. Wei, S. H. Low, G. Buhrmaster, J. Bunn, D. H. Choe, R. L. A. Cottrell, J. C. Doyle, W. Feng, O. Martin, H. Newman, F. Paganini, S. Ravot, and S. Singh, "FAST TCP: From Theory to Experiments," Submitted to *IEEE Communications Magazine*, April 1, 2003. Online at: http://netlab.caltech.edu/FAST/publications.html

69. R. Johari and J. N. Tsitsiklis, "Routing and peering in a competitive Internet," *Proceedings of the 2004 IEEE Conference on Decision and Control*, Bahamas, December 2004.
Extended version: MIT technical report LIDS-P-2570, January 2003. Online at: http://citeseer.ist.psu.edu/572286.html

70. D. Kahneman, *Attention and Effort*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1973.

71. S. Kandula et al. "Walking the tightrope: Responsive yet stable traffic engineering," *Proceedings of the ACM SIGCOMM 2005 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*.

72. V. Kanodia, C. Li, A. Sabharwal, B. Sadeghi, and E. Knightly, "Distributed priority scheduling and medium access in ad hoc networks," *Wireless Networks*, vol. 8, no. 5, pp. 455-466, 2002.

73. D. Katabi, M. Handley, and C. Rohrs, "Congestion control for future high bandwidth-delay product networks," *Proceedings of the ACM SIGCOMM 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, August 2002.

74. R. Katz, "Adaptation and mobility in wireless information systems," *IEEE Personal Communications*, vol. 1, no. 1, pp. 6-17, 1994.

75. S. Keshav, *An Engineering Approach to Computer Networking: ATM Networks, the Internet, and the Telephone Network*. Addison-Wesley Publ. Co., Reading, MA, 1997.

76. D. Kiwior, J. Kingston, and A. Spratt, "PathMon, a methodology for determining available bandwidth over an unknown network," *Proceedings of the 2004 IEEE/Sarnoff Symposium on Advances in Wired and Wireless Communication*, pp. 27-30, Princeton, NJ, April 2004.

77. S. Keshav and R. Sharma, "Issues and trends in router design," *IEEE Communications Magazine*, vol. 36, no. 5, pp. 144-151, May 1998.

78. V. P. Kumar, T. V. Lakshman, and D. Stiliadis, "Beyond best effort: Router architectures for the differentiated services of tomorrow's Internet," *IEEE Communications Magazine*, vol. 36, no. 5, pp. 152-164, May 1998.

79. J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach*. 5th edition. Pearson Education, Inc. (Addison-Wesley), Boston, MA, 2010.

80. K. Lai and M. Baker, "Measuring bandwidth," *Proceedings of the Conference on Computer Communications (IEEE INFOCOM '99)*, pp. 235-245, New York, NY, March 1999.

81. K. Lai and M. Baker, "Measuring link bandwidths using a deterministic model of packet delay," *Proceedings of the ACM SIGCOMM 2000 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, Stockholm, Sweden, August 2000.

82. E. D. Lazowska, J. Zahorjan, G. S. Graham, and K. C. Sevcik, *Quantitative System Performance: Computer System Analysis Using Queuing Network Models*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1984. Online at: http://www.cs.washington.edu/homes/lazowska/qsp/

83. R. R.-F. Liao and A. T. Campbell, "A utility-based approach for quantitative adaptation in wireless packet networks," *Wireless Networks*, vol. 7, no. 5, pp. 541-557, 2001.

84. W. Liu, "Focusing on desirability: The effect of decision interruption and suspension on preferences," *Journal of Consumer Research*. (Forthcoming), December 2008. Online at: http://www.anderson.ucla.edu/x15549.xml

85. Y. Mao and L. K. Saul, "Modeling distances in large-scale networks by matrix factorization," *Proceedings of the Second Internet Measurement Conference (IMC-04)*, pp. 278-287, Sicily, Italy, 2004. Online at: https://wiki.planet-lab.org/bin/view/Planetlab/AnnotatedBibliography

86. S. Mascolo, C. Casetti, M. Gerla, M. Y. Sanadidi, and R. Wang, "TCP Westwood: Bandwidth estimation for enhanced transport over wireless links," *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking (MobiCom 2001)*, Rome, Italy, pp. 287-296, August 2001.

87. M. L. Massie, B. N. Chun, and D. E. Culler, "The Ganglia distributed monitoring system: Design, implementation, and experience," *Parallel Computing*, vol. 30, no. 7, July 2004.

88. J. M. McQuillan, I. Richer, and E. C. Rosen, "The new routing algorithm for the ARPANet," IEEE *IEEE Transactions on Communications*, vol. 28, no. 5, pp. 711-719, 1980.

89. C. S. R. Murthy and B. S. Manoj, *Ad Hoc Wireless Networks: Architectures and Protocols*. Prentice Hall PTR, Upper Saddle River, NJ, 2004.

90. J. Nagle, "Congestion control in IP/TCP internetworks," IETF Request for Comments RFC-896, January 1984. Online at: http://www.rfc-editor.org/rfc/rfc896.txt

91. J. Nagle, "On packet switches with infinite storage," *IEEE Transactions on Communications*, vol. 35, no. 4, pp. 435-438, April 1987.

92. K. Nahrstedt and J. M. Smith, "The QoS broker," *IEEE Multimedia*, vol. 2, no. 1, pp. 53-67, 1995.

93. K. Nahrstedt, D. Xu, D. Wichadakul, and B. Li, "QoS-aware middleware for ubiquitous and heterogeneous environments," *IEEE Communications Magazine*, vol. 39, no. 11, pp. 140-148, 2001.

94. B. D. Noble and M. Satyanarayanan, "Experience with adaptive mobile applications in Odyssey," *Mobile Networks and Applications (MONET)* (ACM / Kluwer Academic Publishers), vol. 4, pp. 245–254, 1999.

95. J. Padhye, V. Firoiu, D. F. Towsley, and J. F. Kurose, "Modeling TCP Reno performance: A simple model and its empirical validation," *IEEE/ACM Transactions on Networking*, vol. 8, no. 2, pp. 133-145, April 2000.

96. J. Padhye, V. Firoiu, D. F. Towsley, and J. F. Kurose, "Modeling TCP throughput: A simple model and its empirical validation," *Proceedings of the ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM '98)*, Vancouver, British Columbia, Canada, pp. 303-314, August/September 1998.

97. V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, "Framework for IP performance metrics," IETF Request for Comments RFC-2330, May 1998. Online at: http://www.apps.ietf.org/rfc/rfc2330.html

98. A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*. 4th edition, McGraw-Hill, New York, NY, 2001.

99. A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: The single-node case," *IEEE/ACM Transactions on Networking*, vol. 1, no. 3, pp. 344-357, June 1993.

100. A. K. Parekh and R. G. Gallagher, "A generalized processor sharing approach to flow control in integrated services networks: The multiple node case," *IEEE/ACM Transactions on Networking*, vol. 2, no. 2, pp. 137-150, April 1994.

101. V. Paxson, "Measurements and Analysis of End-to-End Internet Dynamics," Ph.D. Thesis, University of California, Berkeley, April 1997. (a)

102. V. Paxson, "End-to-end routing behavior in the Internet," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 601-615, October 1997. (b)

103. V. Paxson, "Strategies for sound Internet measurement," *Proceedings of the ACM IMC*, October 2004.

104. V. Paxson and M. Allman, "Computing TCP's Retransmission Timer," IETF Request for Comments RFC-2988, November 2000. Online at: http://www.rfc-editor.org/rfc/rfc2988.txt

105. E. Perahia and R. Stacey, *Next Generation Wireless LANs: Throughput, Robustness, and Reliability in 802.11n*. Cambridge University Press, Cambridge, UK, 2008.

106. C. E. Perkins (Editor), *Ad Hoc Networks*. Addison-Wesley, Upper Saddle River, NJ, 2001.

107. L. L. Peterson and B. S. Davie, *Computer Networks: A Systems Approach*. 4th edition. Morgan Kaufmann Publishers (Elsevier, Inc.), San Francisco, CA, 2007.

108. R. Rajkumar, C. Lee, J. Lehoczky, and D. Siewiorek, "A resource allocation model for QoS management," *Proceedings of the IEEE Real-Time Systems Symposium*, pp.298-307, December 1997.

109. D. P. Reed, J. H. Saltzer, and D. D. Clark, "Comment on active networking and end-to-end arguments," *IEEE Network*, vol. 12, no. 3, pp. 69-71, May/June 1998.

110. Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF Request for Comments RFC-4271, January 2006. Online at: http://www.apps.ietf.org/rfc/rfc4271.html

111. J. H. Saltzer, D. P. Reed, D. D. Clark, "End-to-end arguments in system design," *ACM Transactions on Computer Systems*, vol. 2, no. 4, pp. 277-288, November 1984.

112. C. U. Saraydar, N. B. Mandayam, and D. J. Goodman, "Efficient power control via pricing in wireless data networks," *IEEE Transactions on Communications*, vol. 50, no. 2, pp. 291-303, February 2002.

113. A. Sears and J. A. Jacko, "Understanding the relationship between network quality of service and the usability of distributed multimedia documents," *Human-Computer Interaction*, vol. 15, pp. 43-68, 2000.

114. S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1233-1245, December 2006.

115. C. E. Shannon and W. Weaver, *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, IL, 1949.

116. S. Shenker, "Fundamental design issues for the future Internet," *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 7, pp. 1176-1188, September 1995.

117. S. Shenker, "Making greed work in networks: A game-theoretic analysis of switch service disciplines," *Proceedings of the ACM SIGCOMM 1994 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pp. 47-57, 1994.

118. S. Shenker, C. Partridge, and R. Guerin, "Specification of Guaranteed Quality of Service," IETF Request for Comments RFC-2212, September 1997. Online at: http://www.apps.ietf.org/rfc/rfc2212.html

119. G. Shrimali and S. Kumar, "Paid peering among Internet service providers," *Proceeding from the 2006 Workshop on Game Theory for Communications and Networks (GameNets '06)*, ACM International Conference Proceeding Series, vol. 199, article no. 11, Pisa, Italy, October 2006.

120. W. Simpson, "The Point-to-Point Protocol (PPP)," IETF Request for Comments RFC-1661, July 1994. Online at: http://www.apps.ietf.org/rfc/rfc1661.html

121. R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhäuser, Boston, MA, 2004.

122. P. Srisuresh and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)," IETF Request for Comments RFC-3022, January 2001. Online at: http://www.apps.ietf.org/rfc/rfc3022.html

123. P. Srisuresh and M. Holdrege, "IP Network Address Translator (NAT) terminology and considerations," IETF Request for Comments RFC-2663, August 1999. Online at: http://www.apps.ietf.org/rfc/rfc2663.html

124. W. R. Stevens, *TCP/IP Illustrated, Volume 1: The Protocols*. Addison-Wesley Publ. Co., Reading, MA, 1994.

125. J. W. Stewart III, *BGP: Inter-Domain Routing in the Internet*. Addison-Wesley Publ. Co., Reading, MA, 1999.

126. W. R. Stevens, "TCP slow start, congestion avoidance, fast retransmit, and fast recovery algorithms," IETF Request for Comments RFC-2001, January 1997. Online at: http://www.apps.ietf.org/rfc/rfc2001.html

127. D. E. Taylor, "Survey and taxonomy of packet classification techniques," *ACM Computing Surveys*, vol. 37, no. 3, pp. 238-275, September 2005.

128. C. A. Thekkath, T. D. Nguyen, E. Moy, and E. D. Lazowska, "Implementing network protocols at user level," *Proceedings of the ACM SIGCOMM*, San Francisco, CA, September 1993.

129. K. Thomson, G. J. Miller, and R. Wilder, "Wide-area traffic patterns and characteristics," *IEEE Network*, December 1997.

130. P. Thornycroft, "Designed for Speed: Network Infrastructure in an 802.11n World," white paper, Aruba Networks, Inc. 2009. Online at: http://www.arubanetworks.com/pdf/technology/whitepapers/wp_Designed_Speed_802.11n.pdf

131. J. S. Turner, "New directions in communications (or which way to the information age?)." *IEEE Communications Magazine*, vol. 24, no. 10, pp. 8-15, 1986.

132. C.-Y. Wang and H.-Y. Wei, "IEEE 802.11n MAC enhancement and performance evaluation," *ACM Mobile Networks and Applications (MONET)*, vol. 14, no. 6, pp. 760-771, December 2009.

133. Z. Wang, *Internet QoS: Architectures and Mechanisms for Quality of Service*, Morgan Kaufmann Publishers (Academic Press), San Francisco, CA, 2001.

134. S. Weinstein, "The mobile Internet: Wireless LAN vs. 3G cellular mobile," *IEEE Communications Magazine*, pp.26-28, February 2002.

135. A. Wolisz and F. H. P. Fitzek, "QoS support in wireless networks using simultaneous MAC packet transmission (SMPT)," in ATS, April 1999.

136. Y. Xiao, "IEEE 802.11n: Enhancements for higher throughput in WLANs," *IEEE Wireless Communications Magazine*, vol. 12, no. 6, pp. 82-91, December 2005.

137. Y. Xiao and J. Rosdahl, "Performance analysis and enhancement for the current and future IEEE 802.11 MAC protocols," *ACM SIGMOBILE Mobile Computing and Communications Review (MC2R)*, special issue on Wireless Home Networks, vol. 7, no. 2, pp. 6-19, April 2003.

138. Y. Xiao and J. Rosdahl, "Throughput and delay limits of IEEE 802.11," *IEEE Communications Letters*, vol. 6, no. 8, pp. 355-357, August 2002.

139. R. D. Yates and D. J. Goodman, *Probability and Stochastic Processes: A Friendly Introduction for Electrical and Computer Engineers*. 2nd edition, John Wiley & Sons, Inc., New York, NY, 2004.

140. Q. Zhang, W. Zhu, and Y.-Q. Zhang, "Resource allocation for multimedia streaming over the Internet," *IEEE Transactions on Multimedia*, vol. 3, no. 3, pp. 339-355, September 2001.

141. M. Ziegelmann, *Constrained Shortest Paths and Related Problems: Constrained Network Optimization*, VDM Verlag Dr. Müller, Saarbrücken, Germany, 2007.

Acronyms and Abbreviations

3G — Third Generation (wireless networks)
4G — Fourth Generation (wireless networks)
ABR — Available Bit-Rate
ACK — Acknowledgement
ADDBA — Add Block Acknowledgment
AIMD — Additive Increase/Multiplicative Decrease
AODV — Ad Hoc On-Demand Distance-Vector
AQM — Active Queue Management
AP — Access Point
AF — Assumed Forwarding
API — Application Programming Interface
ARP — Address Resolution Protocol
ARQ — Automatic Repeat Request
ASCII — American Standard Code for Information Interchange
ASIC — Application Specific Integrated Circuit
ASN — Autonomous System Number
ATM — Asynchronous Transfer Mode
AWGN — Additive White Gaussian Noise
BACK — Block Acknowledgment
BAN — Body Area Network
BDPDR — Bounded Delay Packet Delivery Ratio
BER — Bit Error Rate
BGP — Border Gateway Protocol
bps — bits per second
BS — Base Station
BSS — Basic Service Set
CBR — Constant Bit-Rate
CBT — Core Based Tree
CCA — Clear Channel Assessment
CDMA — Code Division Multiple Access
CDN — Content Distribution Network
CIDR — Classless Interdomain Routing
COA — Care-Of Address
CORBA — Common Object Request Broker Architecture
CoS — Class of Service
CPU — Central Processing Unit
CQS — Classify, Queue, and Schedule
CRC — Cyclic Redundancy Check
CSMA — Carrier-Sense Multiple Access
  CSMA/CA — CSMA / Collision Avoidance
  CSMA/CD — CSMA / Collision Detection
CSPF — Constrained Shortest Path First
CTS — Clear To Send
DBS — Direct Broadcast Satellite
DCF — Distributed Coordination Function
DELBA — Delete Block Acknowledgment
DHCP — Dynamic Host Configuration Protocol
DiffServ — Differentiated Services (alternative: DS)
DIFS — DCF (or Distributed) Inter Frame Space
DNS — Domain Name System
DPI — Deep Packet Inspection
DSR — Dynamic Source Routing
DTN — Disruption-Tolerant Networking
dupACK — Duplicate Acknowledgement
DV — Distance Vector
DVMRP — Distance Vector Multicast Routing Protocol
ECN — Explicit Congestion Notification
EF — Expedited Forwarding
EGP — Exterior Gateway Protocol
EIFS — Extended Inter Frame Space
ESS — Extended Service Set
EV-DO — EVolution – Data Optimized
EWMA — Exponential Weighted Moving Average
FCFS — First Come First Served
FDM — Frequency Division Multiplexing
FDMA — Frequency Division Multiple Access
FEC — Forward Error Correction;
  also: Forwarding Equivalence Class (in MPLS)
FIB — Forwarding Information Base
FIFO — First In First Out
FIRO — First In Random Out
FPGA — Field-Programmable Gate Array
FQ — Fair Queuing
FSM — Finite State Machine
FTP — File Transfer Protocol
GBN — Go-Back-N
GPS — Generalized Processor Sharing
GUI — Graphical User Interface

HAA — Home Address Agent
HDLC — High-level Data Link Control
HOL — Head Of Line
HSPA — High Speed Packet Access
HT — High Throughput
HTML — HyperText Markup Language
HTTP — HyperText Transport Protocol
IANA — Internet Assigned Numbers Authority
IBSS — Independent Basic Service Set
ICANN — Internet Corporation for Assigned Names and Numbers
ICMP — Internet Control Message Protocol
IEEE — Institute of Electrical and Electronics Engineers
IETF — Internet Engineering Task Force
IFS — Inter Frame Spacing
IGP — Interior Gateway Protocol
IntServ — Integrated Services
IP — Internet Protocol
  IPv4 — Internet Protocol version 4
  IPv6 — Internet Protocol version 6
IPPM — IP Performance Metrics
ISO — International Standards Organization
ISP — Internet Service Provider
j.n.d. — just noticeable difference
Kbps — Kilo bits per second
LAN — Local Area Network
LCFS — Last Come First Served
LCP — Link Control Protocol
LFIB — Label Forwarding Information Base
LIB — Label Information Base
LLC — Logical Link Control
LS — Link State
LSA — Link State Advertisement
LSDB — Link State Database
L-SIG — Legacy Signal (non-high-throughput Signal field of 802.11n physical-layer frame header)
LSP — Label Switched Path (in MPLS);
  also: Link State Packet (in LS routing and OSPF)
LSR — Label Switching Router
LTE — Long Term Evolution (also known as 4G)
MAC — Medium Access Control
MANET — Mobile Ad-hoc Network
Mbps — Mega bits per second
MCS — Modulation and Coding Scheme
MIB — Management Information Base
MIMO — Multiple-Input Multiple-Output
MPDU — MAC Protocol Data Unit
MPEG — Moving Picture Experts Group
MPLS — MultiProtocol Label Switching
MSDU — MAC Service Data Unit
MSS — Maximum Segment Size
MTU — Maximum Transmission Unit
NAK — Negative Acknowledgement
NAT — Network Address Translation
NAV — Network Allocation Vector
NCP — Network Control Protocol
NDP — Neighbor Discovery Protocol
NFE — Network Front-End (Processor)
NIC — Network Interface Card
NLRI — Network Layer Reachability Information
NMS — Network Management System
OLSR — Optimized Link State Routing
OSI — Open Systems Interconnection
OSPF — Open Shortest Path First
P2P — Peer-to-Peer (some would say Pier-to-Pier ☺)
PAN — Personal Area Network
PC — Personal Computer
PCM — Pulse Code Modulation
PCO — Phased Coexistence Operation
PDA — Personal Digital Assistant
PDU — Protocol Data Unit
pdf — probability distribution function
pmf — probability mass function
PER — Packet Error Rate
PHB — Per-Hop Behavior
PHY — Physical Layer
PIFS — PCF (or Priority) Inter Frame Space
PIM — Protocol Independent Multicast
PLCP — Physical Layer Convergence Procedure
PoP — Point-of-Presence
PPDU — PLCP Protocol Data Unit
PPP — Point-to-Point Protocol
PSDU — PLCP Service Data Unit
PSTN — Public Switched Telephone Network
PtMP — Point-to-Multipoint
PtP — Point-to-Point
QoE — Quality of Experience
QoS — Quality of Service
RED — Random Early Detection
RFC — Request For Comments
RFID — Radio Frequency Identification
RIB — Routing Information Base
RIFS — Reduced Inter Frame Space
RIP — Routing Information Protocol
RMON — Remote Monitoring
RPC — Remote Procedure Call
RPF — Reverse Path Forwarding

RSSI — Receive(r) Signal Strength Index/Indication
RSVP — Resource ReSerVation Protocol
RTCP — Real-Time Control Protocol
RTO — Retransmission Time Out
RTP — Real-Time Protocol
RTS — Request To Send
RTSP — Real-Time Streaming Protocol
RTT — Round-Trip Time
SACK — Selective Acknowledgement
SDM — Spatial Division Multiplexing
SDP — Session Description Protocol
SFD — Start Frame Delimiter
SIFS — Short Inter Frame Space
SIP — Session Initiation Protocol
SLA — Service Level Agreement
SMTP — Simple Mail Transfer Protocol
SN — Sequence Number
SNMP — Simple Network Management Protocol
SNR — Signal-to-Noise Ratio
SONET — Synchronous Optical Network
SR — Selective Repeat
SSTresh — Slow-Start Threshold
STP — Spanning Tree Protocol
TC — Traffic Category
TCP — Transmission Control Protocol
TDM — Time Division Multiplexing
TDMA — Time Division Multiple Access
TE — Traffic Engineering
TID — Traffic Identifier
TS — Traffic Stream
TTL — Time To Live
TXOP — Transmit Opportunity
UBR — Unspecified Bit Rate
UDP — User Datagram Protocol
URL — Uniform Resource Locator
VBR — Variable Bit Rate
VLAN — Virtual Local Area Network
VLSI — Very Large Scale Integration
VoIP — Voice over IP
VoWiFi — Voice over Wi-Fi
VPN — Virtual Private Network
W3C — World Wide Web Consortium
WAN — Wide Area Network
WAP — Wireless Access Protocol
WEP — Wired Equivalent Privacy
WFQ — Weighted Fair Queuing
Wi-Fi — Wireless Fidelity (synonym for IEEE 802.11)
WiMAX — Worldwide Interoperability for Microwave Access (synonym for IEEE 802.16)
W-LAN — Wireless Local Area Network
WWW — World Wide Web
ZigBee — *See* http://en.wikipedia.org/wiki/ZigBee for the origins of this name

Index

Q

R

S

T

U

V

W

X

Y

Z