Introduction

Chapter 1: Computer Networks and the Internet

This chapter provides a broad overview of computer networking and the Internet. The chapter serves as an introductory chapter to the rest of this book. But it can also serve as the foundation for a short course on computer networking. We begin the chapter with an overview of the Internet and of networking protocols, introducing several key terms and concepts. We examine the "edge" of a computer network, looking at the end systems and applications, and at the transport service provided to applications running on the end systems. We also examine the "core" of a computer network, examining the links and switches that transport data. We then take a broader view of networking. From a performance standpoint, we study the causes of packet delay and loss in a computer network. We identify key architectural principles in networking, including layering and service models. We provide a brief introduction history of computer networking. Finally, we provide a brief overview of ATM, a networking technology that provides an important contrast with the Internet technologies.

1.1: What Is the Internet?

In this book we use the public Internet, a specific computer network (and one which probably most readers have used), as our principle vehicle for discussing computer networking protocols. But what is the Internet? We would like to be able to give you a one-sentence definition of the Internet, a definition that you can take home and share with your family and friends. Alas, the Internet is very complex, both in terms of its hardware and software components, as well as in the services it provides.

1.1.1: A Nuts and Bolts Description

Instead of giving a one-sentence definition, let's try a more descriptive approach. There are a couple of ways to do this. One way is to describe the nuts and bolts of the Internet, that is, the basic hardware and software components that make up the Internet. Another way is to describe the Internet in terms of a networking infrastructure that provides services to distributed applications. Let's begin with the nuts-and-bolts description, using Figure 1.1 to illustrate our discussion.

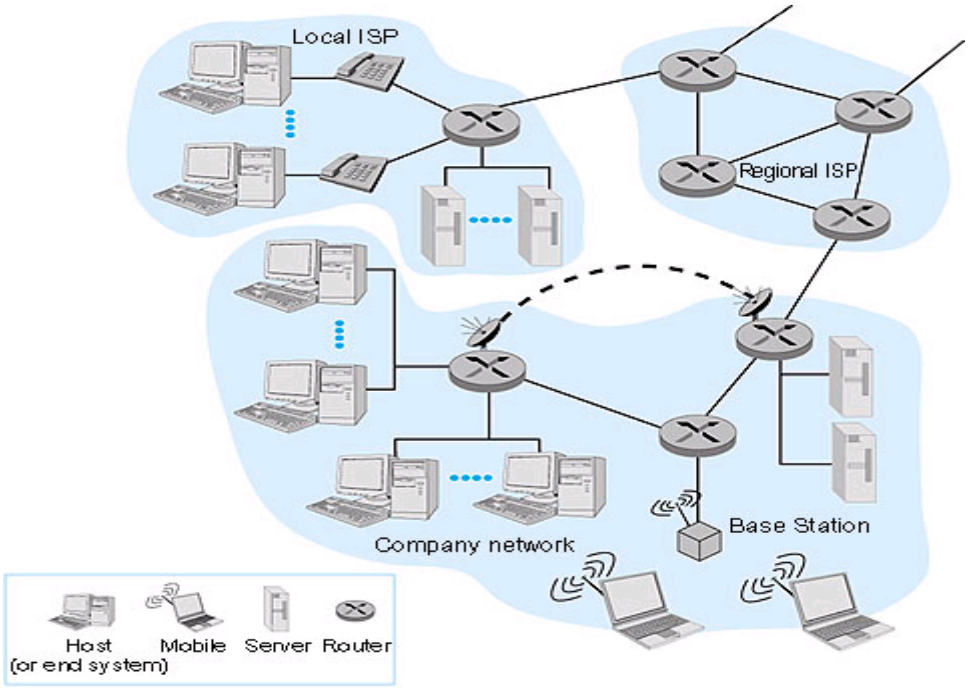

Figure 1.1: Some pieces of the Internet

The public Internet is a world-wide **computer network,** that is, a network that interconnects millions of computing devices throughout the world. Most of these computing devices are traditional desktop PCs, Unix-based workstations, and so called servers that store and transmit information such as Web (WWW) pages and e-mail messages. Increasingly, nontraditional computing devices such as Web TVs, mobile computers, pagers, and toasters are being connected to the Internet. (Toasters are not the only rather unusual devices to have been hooked up to the Internet; see the "The Future of the Living Room" [Greenberg 1997].) In the Internet jargon, all of these devices are called **hosts** or **end systems.** The Internet applications with which many of us are familiar, such as the Web and e-mail, are **network application programs** that run on such end systems. We will look into Internet end systems in more detail in Section 1.3 and then delve deeply into the study of network applications in Chapter 2.

End systems, as well as most other "pieces" of the Internet, run **protocols** that control the sending and receiving of information within the Internet. **TCP** (the Transmission Control Protocol) and **IP** (the Internet Protocol) are two of the most important protocols in the Internet. The Internet's principal protocols are collectively known as **TCP/IP.** We begin looking into protocols in Section 1.2. But that's just a start--much of this book is concerned with computer network protocols!

End systems are connected together by **communication links.** We'll see in Section 1.5 that there are many types of communication links. Links are made up of different types of **physical media,** including coaxial cable, copper wire, fiber optics, and radio spectrum. Different links can transmit data at different rates. The link transmission rate is often called the **link bandwidth** and is typically measured in bits/second.

Usually, end systems are not directly attached to each other via a single communication link. Instead, they are indirectly connected to each other through intermediate switching devices known as **routers.** A router takes information arriving on one of its incoming communication links and then forwards that information on one of its outgoing communication links. The **IP protocol** specifies the format of the information that is sent and received among routers and end systems. The path that transmitted information takes from the sending end system, through a series of communications links and routers, to the receiving end system is known as a **route** or **path** through the network. We introduce routing in more detail in Section 1.4 and study the algorithms used to determine routes, as well as the internal structure of a router itself, in Chapter 4.

Rather than provide a *dedicated* path between communicating end systems, the Internet uses a technique known as **packet switching** that allows multiple communicating end systems to share a path, or parts of a path, at the same time. The earliest ancestors of the Internet were the first packet-switched networks.

The Internet is really a network of networks. That is, the Internet is an interconnected set of privately and publicly owned and managed networks. Any network connected to the Internet must run the IP protocol and conform to certain naming and addressing conventions. Other than these few constraints, however, a network operator can configure and run its network (that is, its little piece of the Internet) however it chooses. Because of the universal use of the IP protocol in the Internet, the IP protocol is sometimes referred to as the **Internet dial tone.**

The topology of the Internet, that is, the structure of the interconnection among the various pieces of the Internet, is loosely hierarchical. Roughly speaking, from bottom-to-top, the hierarchy consists of end systems connected to local **Internet service providers (ISPs)** through **access networks.** An access network may be a so-called local area network within a company or university, a dial telephone line with a modem, or a high-speed cable-based or phone-based access network. Local ISPs are in turn connected to regional ISPs, which are in turn connected to national and international ISPs. The national and international ISPs are connected together at the highest tier in the hierarchy. New tiers and branches (that is, new networks, and new networks of networks) can be added just as a new piece of Lego can be attached to an existing Lego construction. In the first half of 1996, approximately 40,000 *new* networks were added to the Internet [Network 1996]--an astounding growth rate.

At the technical and developmental level, the Internet is made possible through creation, testing, and implementation of **Internet standards.** These standards are developed by the Internet Engineering Task Force (IETF). The IETF standards documents are called **RFCs** (request for comments). RFCs started out as general request for comments (hence the name) to resolve architecture problems that faced the precursor to the Internet. RFCs, though not formally standards, have evolved to the point where they are cited as such. RFCs tend to be quite technical and detailed. They define protocols such as TCP, IP, HTTP (for the Web), and SMTP (for open-standards e-mail). There are more than 2,000 different RFCs.

The public Internet (that is, the global network of networks discussed above) is the network that one typically refers to as *the* Internet. There are also many private networks, such as certain corporate and government networks, whose hosts are not accessible from (that is, they cannot exchange messages with) hosts outside of that private network. These private networks are often referred to as **intranets,** as they often use the same Internet technology (for example, the same types of host, routers, links, protocols, and standards) as the public Internet.