Interconnections

Bridges and Routers

Radia Perlman

Chapter 1
Essential Networking Concepts

This chapter introduces some concepts that are essential to understanding the specific subfield of computer networking that includes bridges and routers. It covers the International Standard Organization's reference model, including layering and service models. It also discusses various dimensions along which network designs can differ, such as scope, scalability, robustness, and autoconfigurability. Chapter 1 also describes the typical techniques involved in providing reliable two-party communication because some of the techniques used by routers can interact with techniques used by other layers.

1.1   Layers

Understanding, designing, and building a computer network would be too difficult a task unless the problem were partitioned into smaller subtasks. This has traditionally been done by considering the problem as being divided into several layers. The idea behind layering is that each layer is responsible for providing some service to the layer above and does this by using the services of the layer below.

Each layer communicates with its *peer* layer in another node through the use of a *protocol*. This communication is accomplished through direct communication with the layer below. The communication between layer n and layer $n-1$ is known as an *interface*.

The OSI (Open Systems Interconnection) Reference Model defines seven layers. There is nothing magic about the number seven or the functionality in the layers, however. The reference model was designed before the protocols themselves, and then committees were set up to design each of the layers. Many of the layers were subsequently subdivided into further layers. The distinction between the layers is not always clear. Bridges and routers are a good example of a case in which people should rightfully be confused about which layers are which. But semantic arguments about layers are not very productive. Instead, the layering should be viewed as a useful framework for discussion, not as a bible.

1
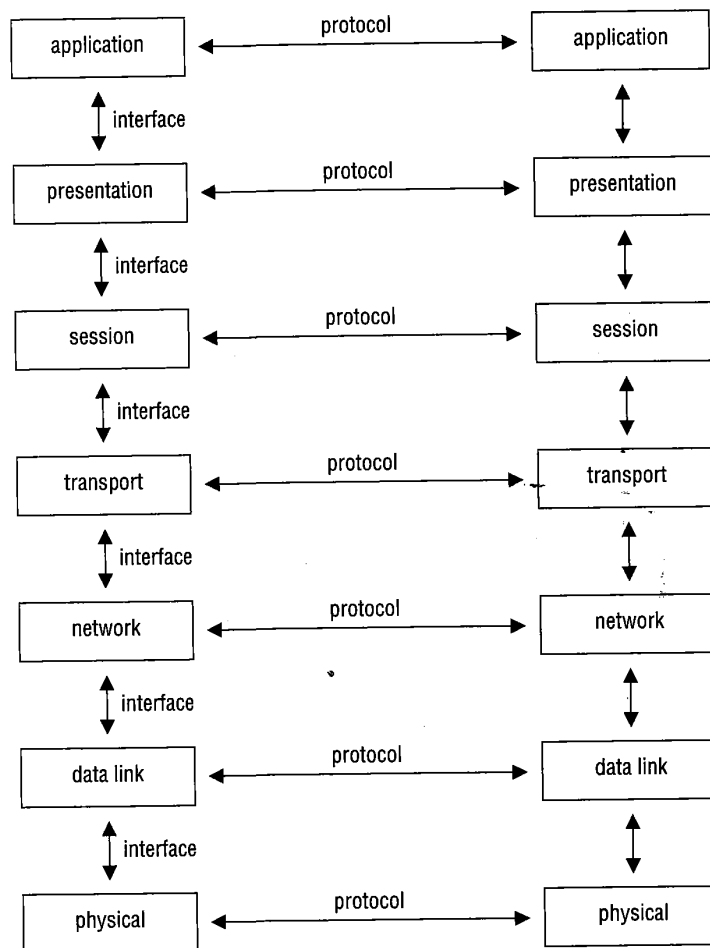
Figure 1.1

The layers defined by the ISO (International Standards Organization) are:

1. *Physical layer:* The responsibility of the physical layer is to transmit unstructured bits of information across a link. It deals with such problems as size and shape of connectors, assignment of functions to pins, conversion of bits to electrical signals, and bit-level synchronization. It is usual for several different types of physical layers to exist within a network and even for multiple different types of physical layers to exist within a node, because each technology requires its own physical layer. For instance, a node with an attachment to FDDI (fiber distributed data interface) and an attachment to a 56-KB synchronous line will have implemented two different physical layers.

2. *Data link layer:* The responsibility of the data link layer (sometimes called the link layer) is to transmit chunks of information across a link. It deals with such problems as checksumming to detect data corruption; orderly coordination of the use of shared media, as in a LAN (local area network); and addressing when multiple systems are reachable, as in a LAN. Again, it is common for different links to implement different data link layers and for a node to implement several data link layer protocols, one to support each of the different types of links to which the node is attached.

3. *Network layer:* The responsibility of the network layer is to enable any pair of systems in the network to communicate with each other. A "fully connected" network is one in which every pair of nodes has a direct link between them, but this kind of topology does not scale beyond a few nodes. Therefore, in the more typical case, the network layer must find a path through a series of connected nodes, and nodes along the path must forward packets in the appropriate direction. The network layer deals with such problems as route calculation, packet fragmentation and reassembly (when different links in the network have different maximum packet sizes), and congestion control.

4. *Transport layer:* The responsibility of the transport layer is to establish a reliable communication stream between a pair of systems. It deals with errors that can be introduced by the network layer, such as lost packets, duplicated packets, packet reordering, and fragmentation and reassembly (so that the user of the transport layer can deal with larger-size messages and so that less efficient network layer fragmentation and reassembly might be avoided).

5. *Session layer:* The responsibility of OSI's session layer is to offer services above the simple full-duplex reliable communication stream provided by transport, such as dialogue control (enforcing a particular pattern of communication between systems) and chaining (combining groups of packets so that either all or none in the group gets delivered).

6. *Presentation layer:* The responsibility of OSI's presentation layer is to provide a means by which OSI applications can agree on representations for data.

7. *Application layer:* Many OSI applications are currently, or are soon to become, standard, such as FTAM (file transfer, access, and management services) and VT (virtual terminal services). It is common for multiple applications to be running concurrently in a node.

In this book, the data link layer is relevant because bridges operate within the data link layer and because the service provided by the data link layer is relevant to routers, which operate at the network layer, thereby making the network layer obviously relevant as well. The transport layer is somewhat relevant because it is a user of the network layer and certain decisions that the network layer might make (such as whether to allow traffic to be split among several equivalent paths) affect the transport layer. The layers above transport are pretty much irrelevant to the study of bridges and routers.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.