

# *Models of Distributed Computation*

## CHAPTER 15

A distributed system is a collection of processors that do not share memory or a clock. Instead, each processor has its own local memory. The processors communicate with one another through various communication networks, such as high-speed buses or telephone lines. In this chapter, we discuss the general structure of distributed systems and the networks that interconnect them. We contrast the main differences in operating-system design between these systems and centralized systems. In Section xxx, we go on to discuss distributed file systems. Then, in Section xxx, we describe the methods necessary for distributed operating systems to coordinate their actions.

### CHAPTER OBJECTIVES

- To provide a high-level overview of distributed systems and the networks that interconnect them.
- To discuss the general structure of distributed operating systems.
- To describe various methods for achieving mutual exclusion in a distributed system.
- To explain how atomic transactions can be implemented in a distributed system.
- To show how some of the concurrency-control schemes discussed in Chapter xxx can be modified for use in a distributed environment.
- To present schemes for handling deadlock prevention, deadlock avoidance, and deadlock detection in a distributed system.

**15.1 Motivation**

**15.2 Network Structure**

**15.3 Communication Protocols**

**15.4 Robustness**

**15.5 Mutual Exclusion**

**15.6 Atomicity**

**15.7 Reaching Agreement**

**15.8 Summary**

**Exercises**

15.1 xx

**Bibliographical Notes**

Tanenbaum [2003], Stallings [2000a], and Kurose and Ross [2005] provided general overviews of computer networks. Williams [2001] covered computer networking from a computer-architecture viewpoint.