

Hardware Performance Factors

by David G. Messerschmitt

Supplementary section for Understanding Networked Applications: A First Course, Morgan Kaufmann, 1999.

Copyright notice: Permission is granted to copy and distribute this material for educational purposes only, provided that this copyright notice remains attached.

Performance and quality metrics are affected by the underlying technologies of computing and communications, by the equipment and software infrastructure, and by the application architecture.

Electronics and Photonics

The hardware technologies underlying networked computing have advanced remarkably over the past decades, including processing, memory, storage, and communications.

One major driver for both computing and networking is semiconductor electronics. As the size of devices (transistors, interconnect, etc.) on integrated circuits shrink, everything improves: Chips can have more devices, speed increases, and power consumption and cost per unit function decrease. According to Moore's law, the cost for a unit functionality has halved approximately every 18 months, or decreased by a factor of 100 every decade as described in "Semiconductor Manufacturers: Moore's Law" on page 215. Although this can't continue forever, it is expected to persist for at least another decade, and likely longer.

The most direct visible impact of electronics advances is the rapid improvement in processing speeds, the size and speed of semiconductor memory, and the bitrate supported by network switches. One direct measure of processor performance is the *instruction rate*, which aids programs in executing more quickly. Putting more functionality in a single integrated circuit also improves performance, because communicating off chip is slower than on chip.

Another remarkable technology is the fiber optics used in backbone networks and LANs. A fiber is a strand of glass that confines and propagates light over great distances with little loss in intensity. The bits comprising data in a packet are carried through fiber communication links by pulses of light. The performance measure of a communication link is the rate at which bits can be communicated over large distances, called the *bitrate*. This bitrate has improved even more rapidly than Moore's law, and the theoretical capacity of a single fiber is unlimited for practical purposes. Fiber optics is discussed further in Chapter 20.

The other major technology in computing is magnetic storage. In this case, technological advance is reflected by a larger number of stored bits per unit area of storage, correspondingly greater total storage capacity, and the speed with which data can be stored and accessed. Storage density is increasing by about a factor of ten every decade (this is known as "Hoagland's law").

It might seem that every underlying technology is improving, and this is almost the case. One performance characteristic *not* improving is the speed of light. This speed is increasingly a limitation to what can be accomplished in networked computing, and will become increasingly important in the future. This issue is discussed in Chapter 20.

Host Architecture

The instruction rates of processors, which are increasing rapidly in consonance with Moore's law, improve performance. However, other factors such as how processors, memory, storage, etc. are architected, also strongly impact performance. For example, if the communication of data to and from memory is slow, a processor may waste many idle instruction cycles waiting for data to be stored or retrieved. To improve performance, memory and storage are organized in a hierarchy, with the most frequently accessed data stored in a smaller more rapidly accessed memory (see the sidebar "Host Memory and Storage Hierarchy").

One way to improve host performance, particularly servers, is to incorporate more than one processor sharing a common memory and storage system and managed as a unit by the operating system, as shown in Figure 2.. This is called a *shared-memory multiprocessor*. The processors can run processes in parallel, and communicate at high speeds through the common memory. As defined in Chapter 4, a host has a single network access connection, but there is no reason that a host must have a single processor. However, the speed of light communication delays coupled to faster processors make this architecture decreasingly attractive with time [Gra93].

An alternative is a processor *cluster*—also illustrated in Figure 2.—in which a collection of pro-

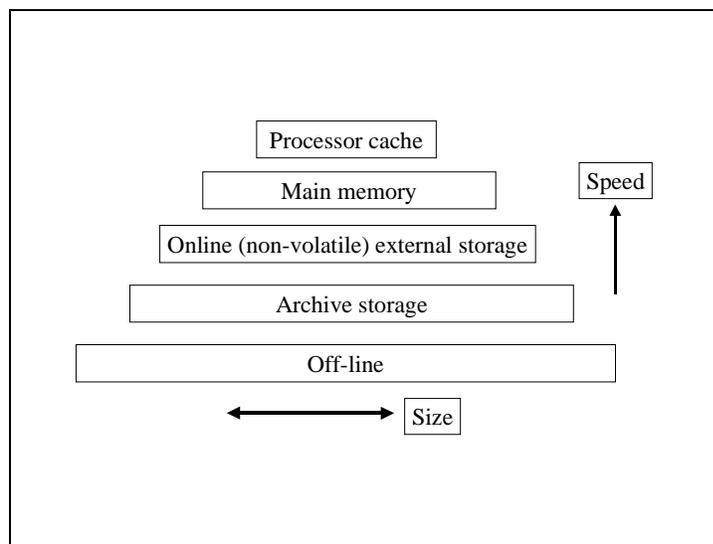


Figure 1. The hierarchy of memory and storage in a processor.

Host Memory and Storage Hierarchy

Figure 1. illustrates the hierarchical decomposition of storage and memory in a host [Gra93]. Moving downward, each level of the memory and storage has much greater capacity, but is also roughly 100 times slower to access. Mass storage media based on magnetics or optics have mechanical components (spinning disks), implying an advantage to retrieving large contiguous blocks of data relative to small amounts of widely dispersed data. The data archiving storage is based on robotic tape handlers, and is considerably slower to access than disks. Archive storage must be retrieved from a human operator, often from a warehouse at a distance. The idea of caching is discussed in Chapter 20.

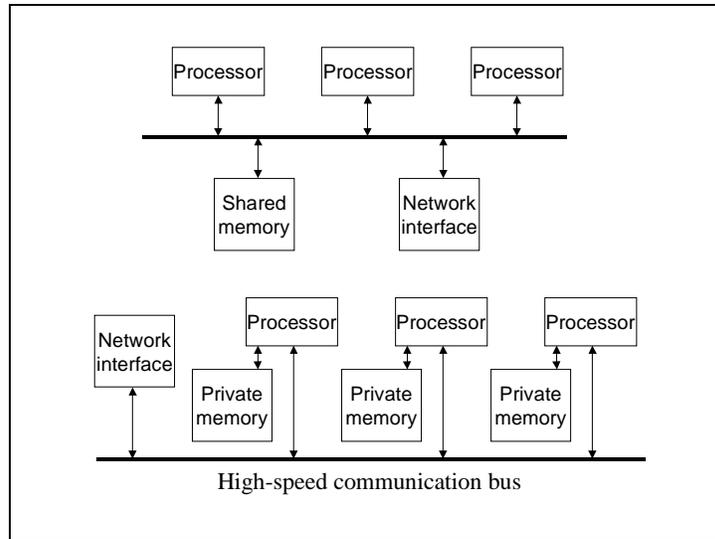


Figure 2. Simplified multiprocessor shared memory and processor cluster architec-

processors each have a dedicated private memory, but are connected with high-bitrate parallel links. The processors in a cluster can communicate more rapidly than over a network, but utilize similar communication services.

Discussion

- D1 What are some other historical examples of rapidly advancing technologies? Can you think of any other technologies that have come as far and as fast as electronics and fiber optics?
- D2 Assume that electronics, fiber optics, and storage continue to improve dramatically in performance characteristics over the next one to two decades. Discuss the implications of this to computing applications.

Review

Generally the performance of a networked computing system improves with time with the geometrical improvement in processor speeds and fiber optics bitrates, with the notable exception of the speed of light delay of a communication link. The performance of a host can be improved by using multiple processors in a common host, with either shared memory or a cluster.

Concepts

Hardware

- Host, multiprocessor, processor clusters
- Moore's law
- Fiber optics, bitrate, speed of light