8

MULTIPLE PROCESSOR SYSTEMS

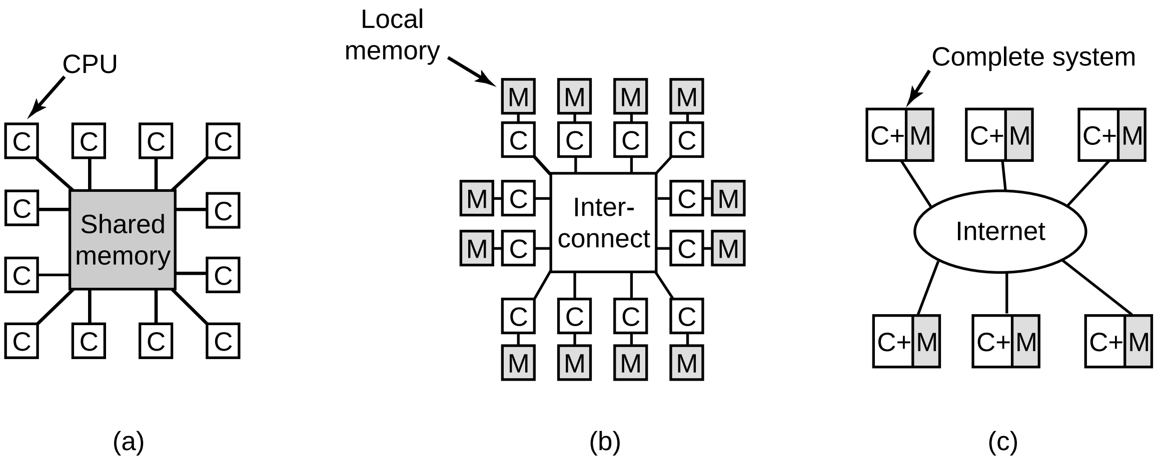

Fig. 8-1. (a) A shared-memory multiprocessor. (b) A message-passing multicomputer. (c) A wide area distributed system.

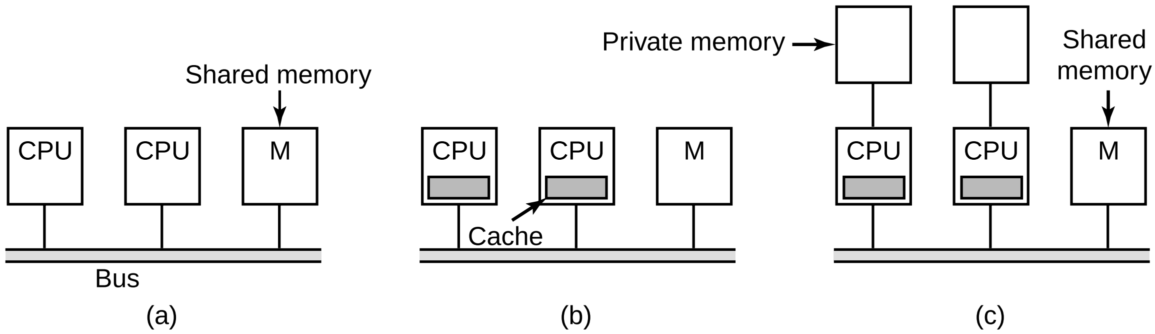

Fig. 8-2. Three bus-based multiprocessors. (a) Without caching. (b) With caching. (c) With caching and private memories.

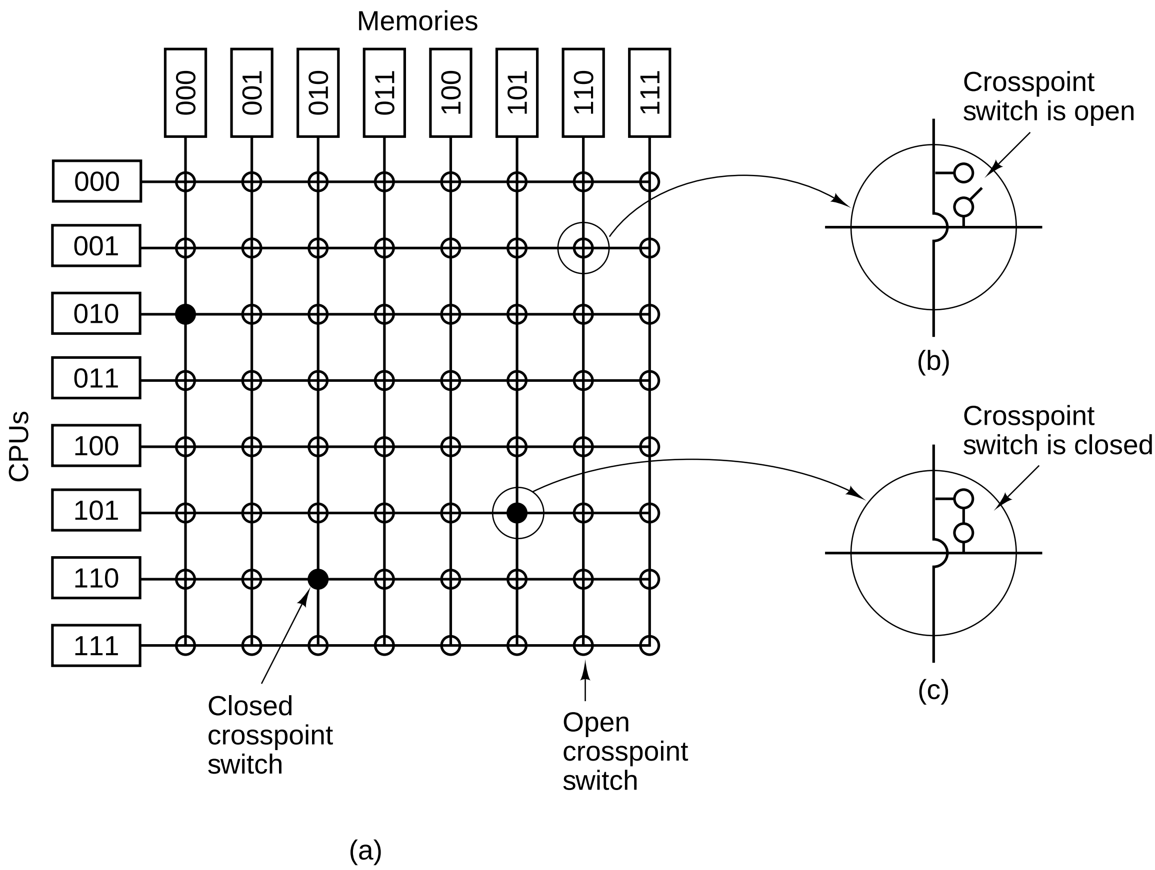

Fig. 8-3. (a) An 8×8 crossbar switch. (b) An open crosspoint. (c) A closed crosspoint.

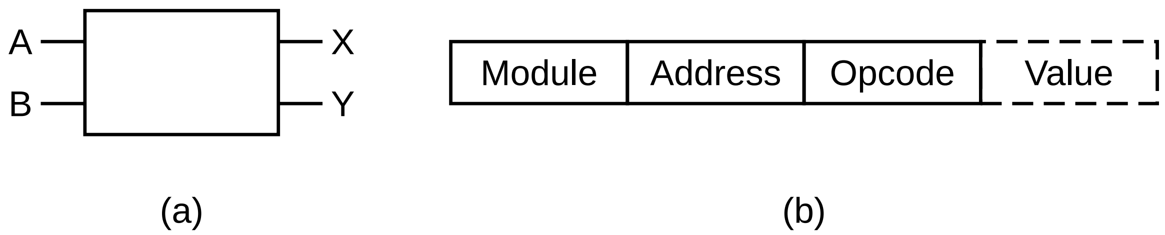

Fig. 8-4. (a) A 2×2 switch. (b) A message format.

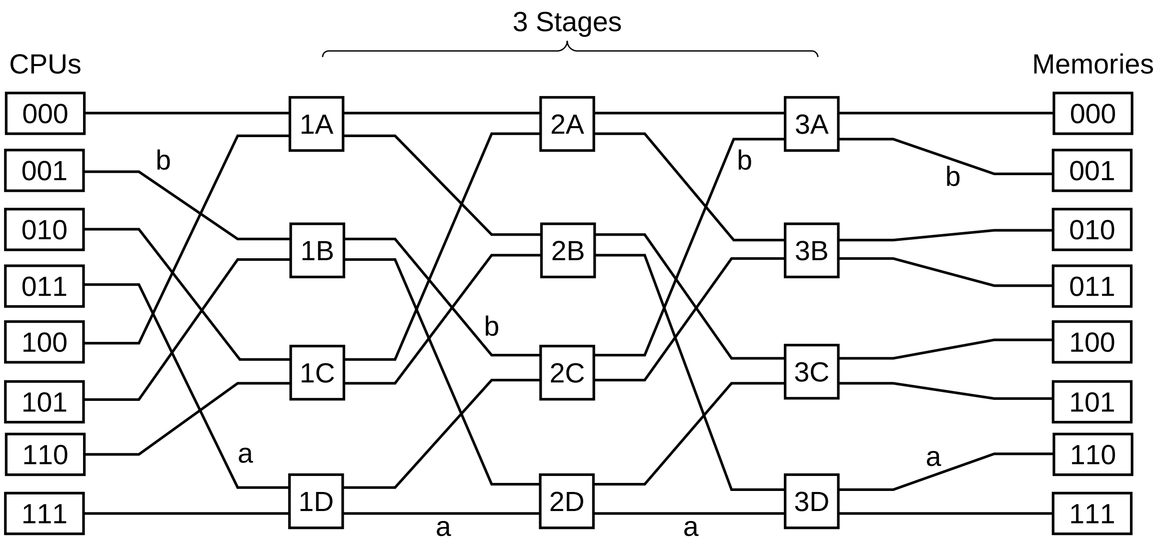

Fig. 8-5. An omega switching network.

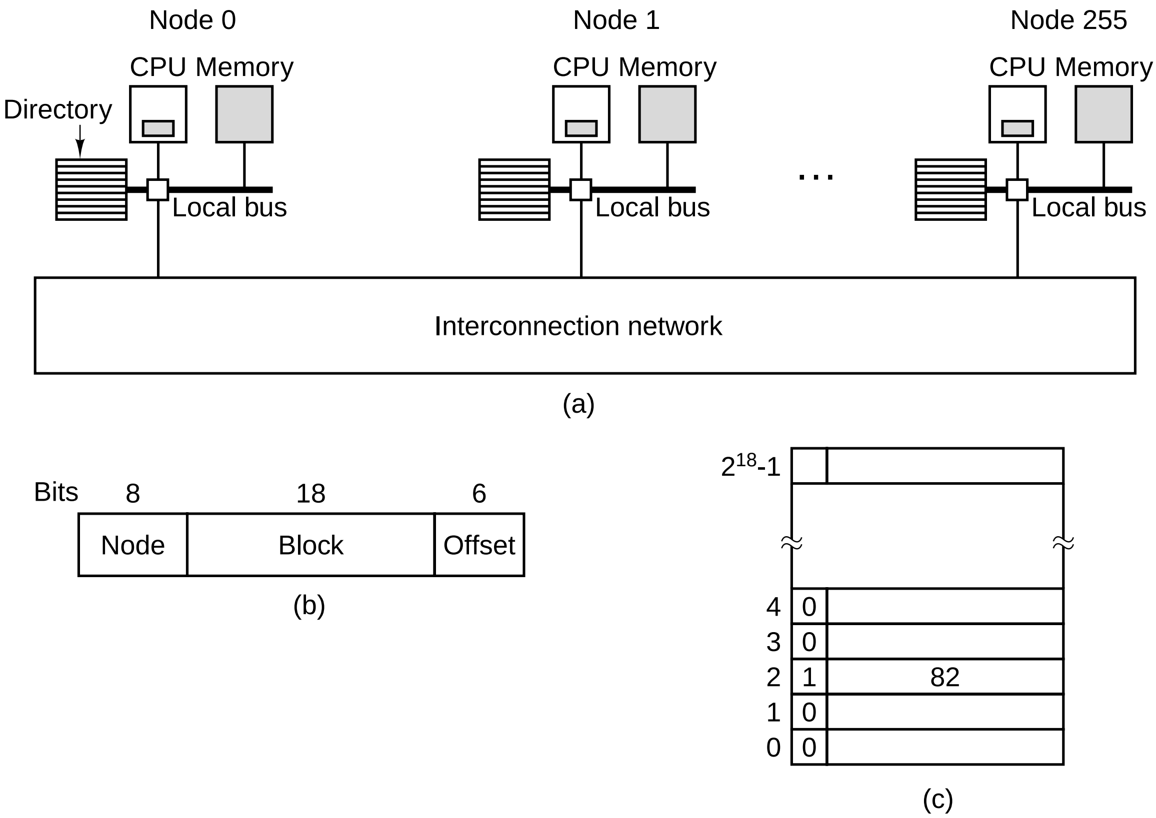

Fig. 8-6. (a) A 256-node directory-based multiprocessor. (b) Division of a 32-bit memory address into fields. (c) The directory at node 36.

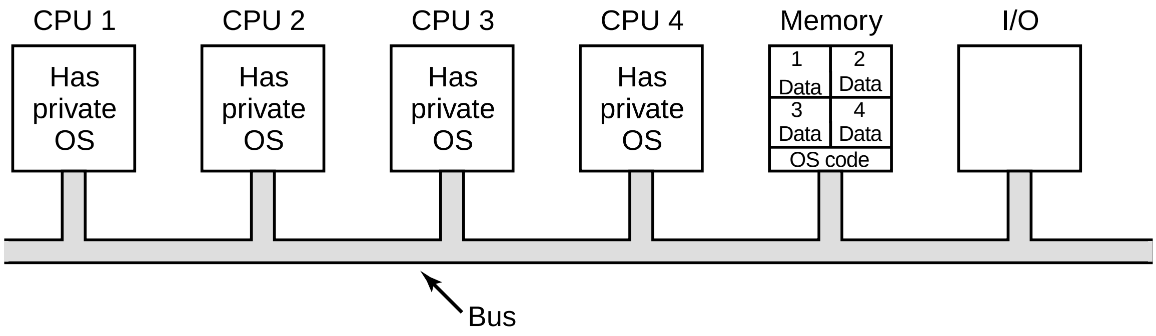

Fig. 8-7. Partitioning multiprocessor memory among four CPUs, but sharing a single copy of the operating system code. The boxes marked Data are the operating system's private data for each CPU.

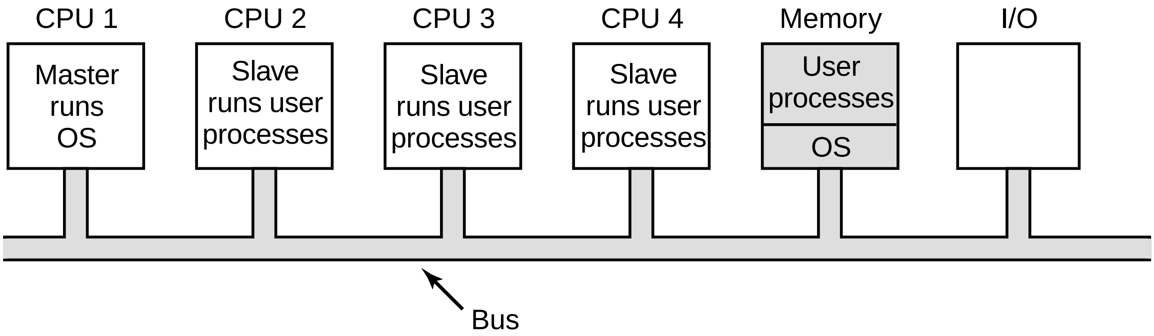

Fig. 8-8. A master-slave multiprocessor model.

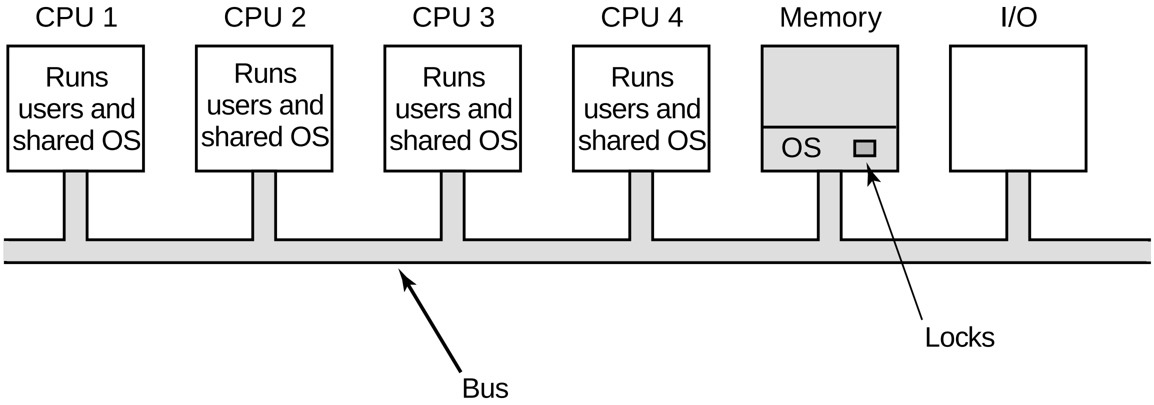

Fig. 8-9. The SMP multiprocessor model.

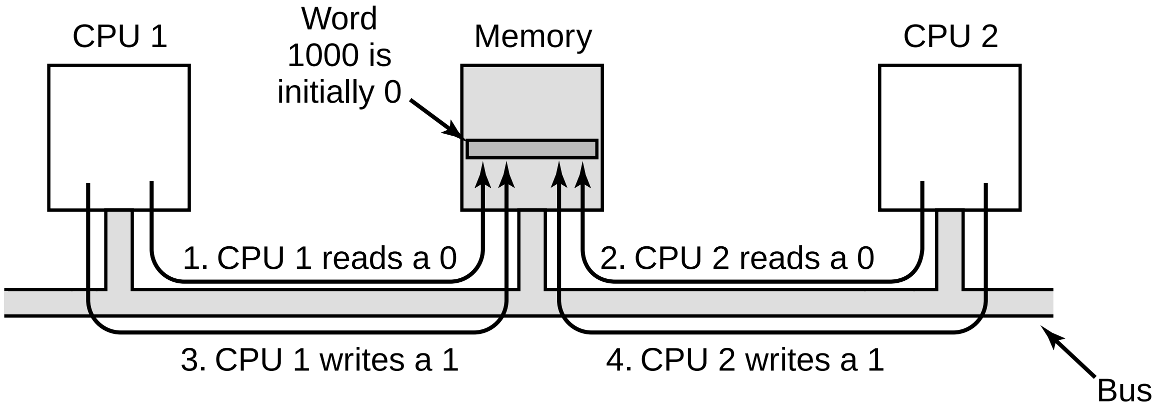

Fig. 8-10. The TSL instruction can fail if the bus cannot be locked. These four steps show a sequence of events where the failure is demonstrated.

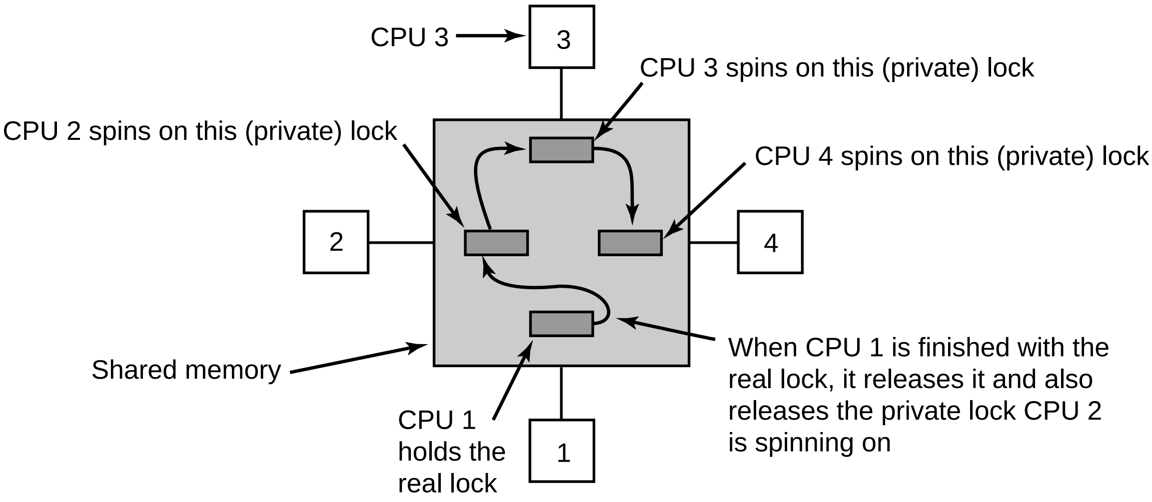

Fig. 8-11. Use of multiple locks to avoid cache thrashing.

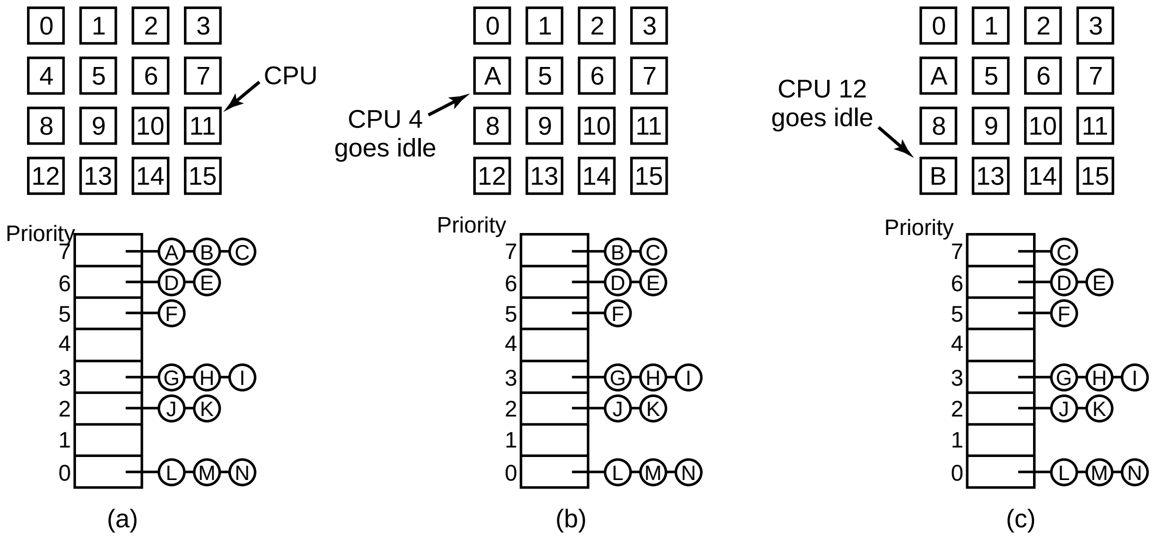

Fig. 8-12. Using a single data structure for scheduling a multiprocessor.

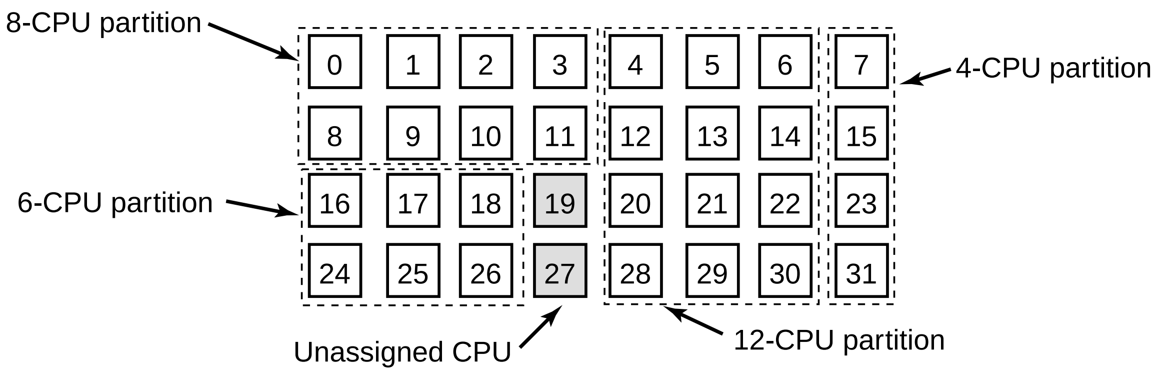

Fig. 8-13. A set of 32 CPUs split into four partitions, with two CPUs available.

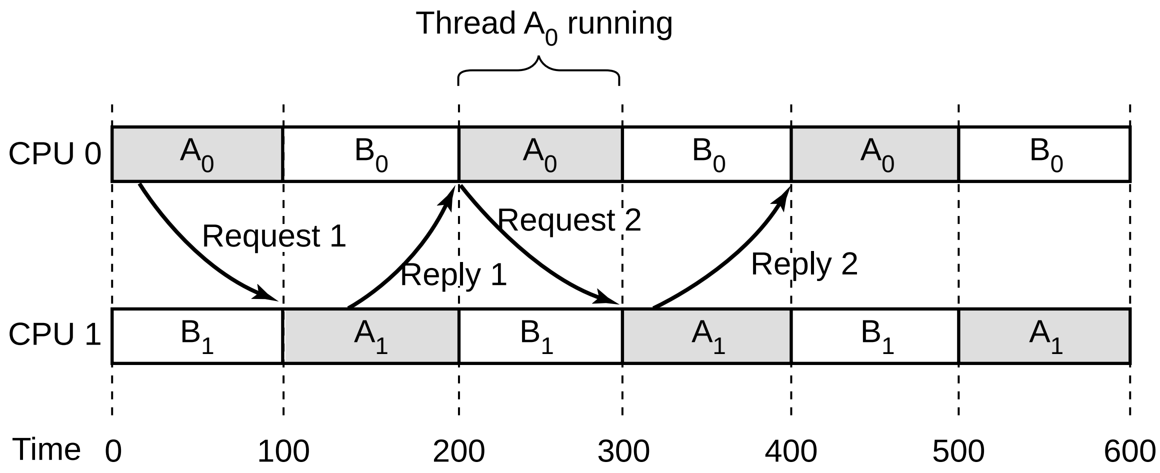

Fig. 8-14. Communication between two threads belonging to process *A* that are running out of phase.

Time slot \ CPU	0	1	2	3	4	5
0	A_0	A_1	A_2	A_3	A_4	A_5
1	B_0	B_1	B_2	C_0	C_1	C_2
2	D_0	D_1	D_2	D_3	D_4	E_0
3	E_1	E_2	E_3	E_4	E_5	E_6
4	A_0	A_1	A_2	A_3	A_4	A_5
5	B_0	B_1	B_2	C_0	C_1	C_2
6	D_0	D_1	D_2	D_3	D_4	E_0
7	E_1	E_2	E_3	E_4	E_5	E_6

Fig. 8-15. Gang scheduling.

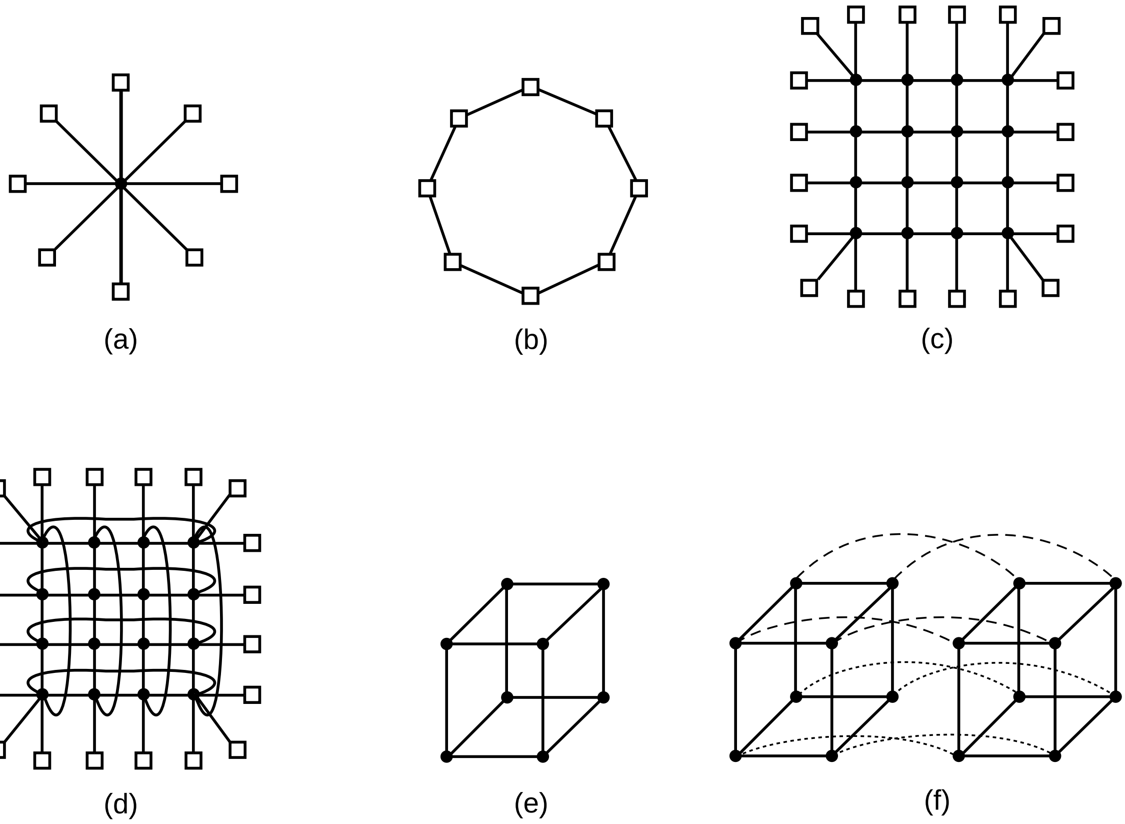

Fig. 8-16. Various interconnect topologies. (a) A single switch. (b) A ring. (c) A grid. (d) A double torus. (e) A cube. (f) A 4D hypercube.

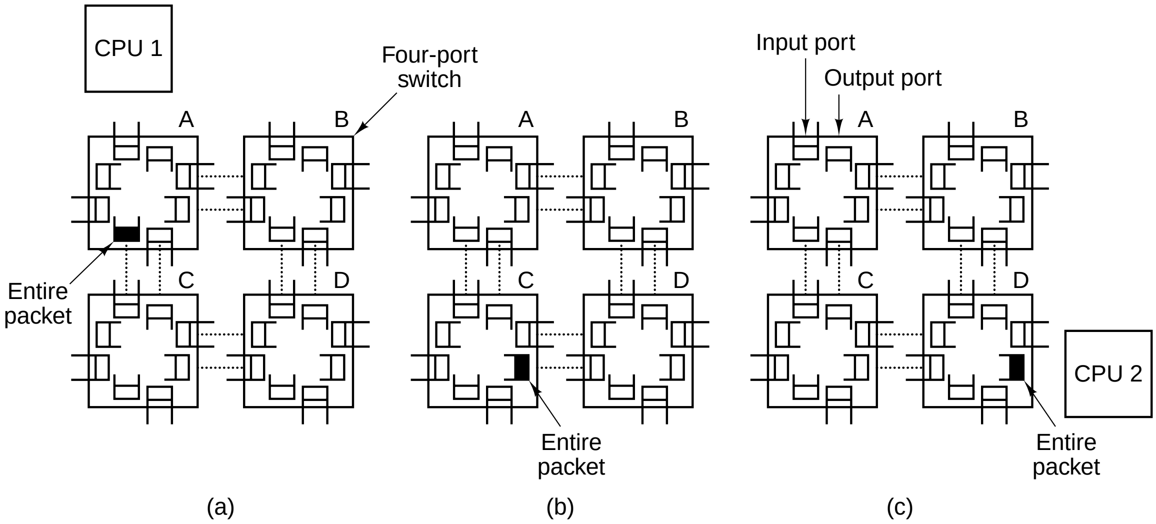

Fig. 8-17. Store-and-forward packet switching.

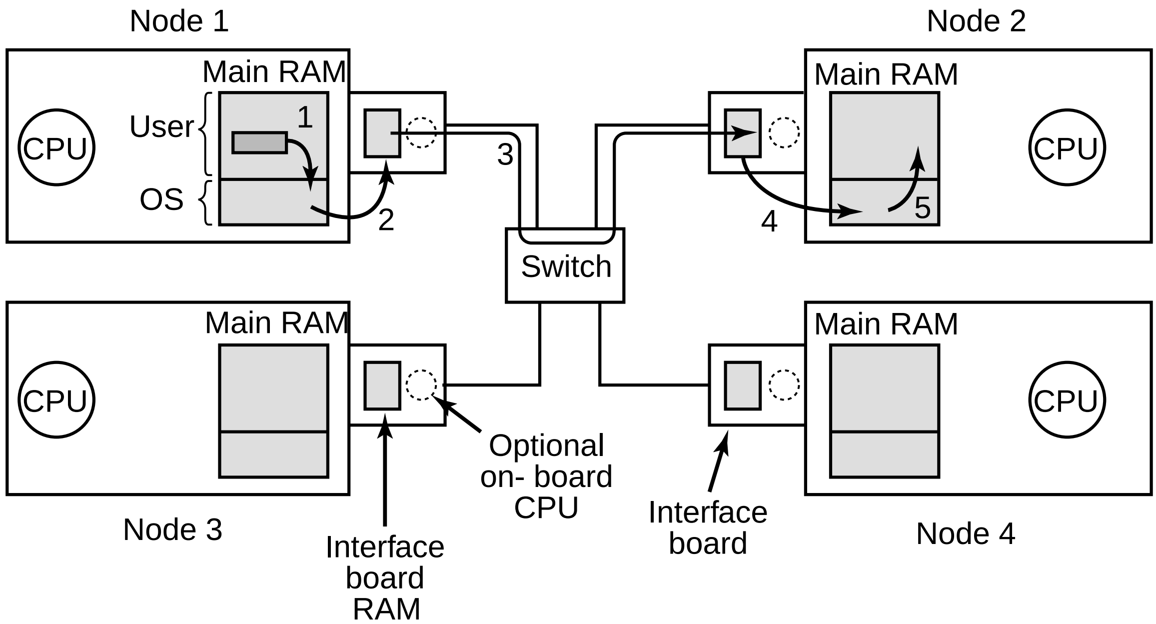

Fig. 8-18. Position of the network interface boards in a multicomputer.

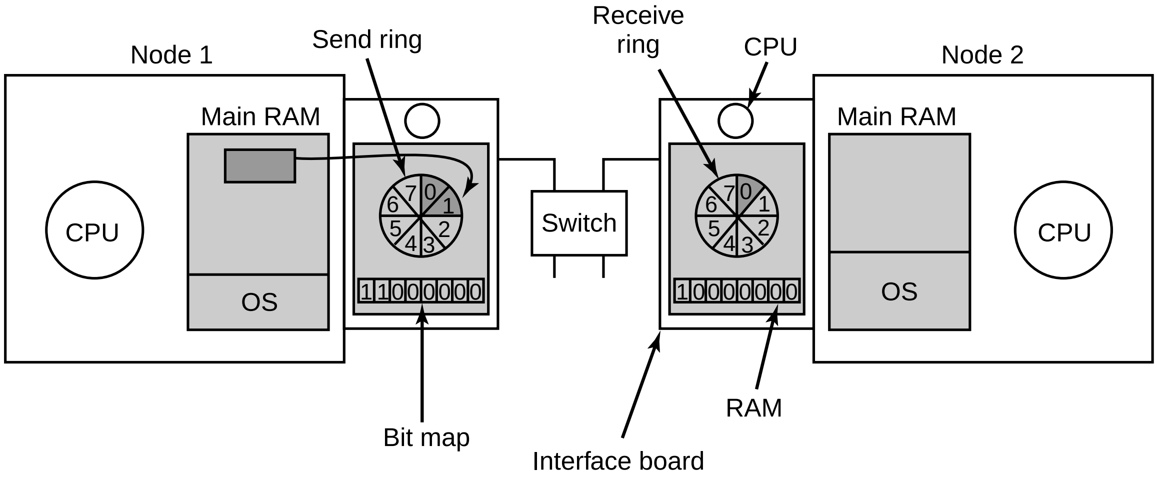

Fig. 8-19. Use of send and receive rings to coordinate the main CPU with the on-board CPU.

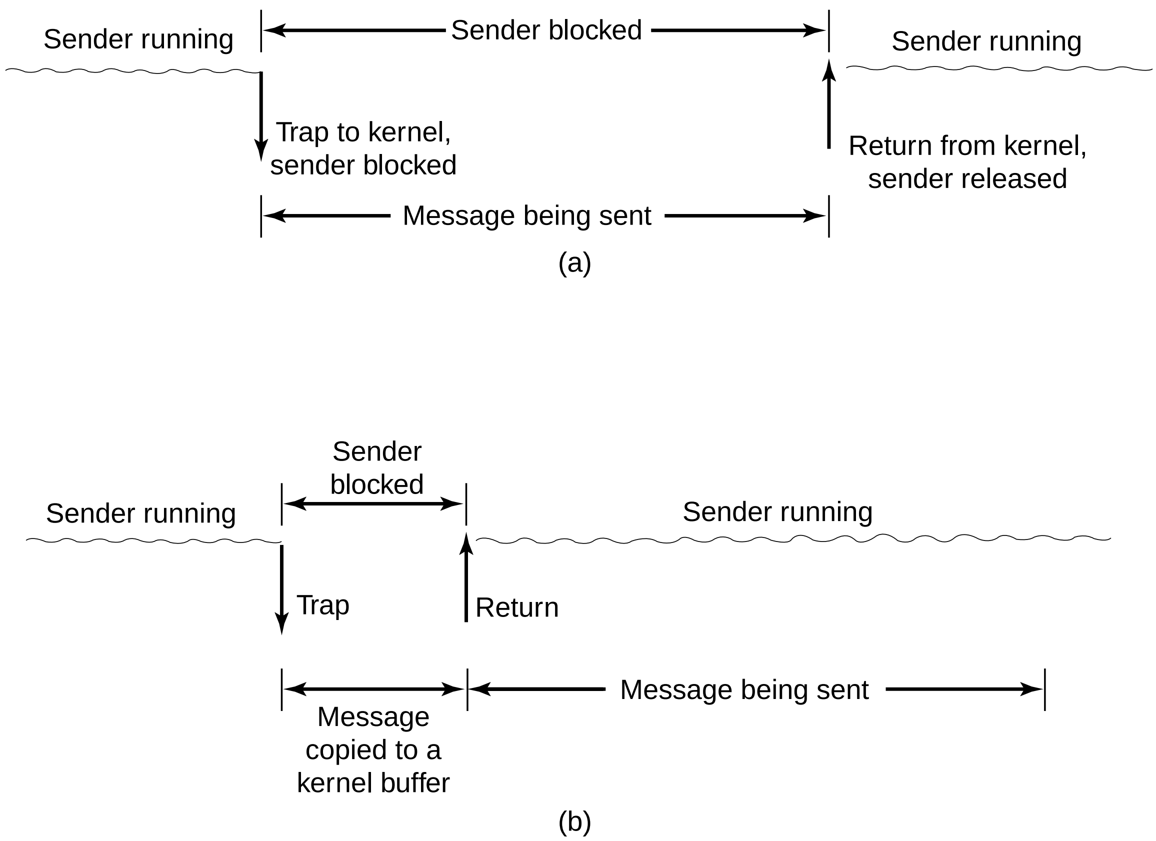

Fig. 8-20. (a) A blocking send call. (b) A nonblocking send call.

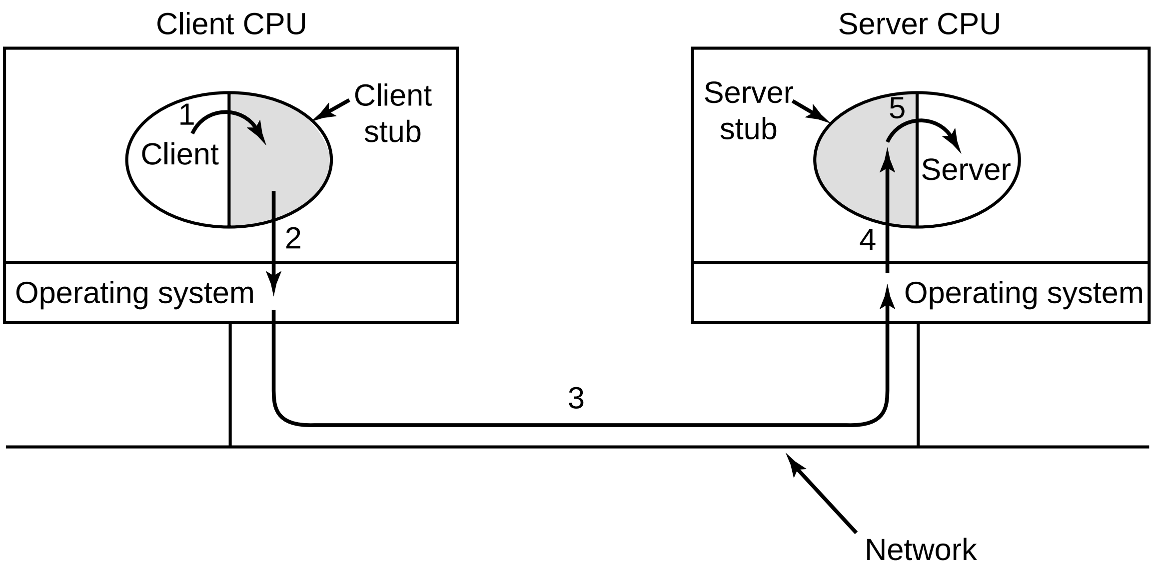

Fig. 8-21. Steps in making a remote procedure call. The stubs are shaded gray.

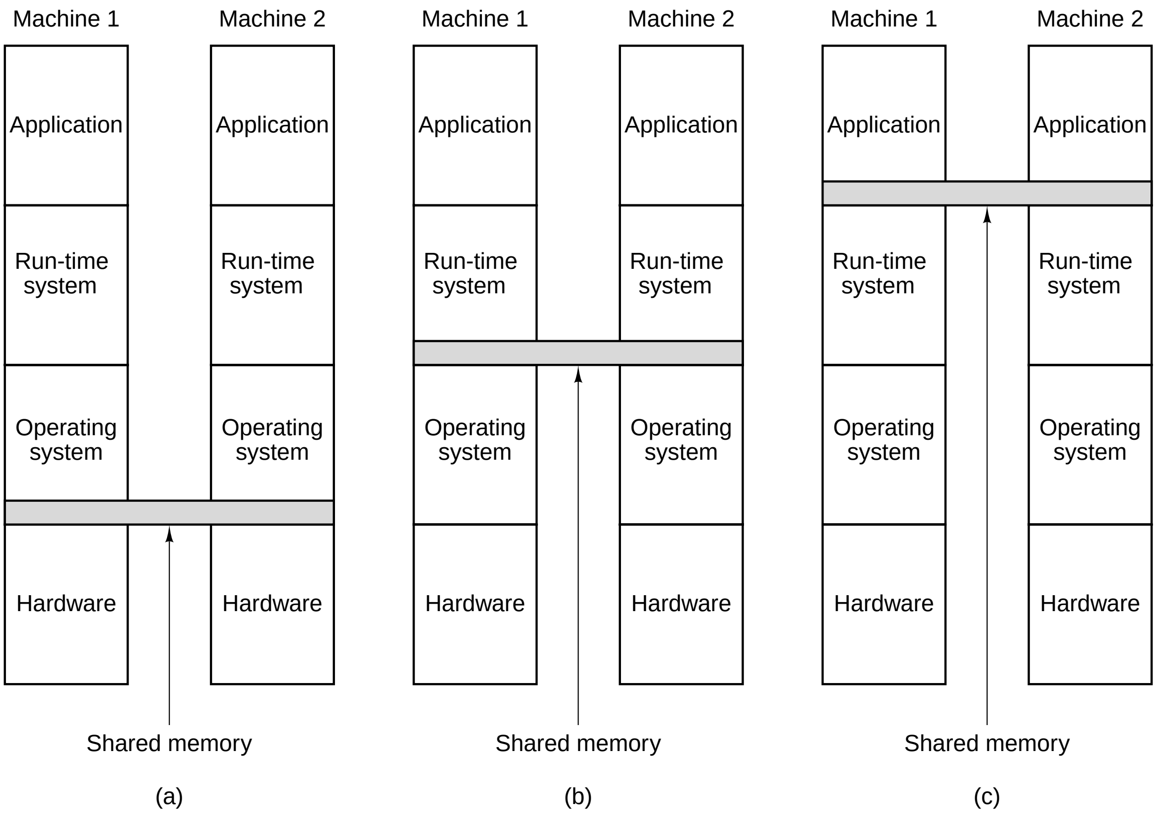

Fig. 8-22. Various layers where shared memory can be implemented. (a) The hardware. (b) The operating system. (c) User-level software.

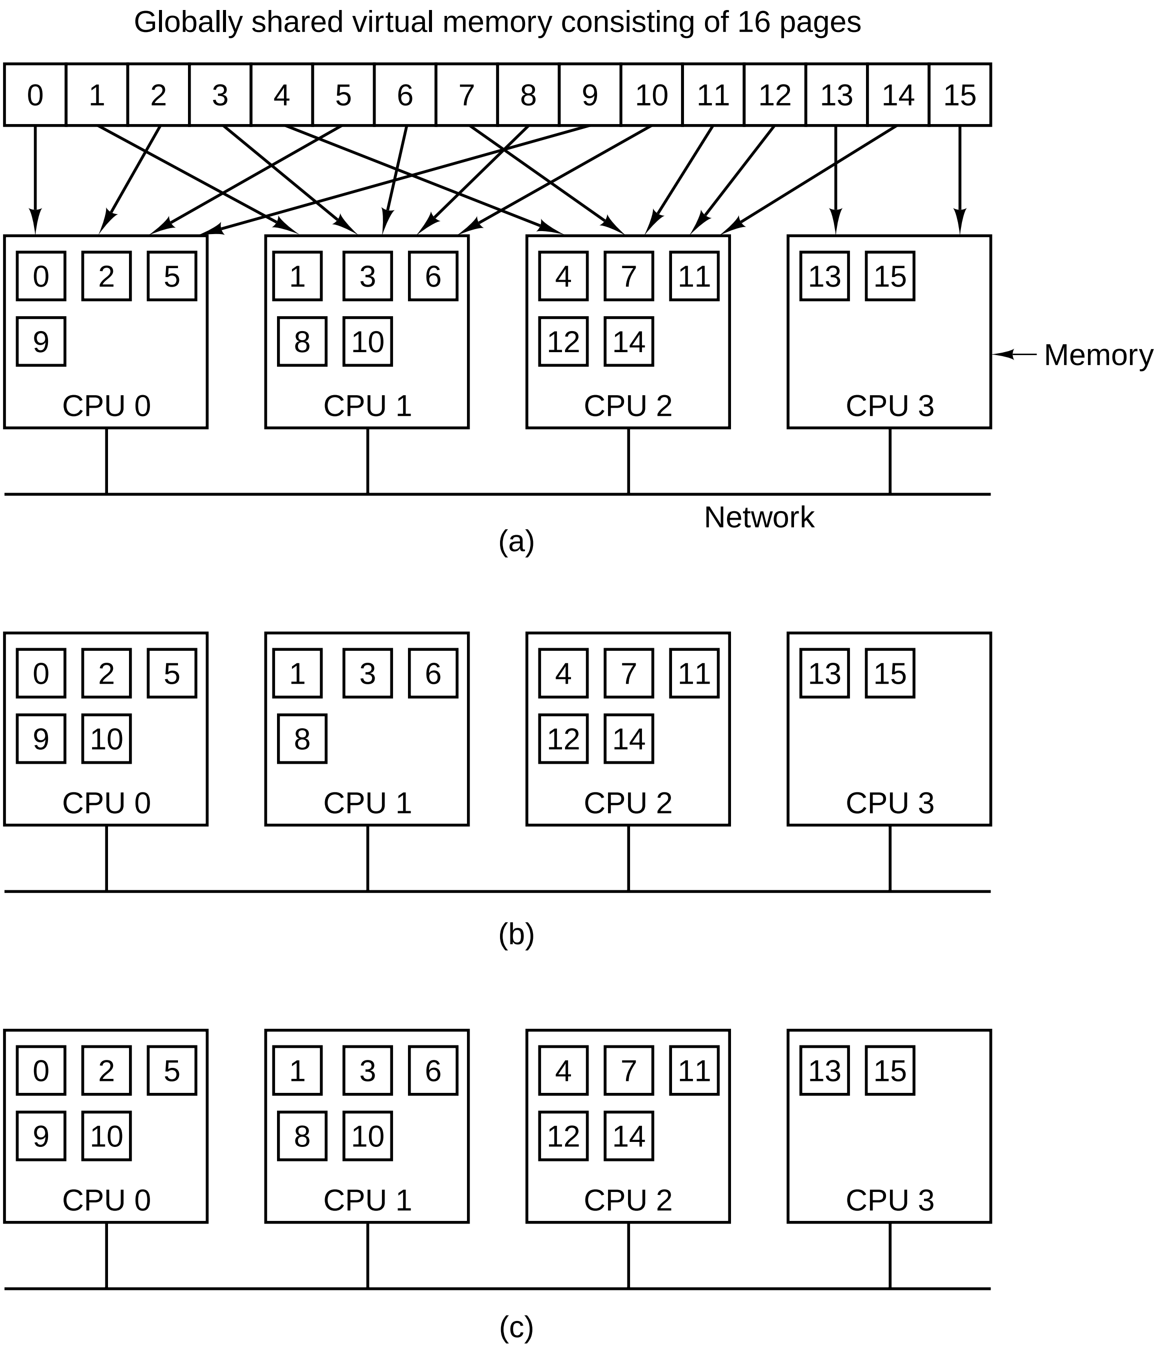

Fig. 8-23. (a) Pages of the address space distributed among four machines. (b) Situation after CPU 1 references page 10. (c) Situation if page 10 is read only and replication is used.

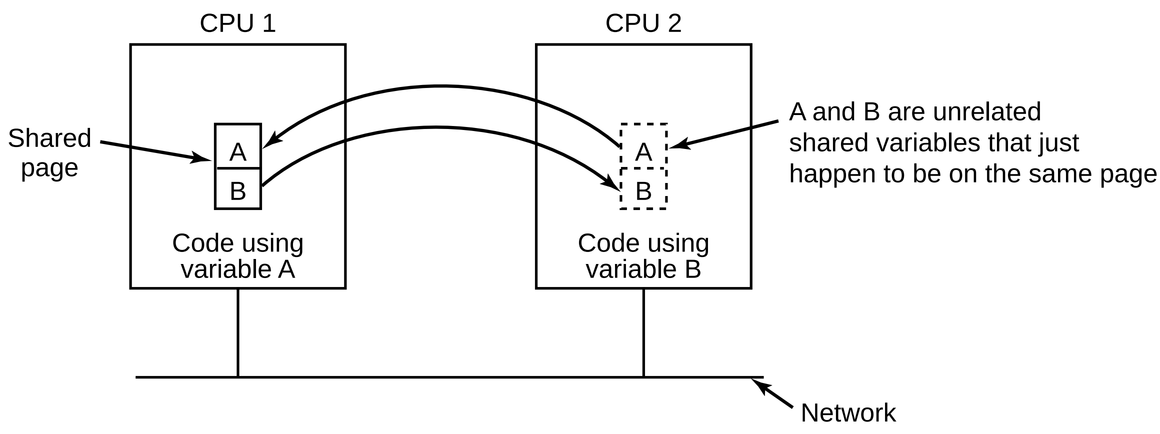

Fig. 8-24. False sharing of a page containing two unrelated variables.

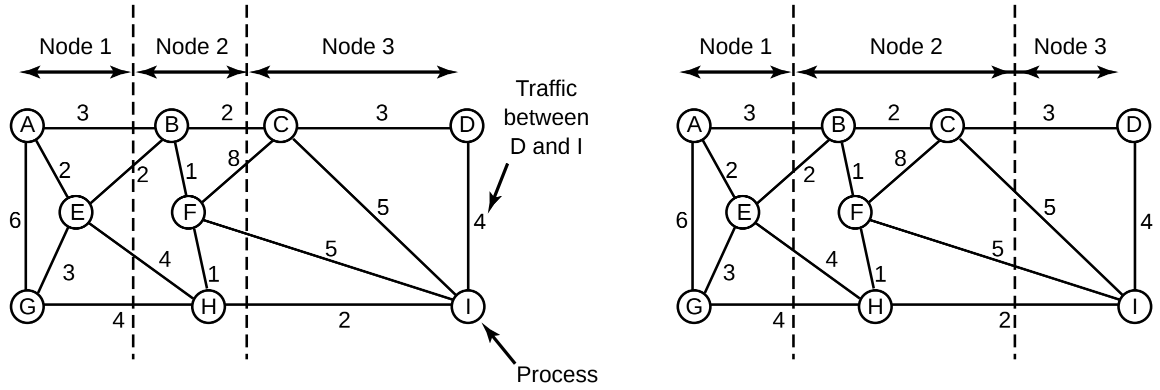

Fig. 8-25. Two ways of allocating nine processes to three nodes.

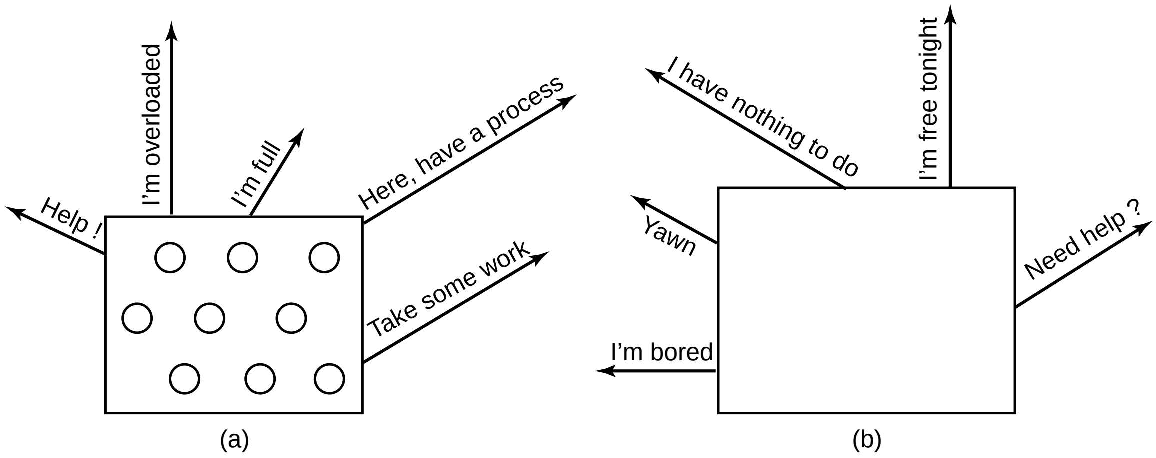

Fig. 8-26. (a) An overloaded node looking for a lightly loaded node to hand off processes to. (b) An empty node looking for work to do.

Item	Multiprocessor	Multicomputer	Distributed System
Node configuration	CPU	CPU, RAM, net interface	Complete computer
Node peripherals	All shared	Shared exc. maybe disk	Full set per node
Location	Same rack	Same room	Possibly worldwide
Internode communication	Shared RAM	Dedicated interconnect	Traditional network
Operating systems	One, shared	Multiple, same	Possibly all different
File systems	One, shared	One, shared	Each node has own
Administration	One organization	One organization	Many organizations

Fig. 8-27. Comparison of three kinds of multiple CPU systems.

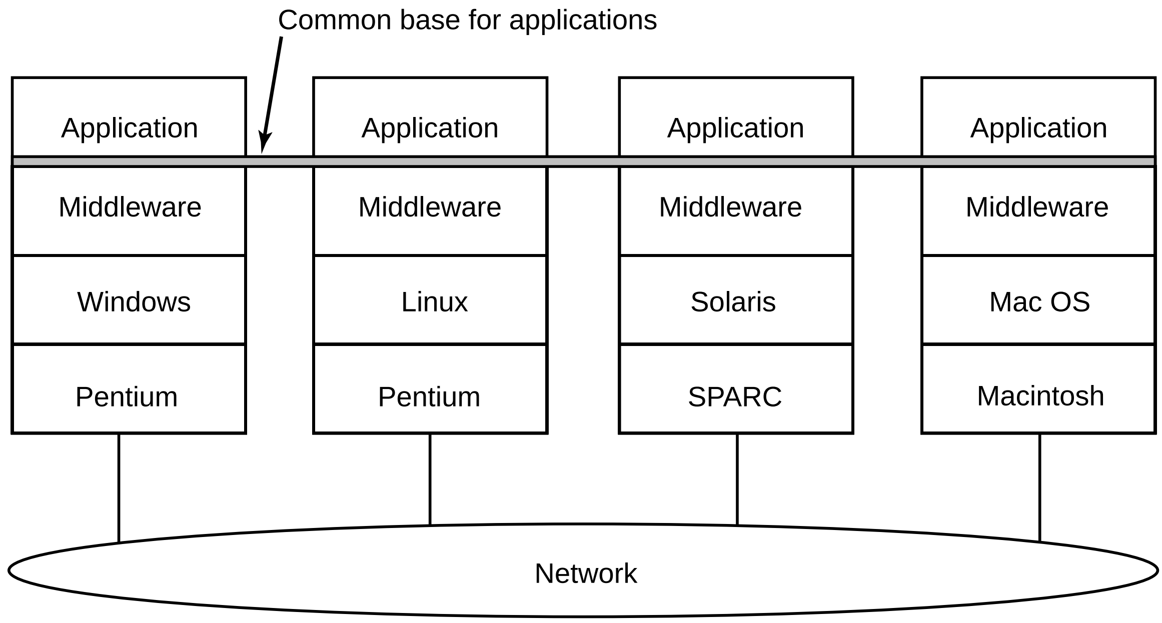

Fig. 8-28. Positioning of middleware in a distributed system.

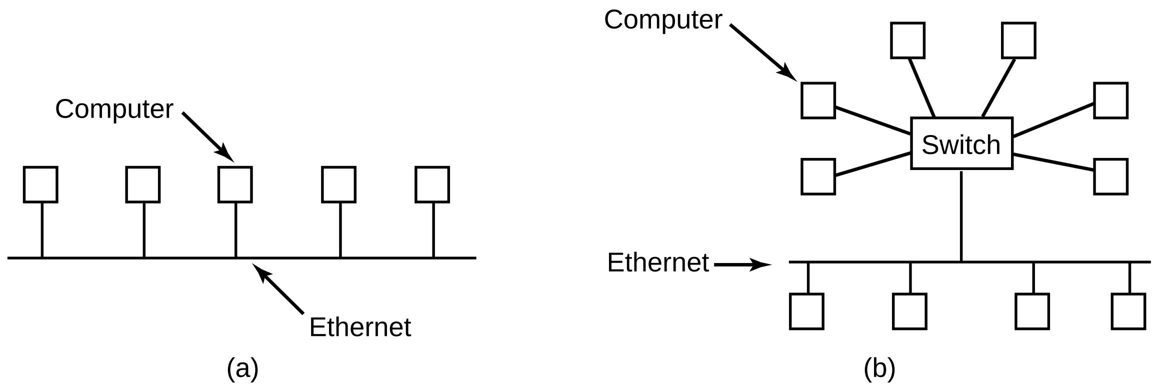

Fig. 8-29. (a) Classic Ethernet. (b) Switched Ethernet.

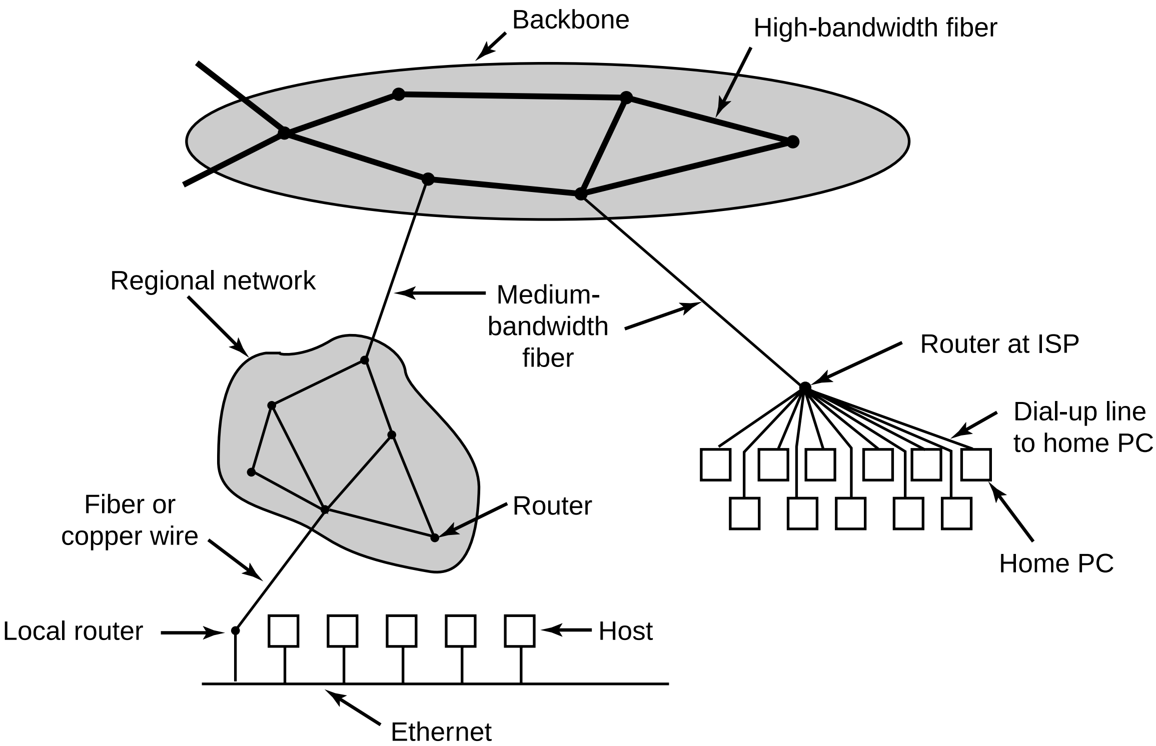

Fig. 8-30. A portion of the Internet.

	Service	Example
Connection-oriented	Reliable message stream	Sequence of pages of a book
	Reliable byte stream	Remote login
	Unreliable connection	Digitized voice
Connectionless	Unreliable datagram	Network test packets
	Acknowledged datagram	Registered mail
	Request-reply	Database query

Fig. 8-31. Six different types of network service.

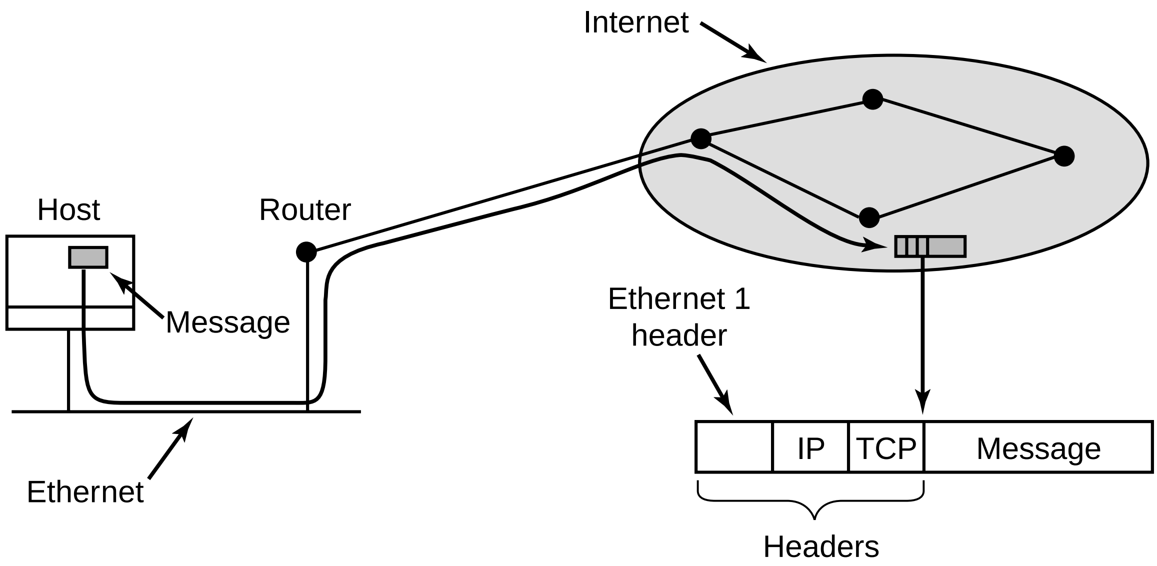

Fig. 8-32. Accumulation of packet headers.

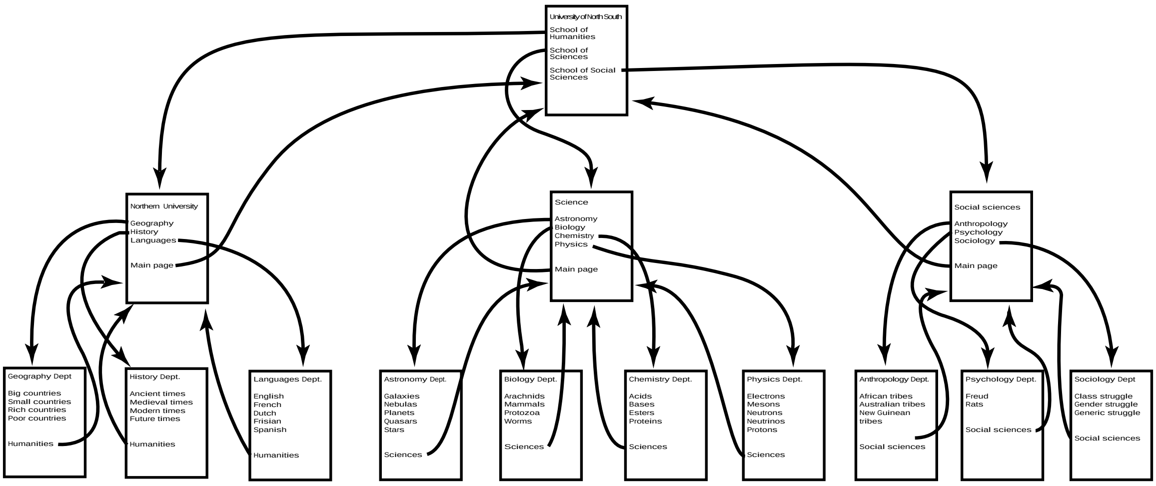

Fig. 8-33. The Web is a big directed graph of documents.

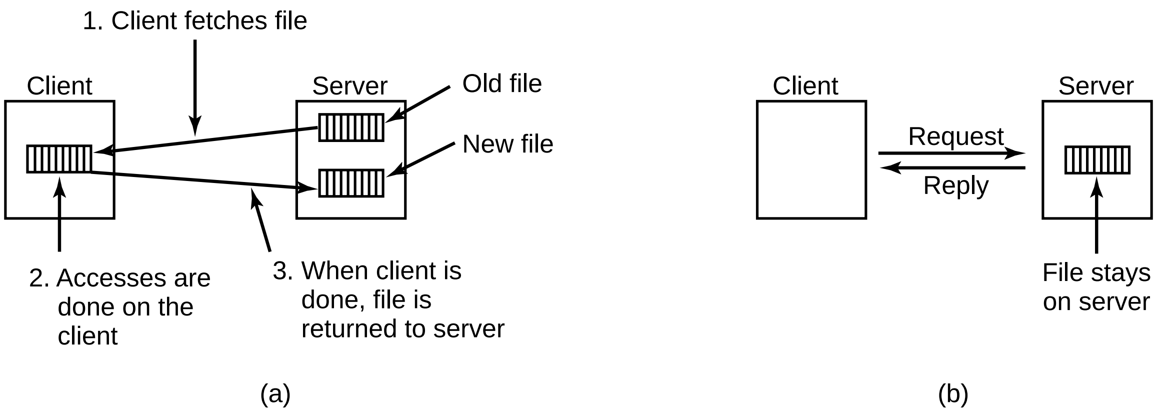

Fig. 8-34. (a) The upload/download model. (b) The remote access model.

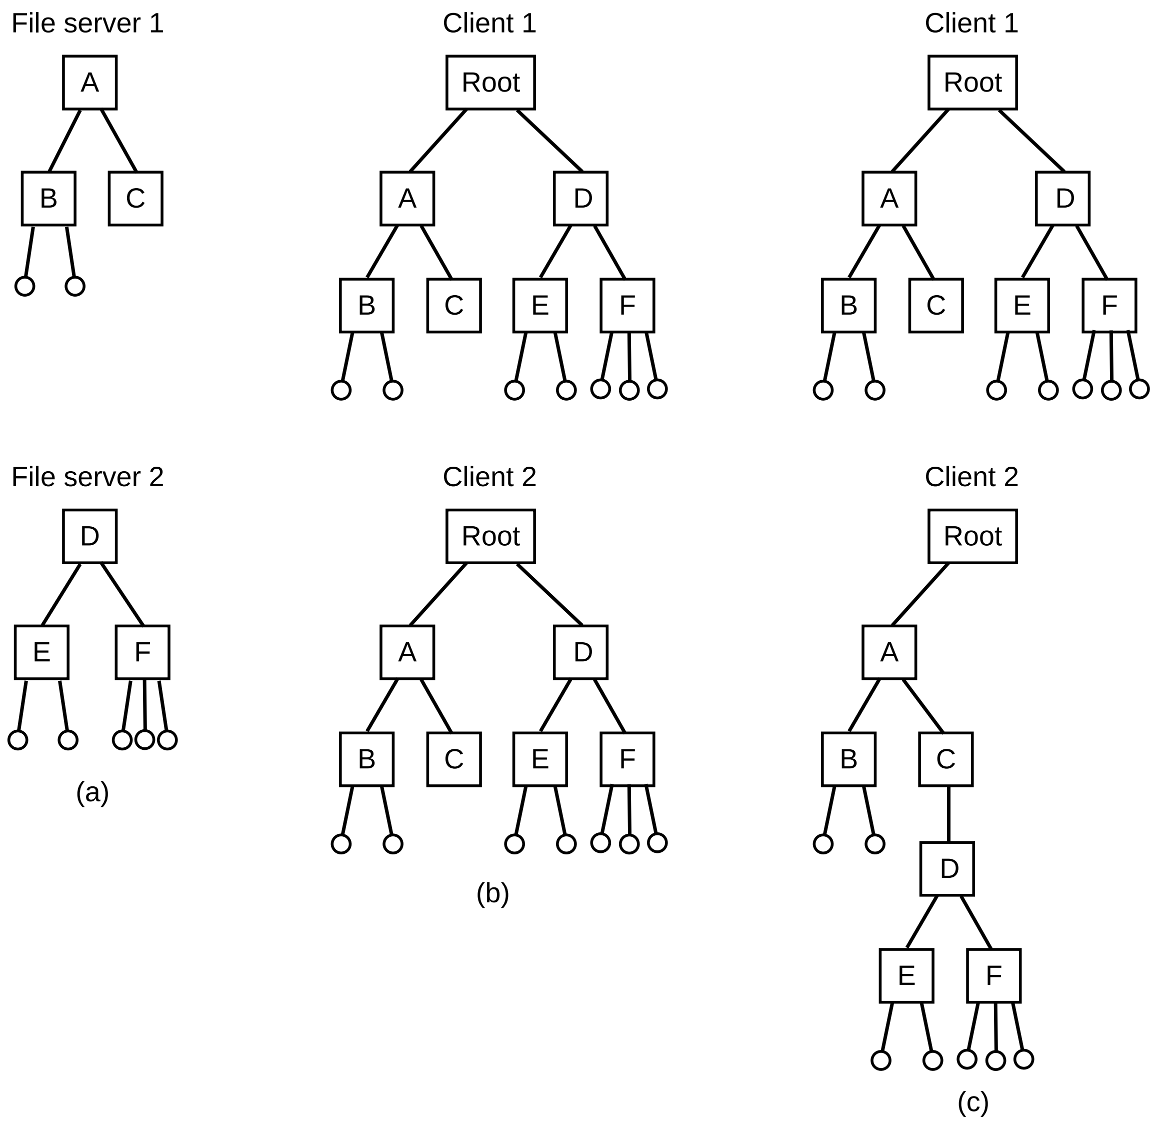

Fig. 8-35. (a) Two file servers. The squares are directories and the circles are files. (b) A system in which all clients have the same view of the file system. (c) A system in which different clients may have different views of the file system.

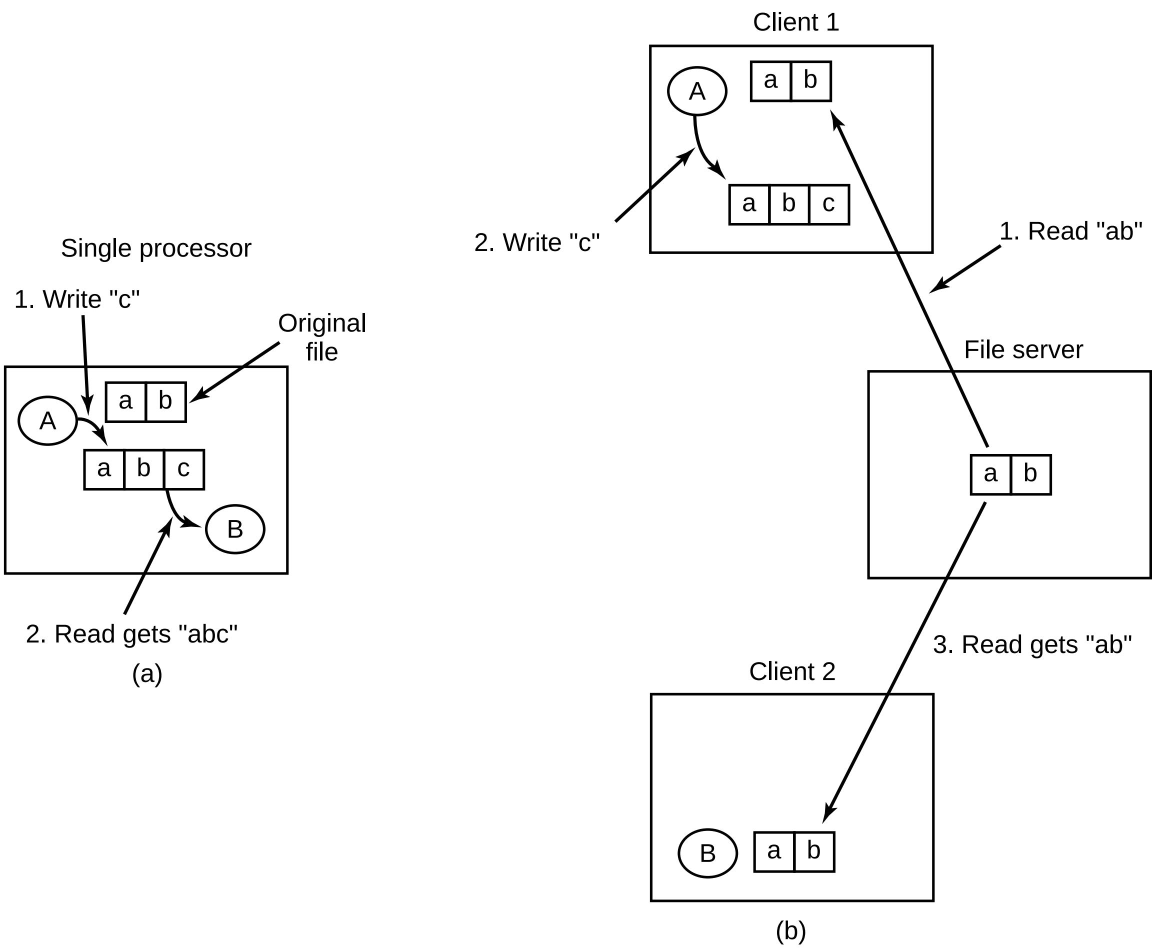

Fig. 8-36. (a) Sequential consistency. (b) In a distributed system with caching, reading a file may return an obsolete value.

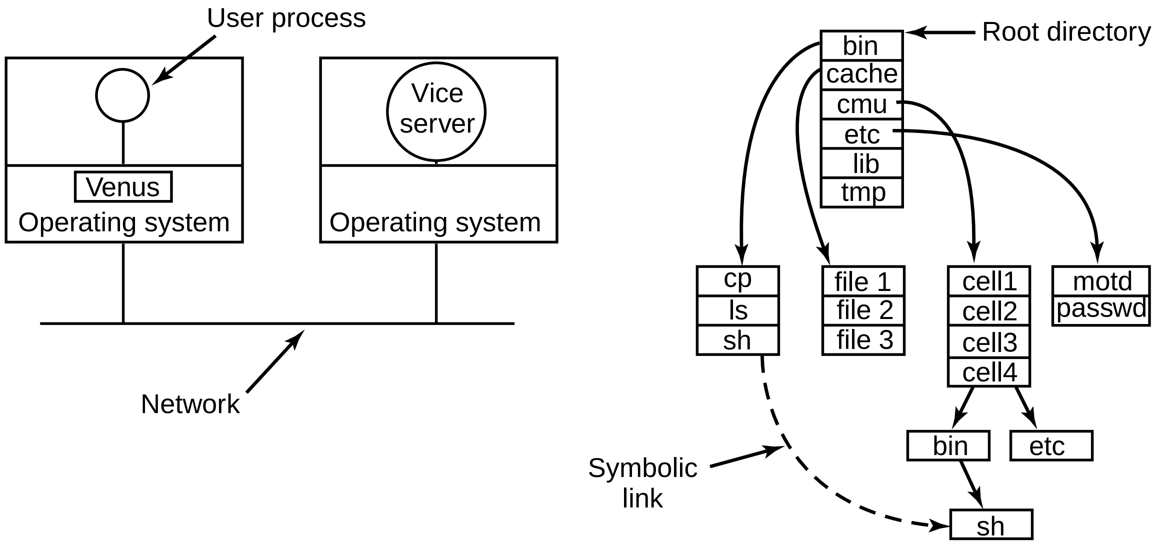

Fig. 8-37. (a) The position of venus and vice in AFS. (b) A client's view of the file system.

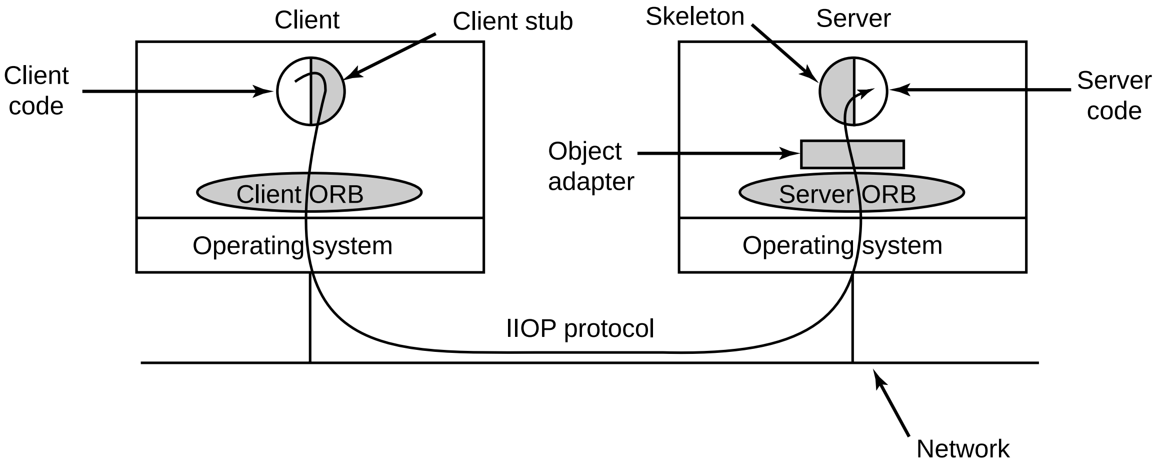

Fig. 8-38. The main elements of a distributed system based on CORBA. The CORBA parts are shown in gray.

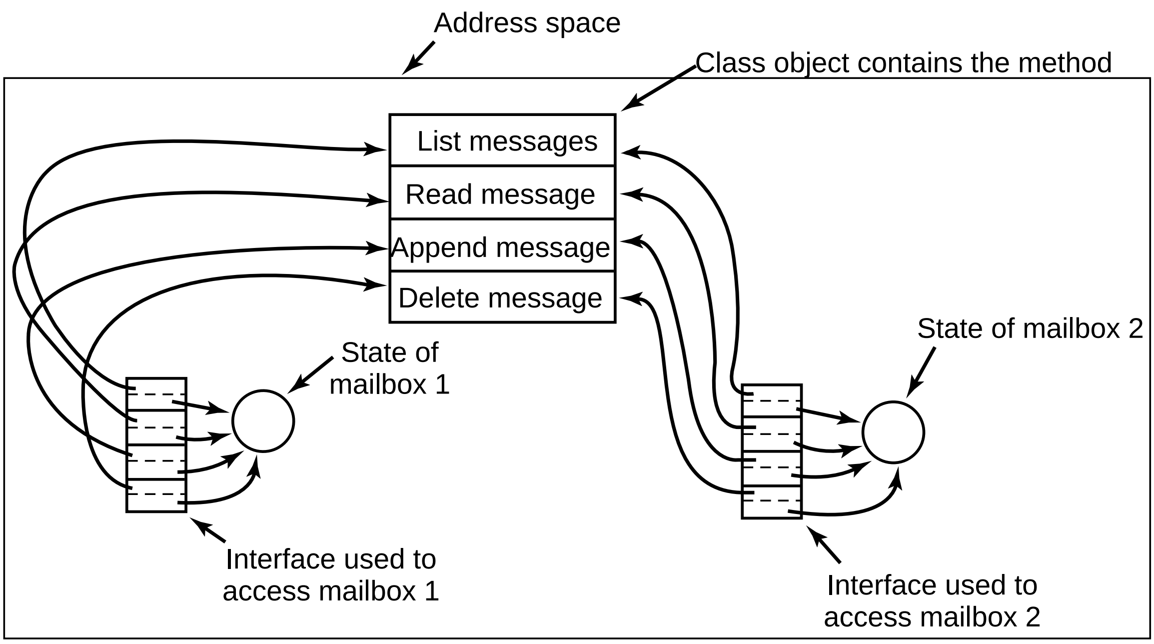

Fig. 8-39. The structure of a Globe object.

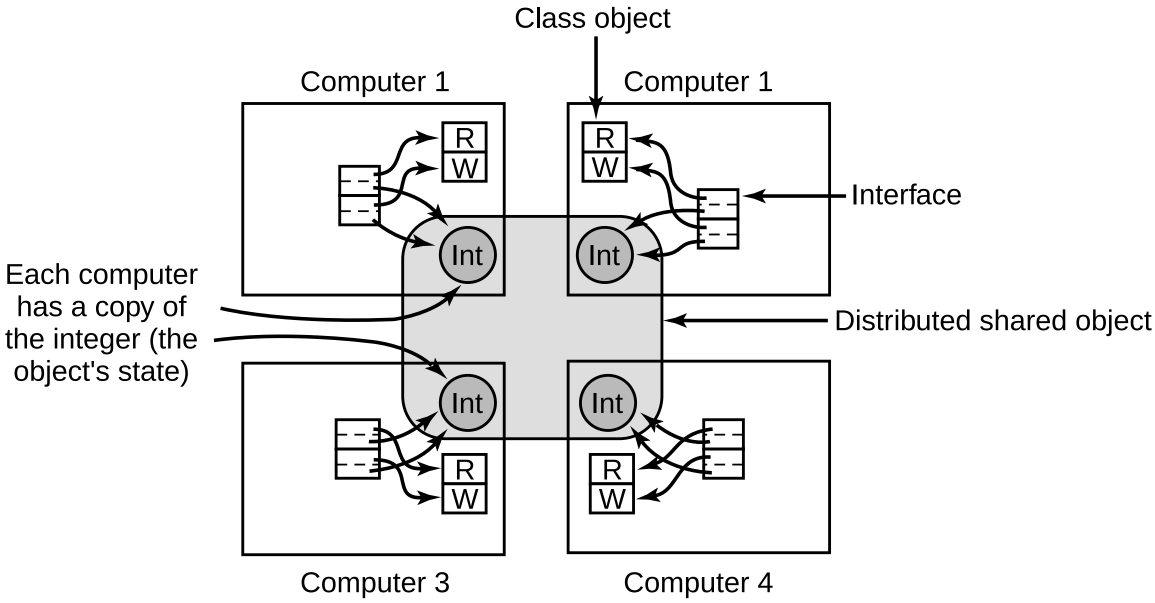

Fig. 8-40. A distributed shared object can have its state copied on multiple computers at once.

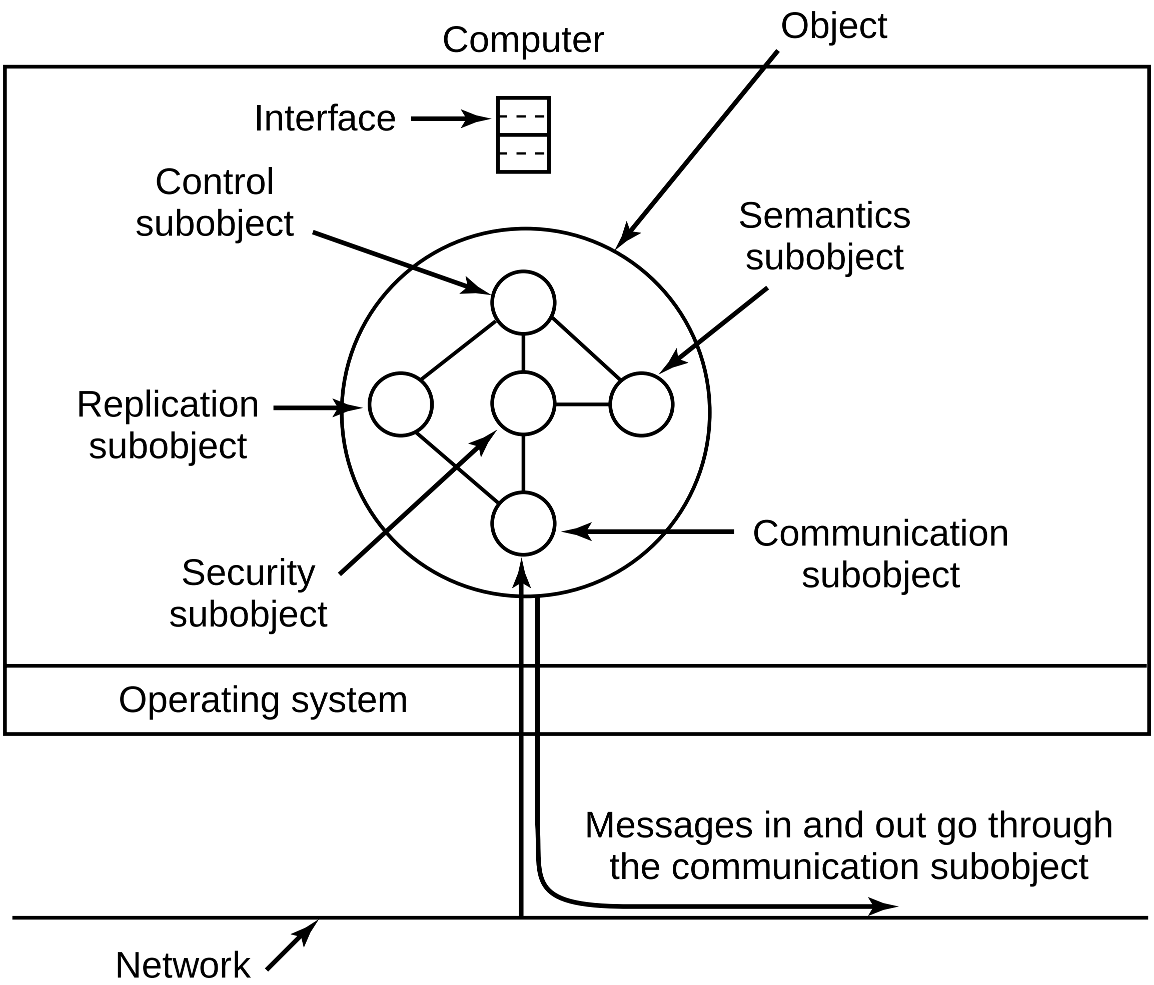

Fig. 8-41. Structure of a Globe object.

```
("abc", 2, 5)
("matrix-1", 1, 6, 3.14)
("family", "is-sister", "Stephany", "Roberta")
```

Fig. 8-42. Three Linda tuples.

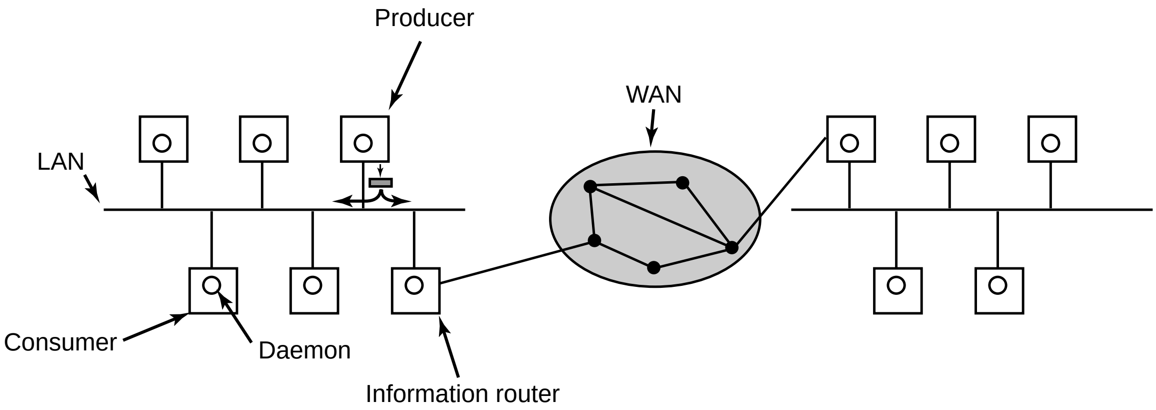

Fig. 8-43. The publish/subscribe architecture.