# Multiprocessors connected by a network

- Restrictions of the bus architecture: high bandwidth, low latency and long length are incompatible
- Shared memory (single address space) vs. multiple private memories
- Centralized memory vs. distributed memory (physically distributed single address space)

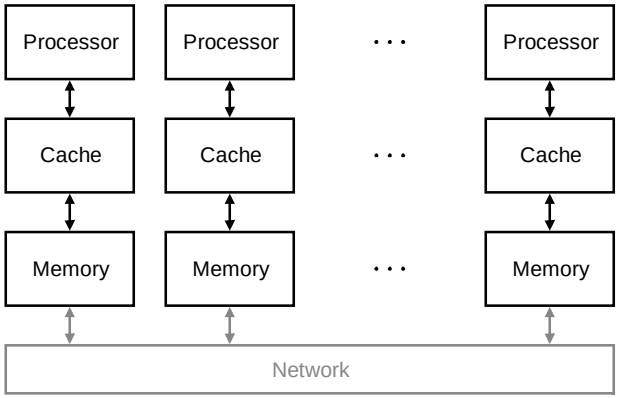

# Parallel programming by message passing

**Summing 100000 numbers by 100 processors**

- The processor containing 100000 numbers sends subsets of 1000 to the others.

- Each processor individually sums the 1000 numbers:

```
sum=0;
for (i=0; i<1000; i++) sum=sum+Al[i];
```

- Adding the 100 partial sums (send(Pr,Val), receive() ):

```
half=100;
limit=100;
repeat
   half=(half+1)/2;
if (Pn>= half && Pn <limit) send (Pn-half, sum)
   if (Pn< limit/2+1) sum=sum+receive()
   limit=half;
until (half==1);
```

# Distributed memory communication

• Explicit communication: *send* and *receive* instead of *load* and *store*
• *Shared virtual memory*: addind a software layer on top of sends and receives to provide a single address space
• Cache coherency (snooping does not work): *directories*. A single directory contains infromation for each block of memory: which caches have copies of the block, whether it is dirty etc.
• When there is a write to a shared block the *directory controller* sends messages to each processor that has a copy of the block.

# Memory allocation (single address space)

• The programmer or the compiler allocates data to the processor that is likely to use it.
• *Cache-only memory*: Add a second level of coherence to the main memory for every processor: moving blocks of memory between the processors.

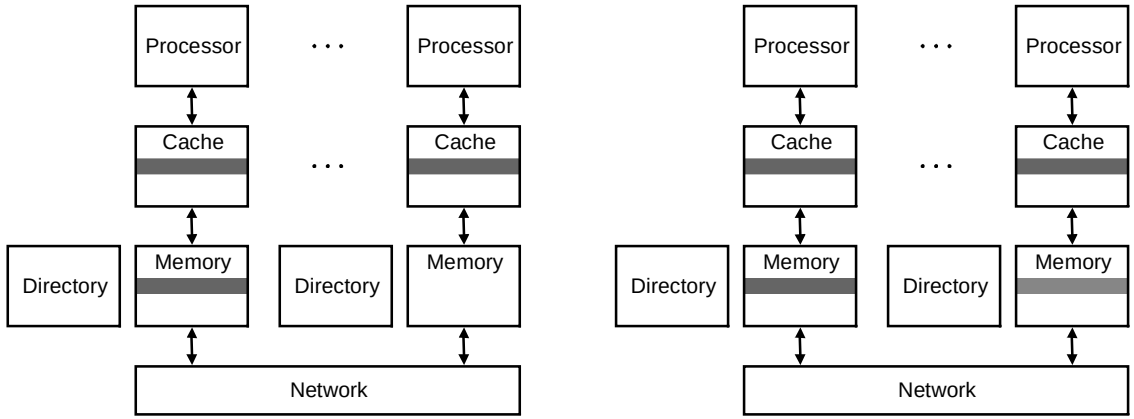

# Cost, performance, and cost/performance

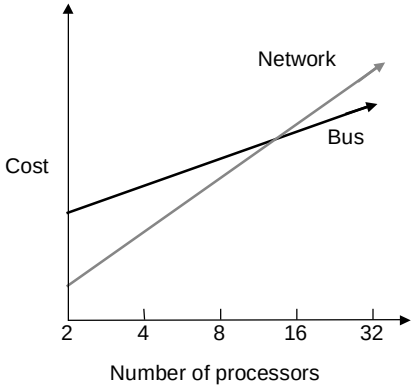


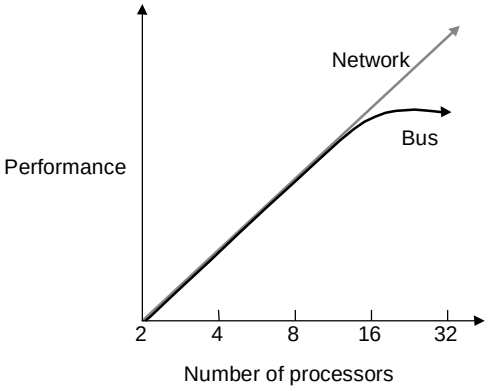


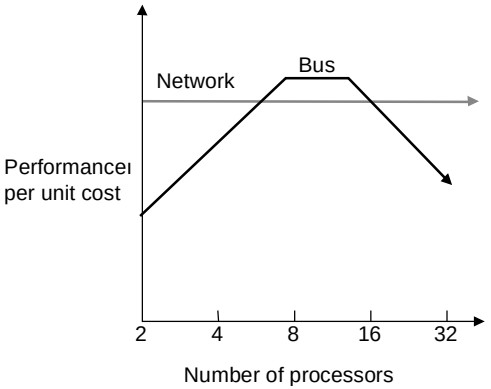

# Clusters

- A network (LAN) of loosely coupled machines
- Drawbacks:
    - High cost of administering
    - Slow connection through the I/O bus;
    - Division of memory
- Hybrid clusters (clustered shared memory)

# Network topology

- *Network cost*: number of switches, number of links on a switch, width (bits) of a link, length of links
- Performance metrics:
  - *Total network bandwidth* (best case)
  - *Bisection bandwidth* (worst case)
- Topologies:
  - *Fully connected network*
  - *Ring*
  - *Multistage networks*

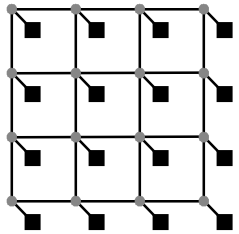

a. 2D grid or mesh of 16 nodes

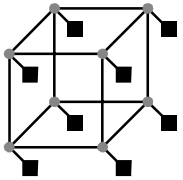

b. n-cube tree of 8 nodes ($8 = 2^3$ so n = 3)

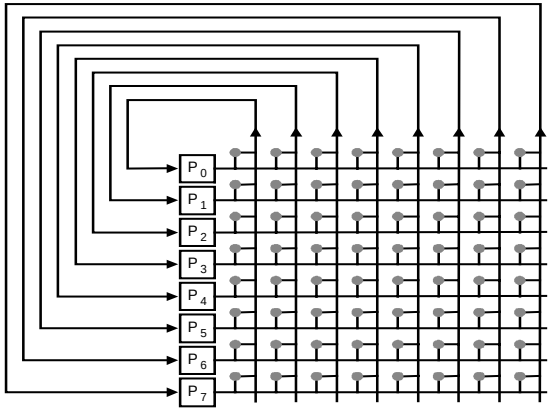


a. Crossbar

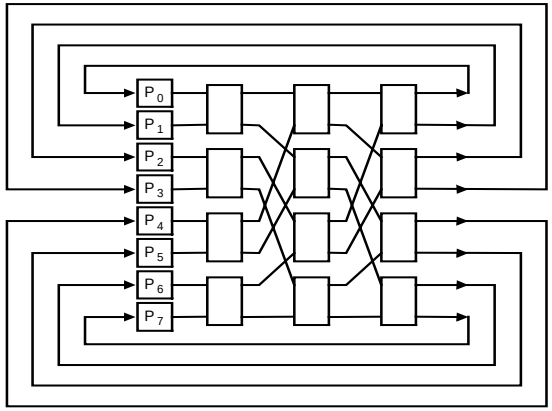


b. Omega network

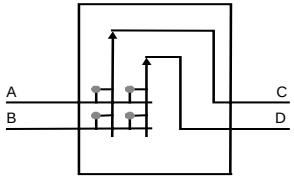


c. Omega network switch box