What is a Distributed System?

- From various textbooks:
  - "A distributed system is a collection of independent computers that appear to the users of the system as a single computer."
  - "A distributed system consists of a collection of autonomous computers linked to a computer network and equipped with distributed system software."
  - "A distributed system is a collection of processors that do not share memory or a clock."
  - "Distributed systems is a term used to define a wide range of computer systems from a weakly-coupled system such as wide area networks, to very strongly coupled systems such as multiprocessor systems."

What is a Distributed System? (cont.)

- A *distributed system* is a set of physically separate processors connected by one or more communication links

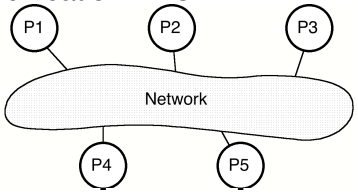

  - Workstation = computer = machine = processor = host = site = node
- Is every system with >2 computers a distributed system??
  - File server, printer server, web server
  - Beowulf-style cluster of workstations
  - 16-processor Cray SV1 at OSC
  - How does a distributed system differ from a parallel system (architecture, OS)?

SIMDs & MIMDs — Parallel or Distributed Architecture?

- Michael Flynn (1966)
  - SISD — single instruction, single data
  - SIMD — single instruction, multiple data
  - MISD — multiple instruction, single data
  - MIMD — multiple instruction, multiple data
- SIMD
  - Many (tens of thousands) of very simple custom cells (processor + memory) in one cabinet, typically organized as a 2D mesh and connected by very high-speed interconnection
  - All processors execute the same instruction concurrently
  - Data parallelism — processor assigned to a unit of data
  - Architecture = parallel, not distributed

SIMDs & MIMDs — Parallel or Distributed Architecture?(cont.)

- Michael Flynn (1966)
  - SIMD — single instruction, multiple data
  - MIMD — multiple instruction, multiple data
- MIMD
  - Few (tens) of powerful commercial RISC processors, connected in some way
    - One cabinet vs. multiple cabinets
    - Distributed memory vs. shared memory
    - Many interconnection alternatives
  - Processors execute different instructions
  - Control parallelism — processor assigned to a unit of code
  - Parallel or distributed??
    - Depends on memory, interconnection
    - Also depends on OS support

SIMDs & MIMDs — Parallel or Distributed Architecture?(cont.)

- More recent (Stallings, 1993)

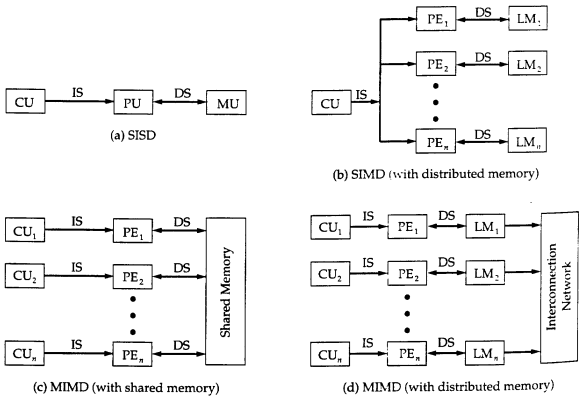

FIGURE 16.16. Alternative Computer Organizations

- Key
  - Contol Unit, Processing Unit, Memory Unit
  - Processing Element, Local Memory
  - Instruction Stream, Data Stream

SIMDs & MIMDs — Parallel or Distributed Architecture? (cont.)

- Tanenbaum (date?)

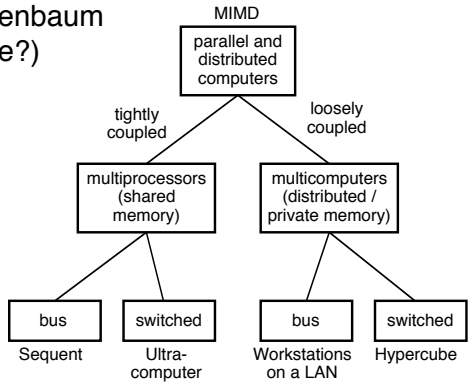

- Few (tens) of powerful commercial RISC processors, connected in some way
  - One cabinet vs. multiple cabinets
  - Distributed memory vs. shared memory
  - Many interconnection alternatives