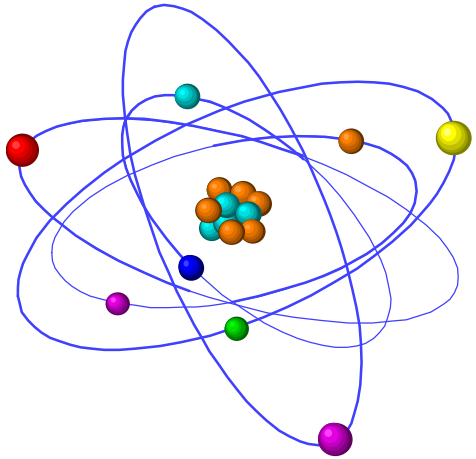

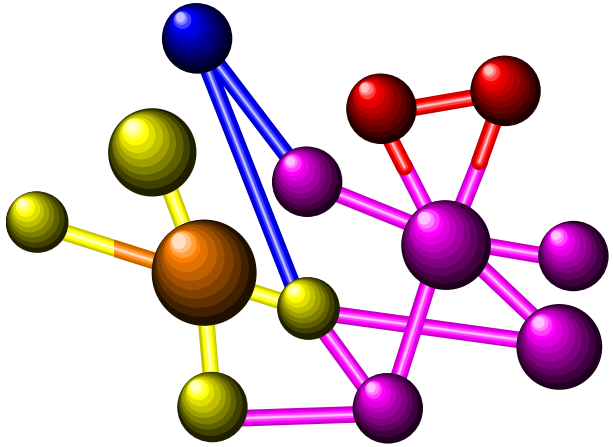

Parallel Computing and Crystallography

Russ Miller

SUNY-Buffalo
Hauptman-Woodward Medical Research Institute

Overview

- **Fine-Grained Parallelism**
- **Coarse-Grained Parallelism**
- **Medium-Grained Parallelism**
- **NOWs** (Networks of Workstations)
- **Examples of Machines**
- **Master/Workers Model**
- **Applications to Crystallography**

Fine-Grained Machines

- **Tens of thousands of Processors**
- **Processors**
  - Slow (bit serial)
  - Small (K bits of RAM)
- **Distributed Memory**
- **Interconnection Networks**
- **Message Passing**
- **Single Instruction Multiple Data (SIMD)**

The SIMD Model

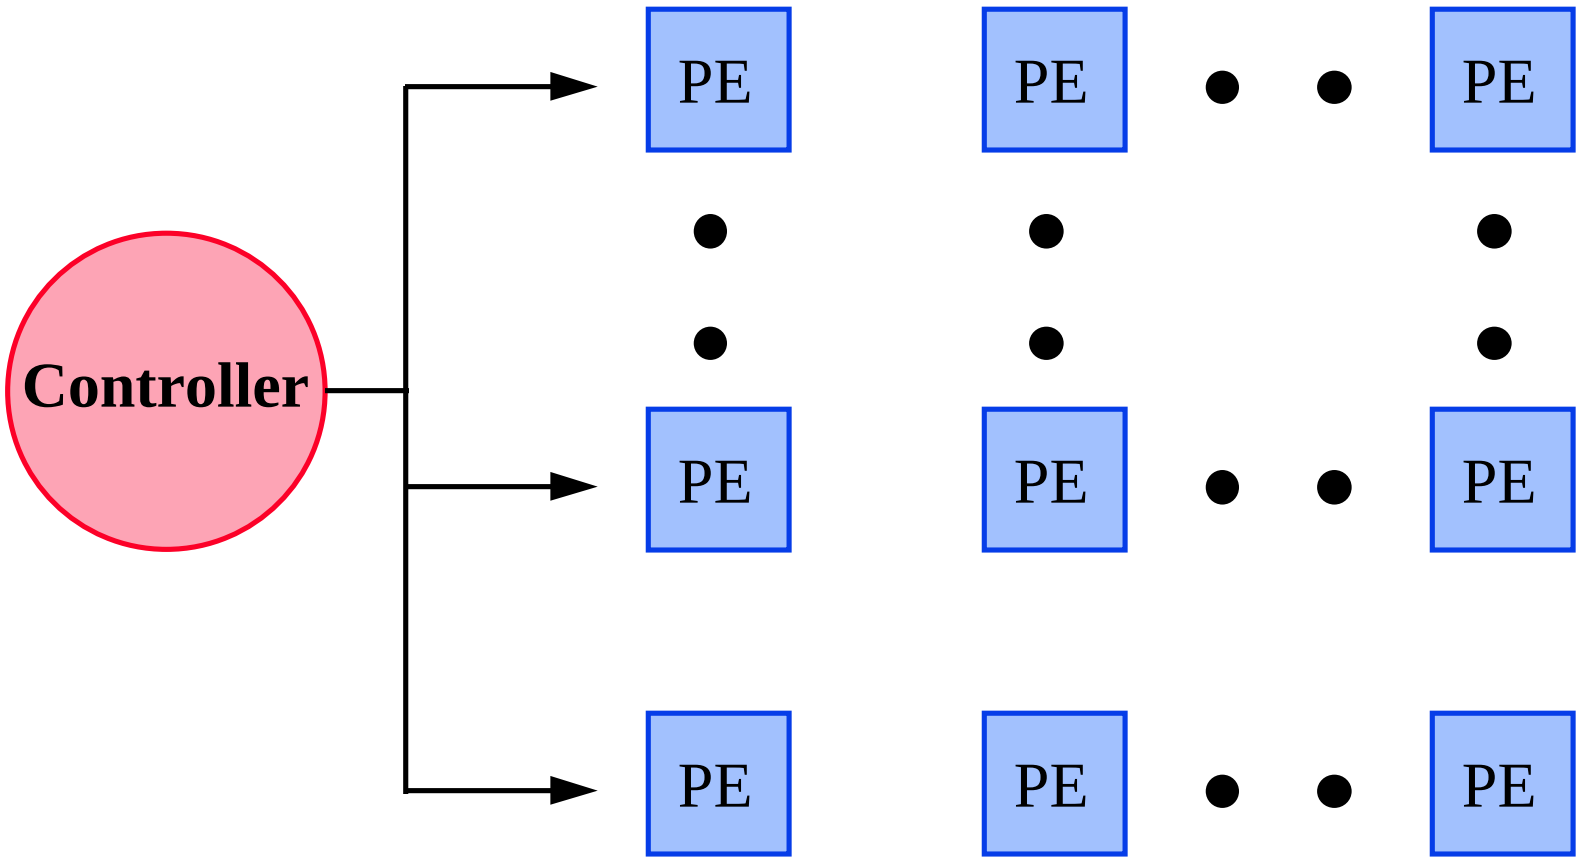

The SIMD Model

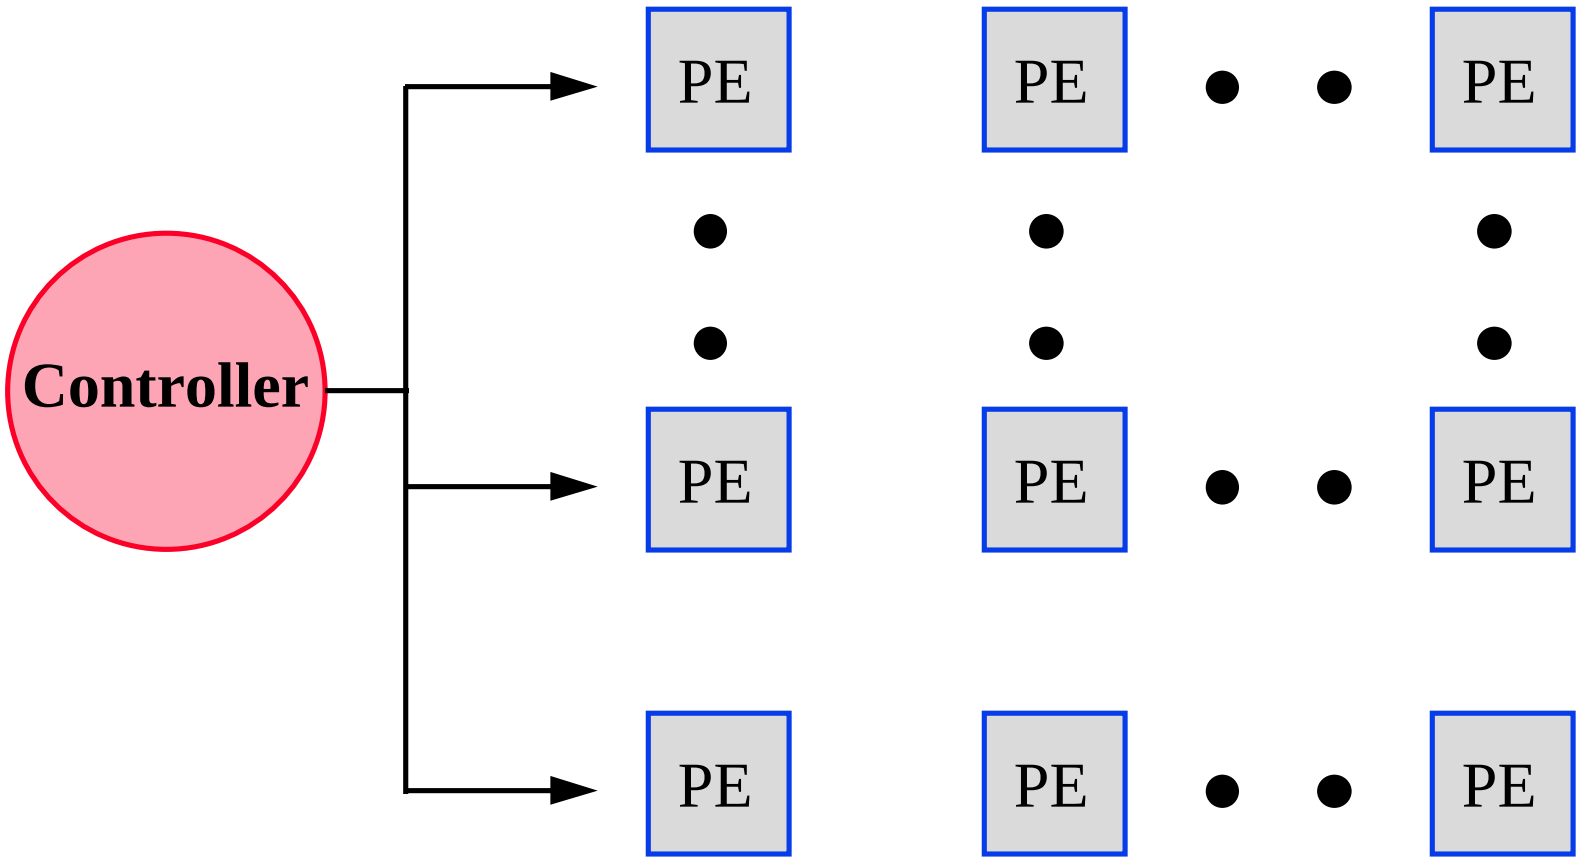

The SIMD Model

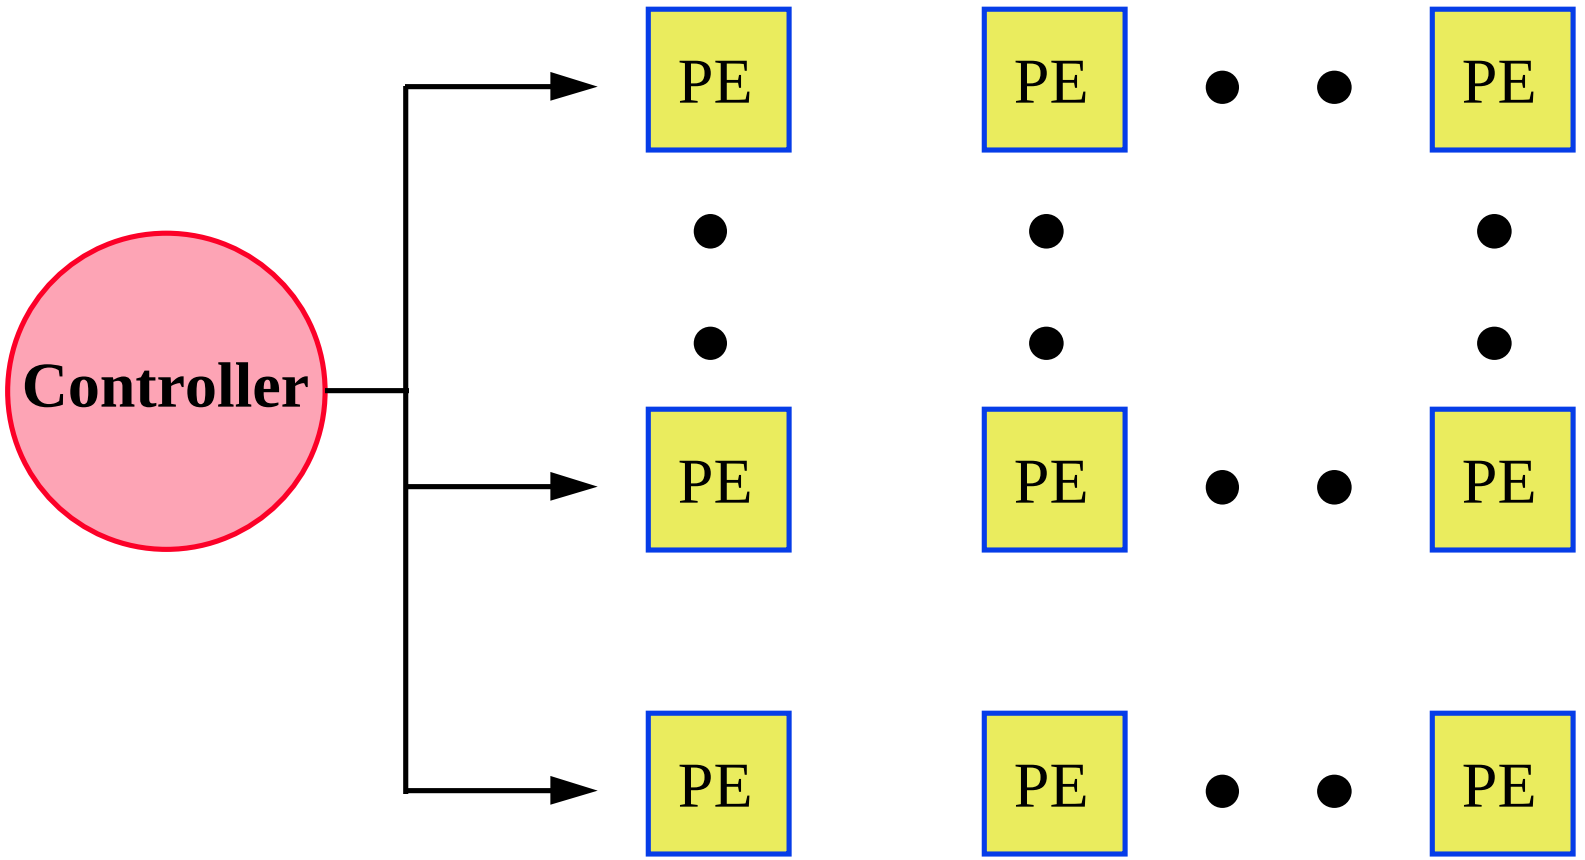

Sample Meshes

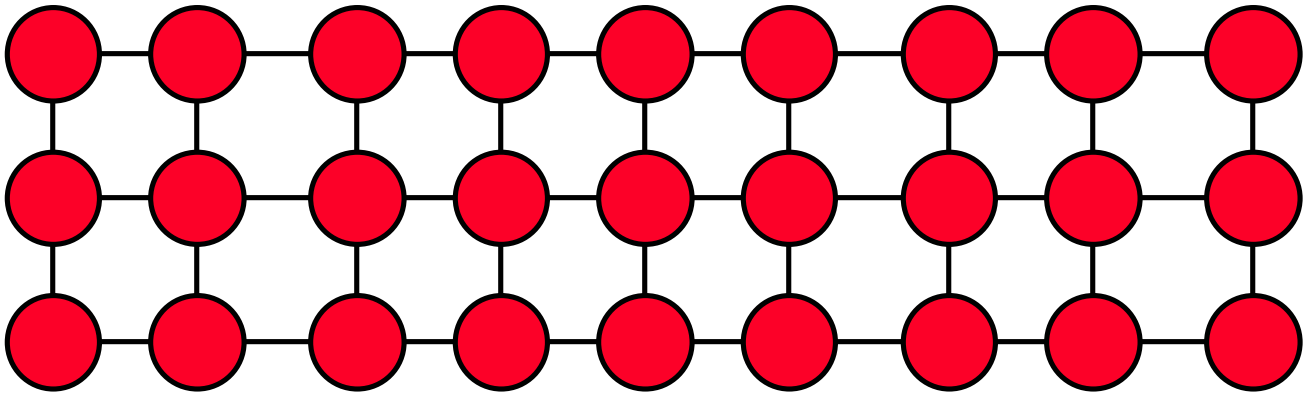

- Massively Parallel Processor (MPP)
- TMC CM-2 (Connection Machine)
- CLIP4 & DAP
- MasPar MP-1/2

MasPar MP-1

- Maximum of 16,384 PEs
- PEs contain a 4-bit ALU
- Peak Performance
  - 1.3 GFLOPS or
  - 200 GIPS (8-bit integer)
- 128×128 array (8-connected)
- Multistage router

The Connection Machine CM-2

- 65,536 1-bit processors
- 12-cube of 4×4 meshes
- 8K bytes of RAM per PE
- 2048 floating point (Weitek) chips
- 32 Gflops peak with Weitek chips

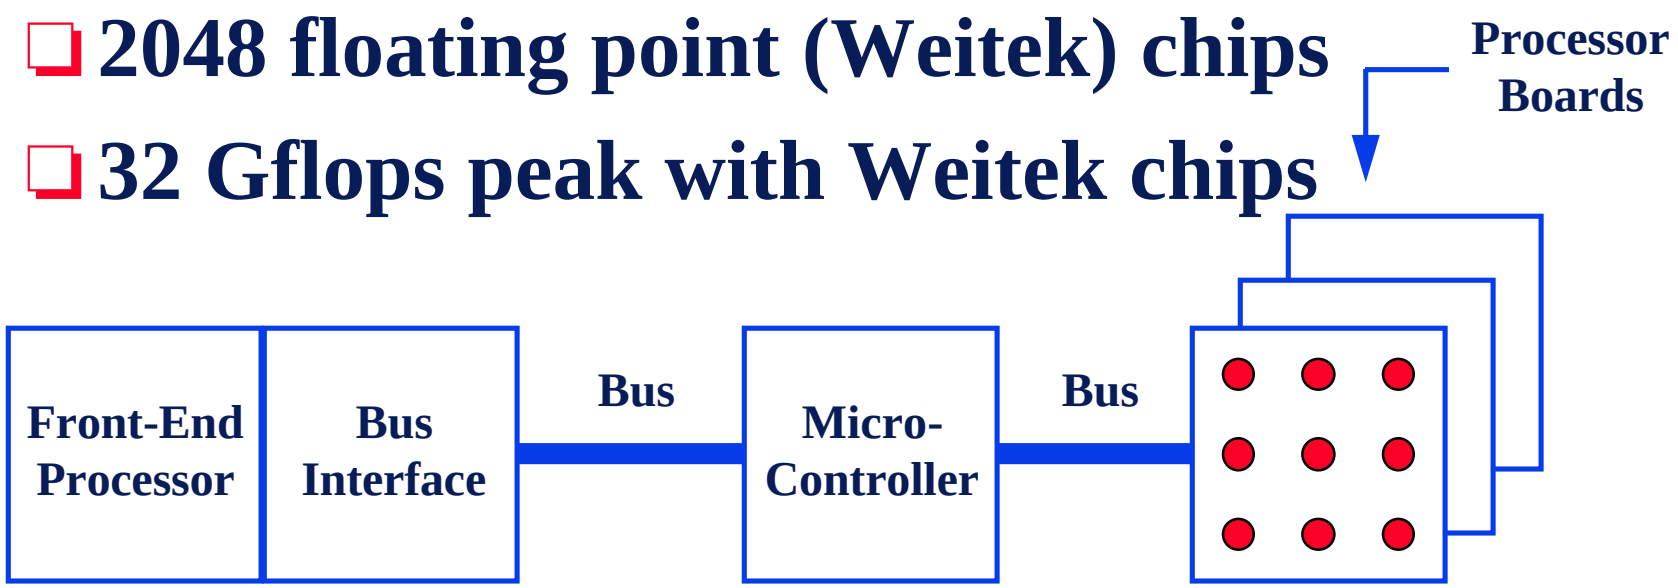

Computing Min on a Tree Machine

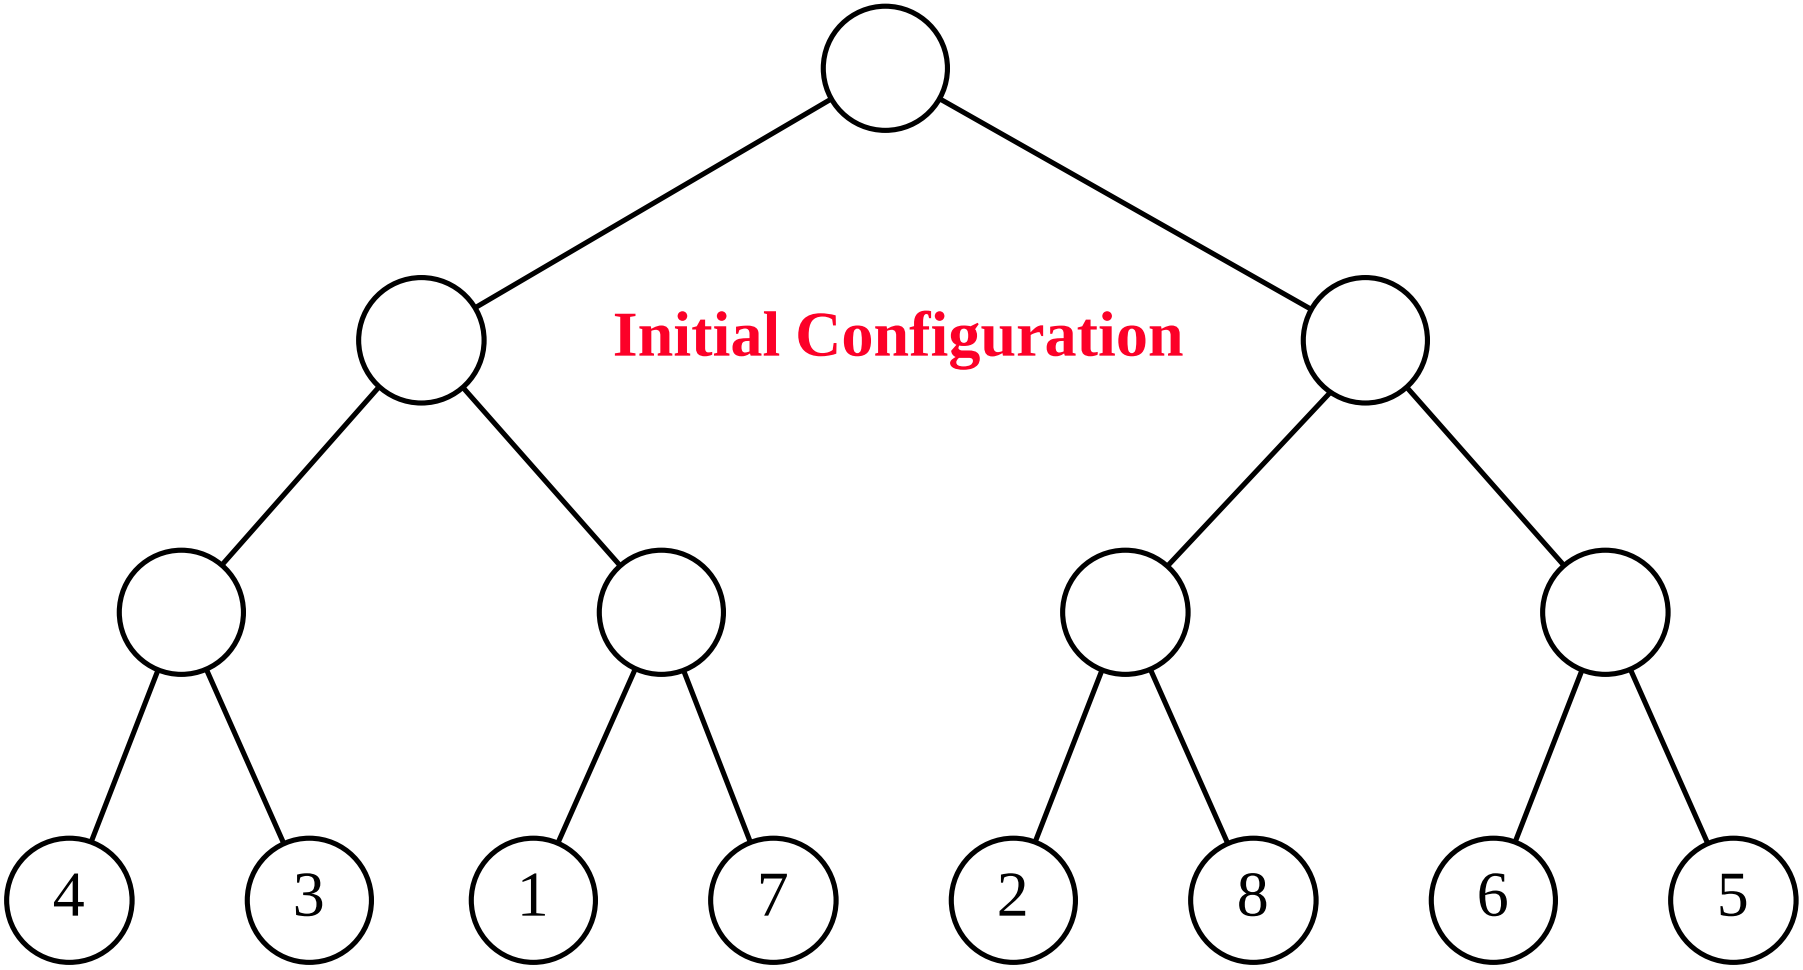

Computing Min on a Tree Machine

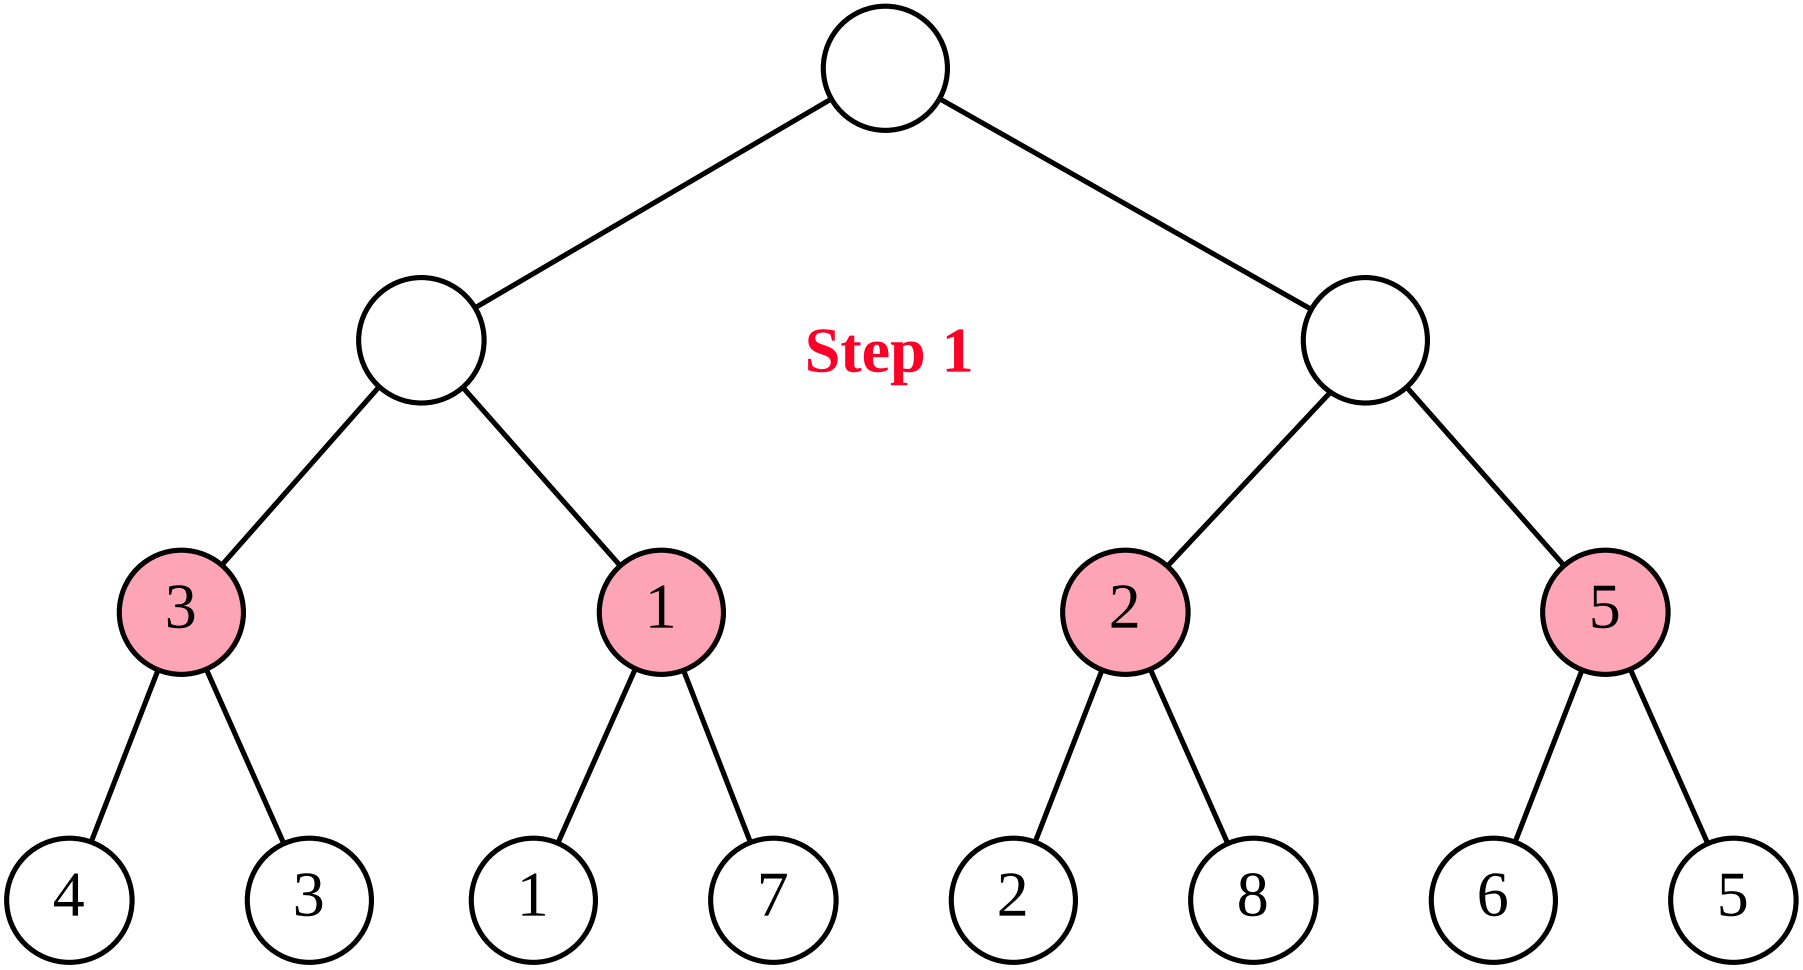

Computing Min on a Tree Machine

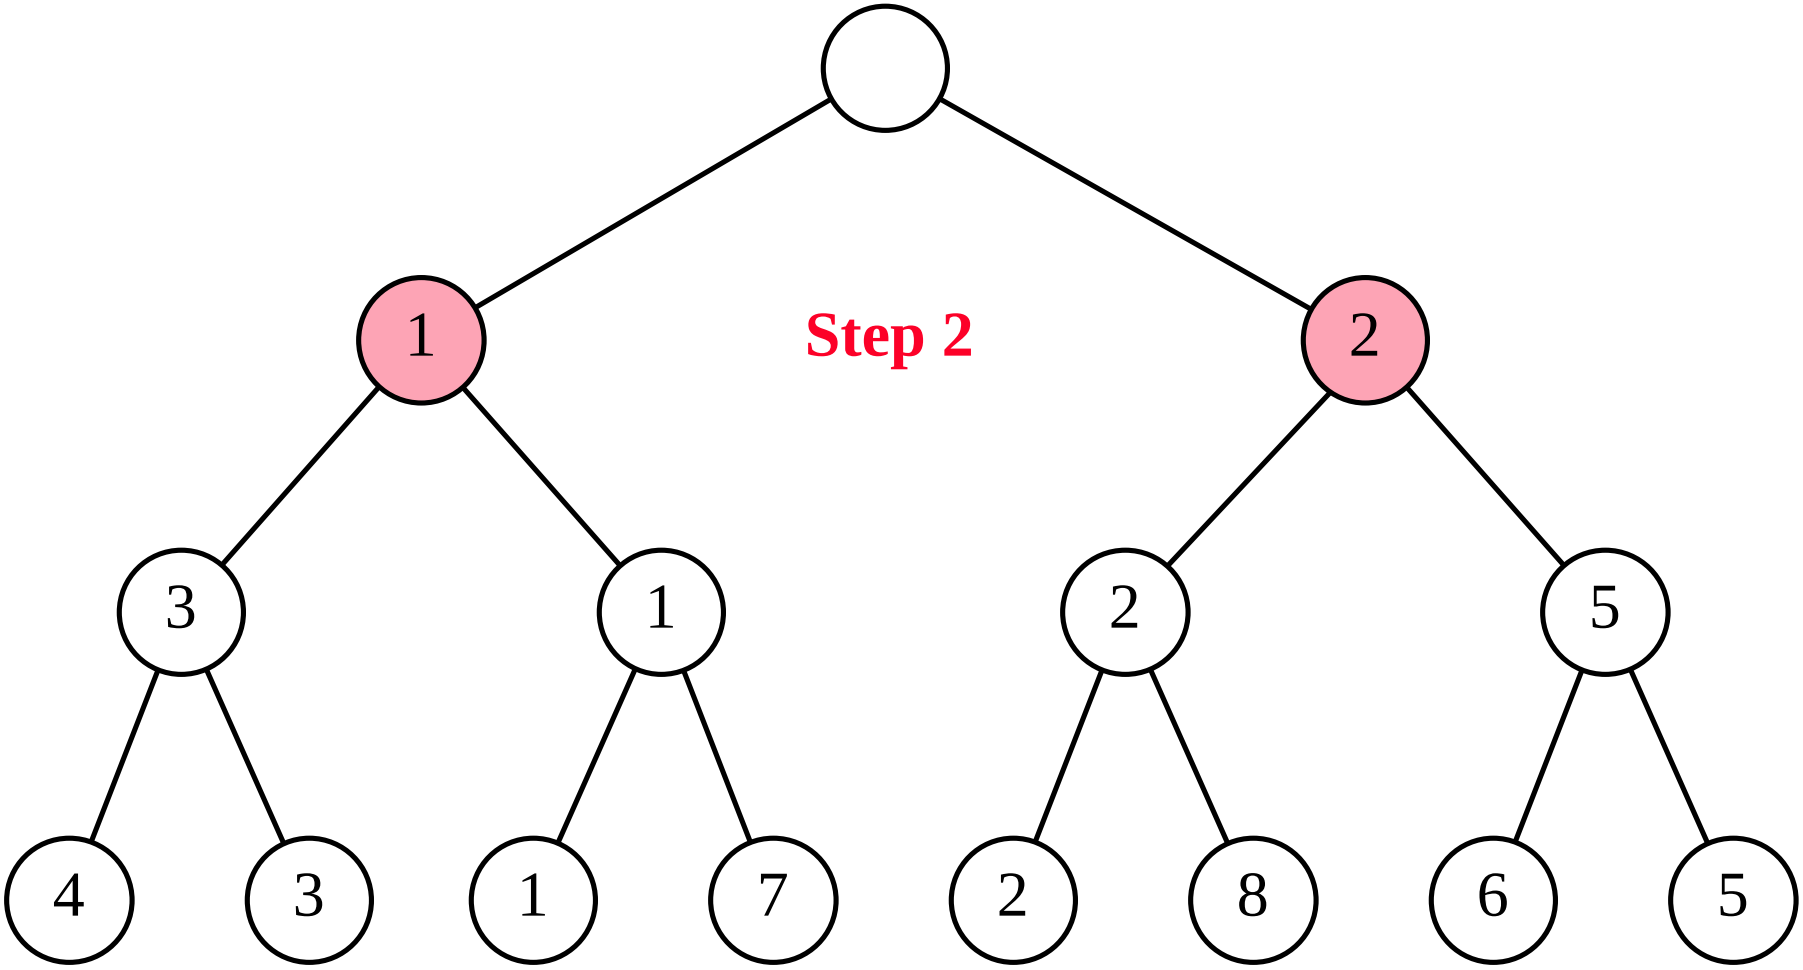

Computing Min on a Tree Machine

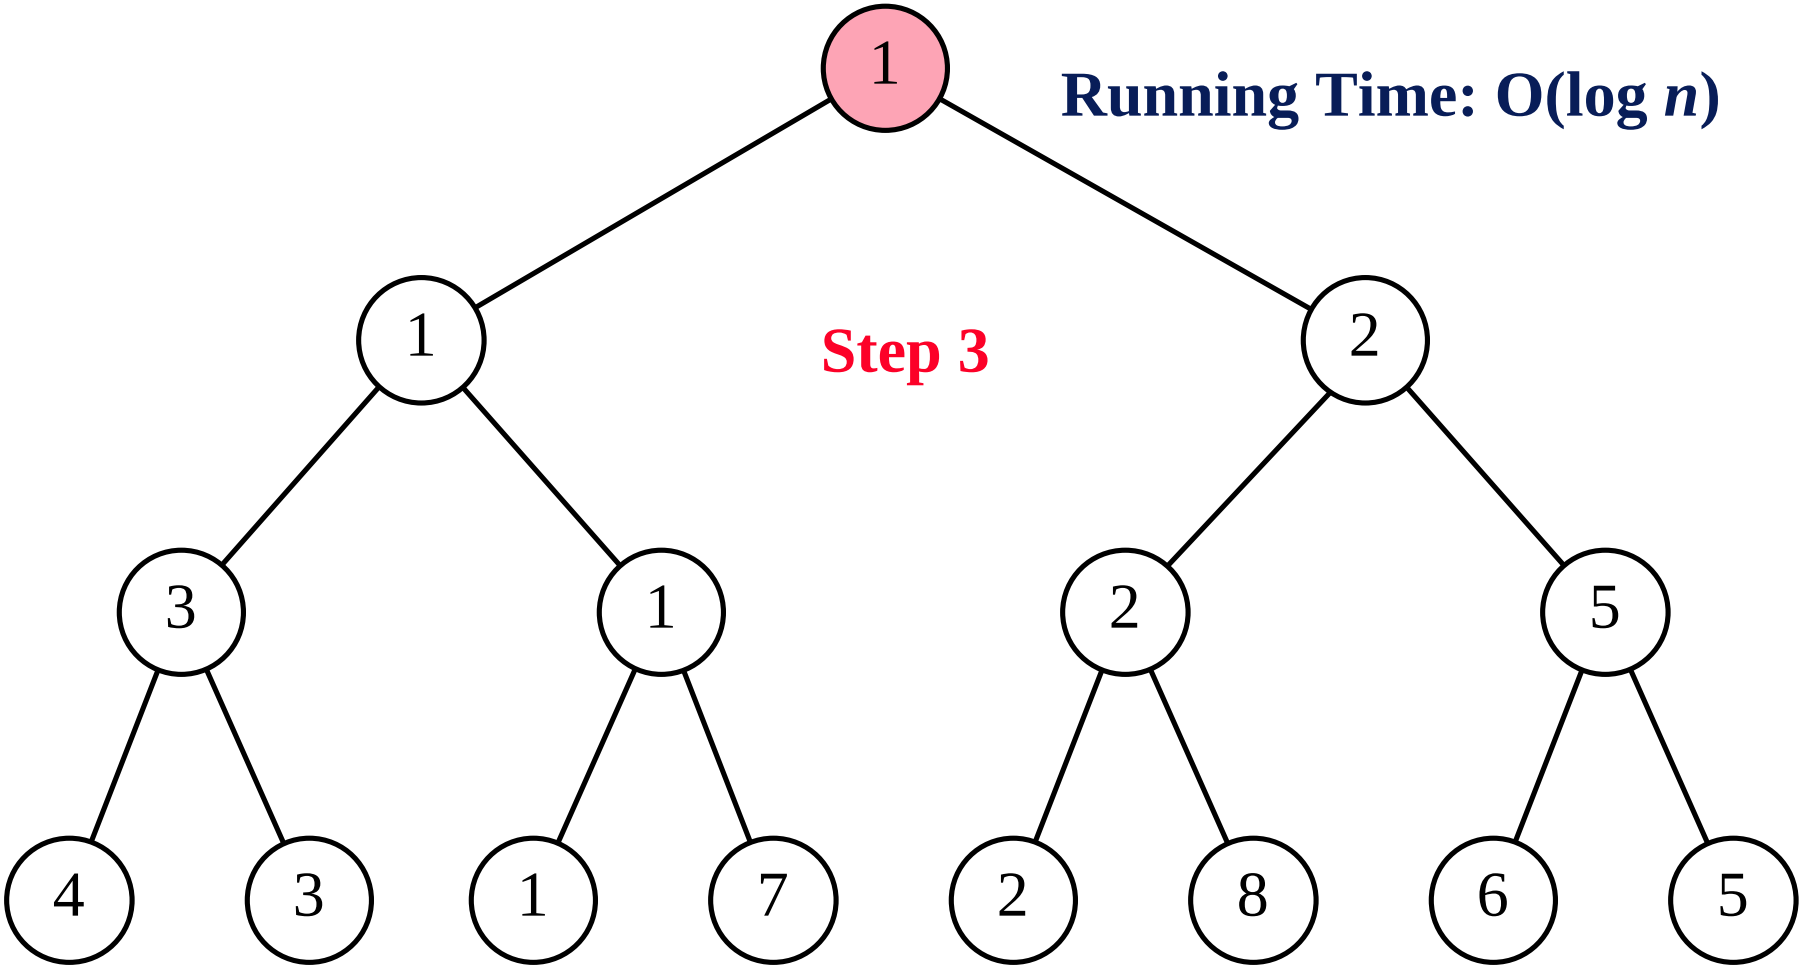

Coarse-Grained Machines

- **Typical Configurations**
  - Hundreds of Processors
  - Processors
    - Powerful (fast CPUs)
    - Large (cache, vectors, multiple fast buses)
  - Memory: Shared or *Distributed-Shared*
  - Multiple Instruction Multiple Data (MIMD)

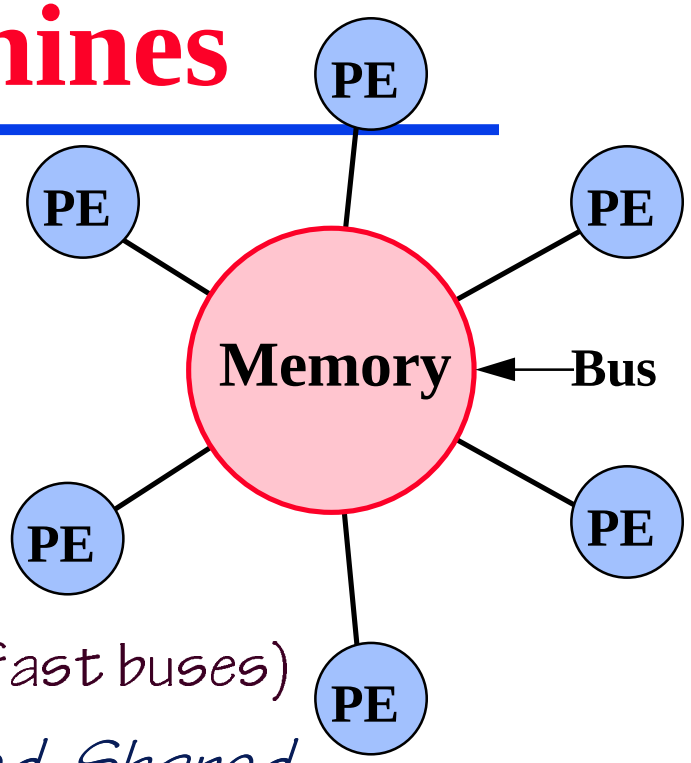

UMA vs. NUMA

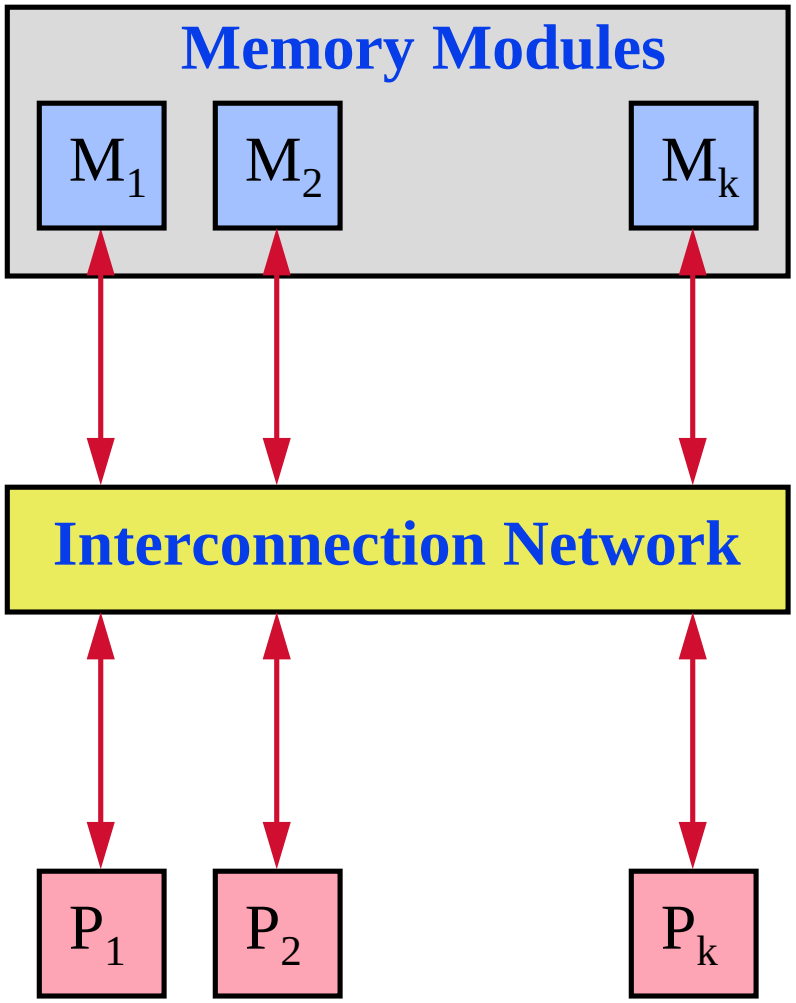
Memory Modules
M1
M2
Mk
Interconnection Network
P1
P2
Pk

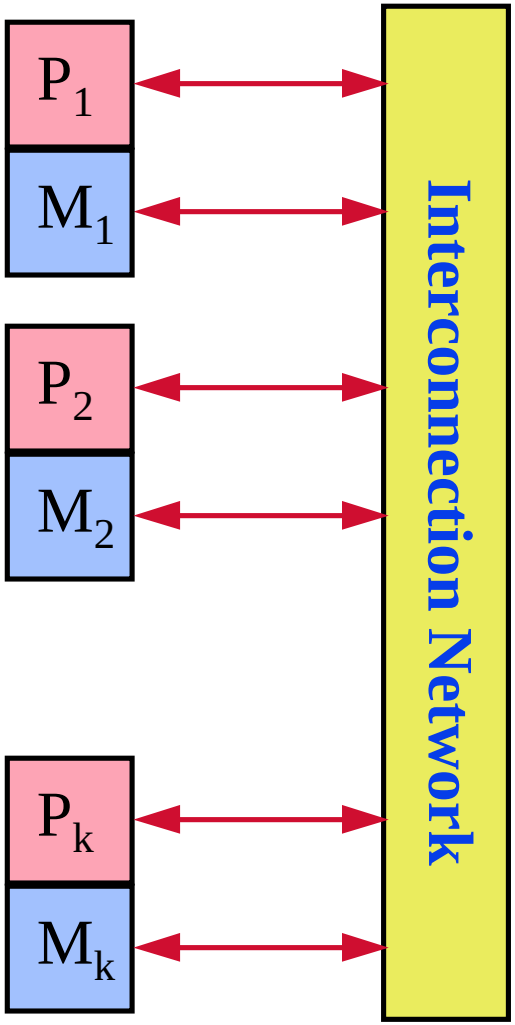
P1
M1
P2
M2
Pk
Mk
Interconnection Network

Coarse-Grained Examples

- **SGI Origin 2000:**
  - PEs (MIPS R10000): Max of 128
  - Peak Performance: 49 Gflops
  - Memory: Max of 256 GBytes
  - Crossbar switches for interconnect
- **HP/Convex Exemplar:**
  - PEs (HP PA-RISC 8000): Max of 64
  - Peak Performance: 46 Gflops
  - Memory: Max of 64 GBytes
  - Distributed crossbar switches for interconnect

Medium-Grained Machines

- **Typical Configurations**
  - Thousands of processors
  - Processors have power between coarse- and fine-grained
  - Investigating shared vs. distributed memory
- **Traditionally: Research Machines**
- **Single Code Multiple Data (SCMD)**

Medium-Grained Example

- **CRAY T3E**
  - Processors
    - DEC Alpha EV5 (600 MFLOPS peak)
    - Max of 2048
  - Peak Performance: 1.2 TFLOPS
  - 3-D torus
  - Memory: 64 MB - 2 GB per processor

NOWs: Networks of Workstations

- Exploit inexpensive Workstations/PCs
- Commodity network
- PVM/MPI
- Condor: “steal” background cycles

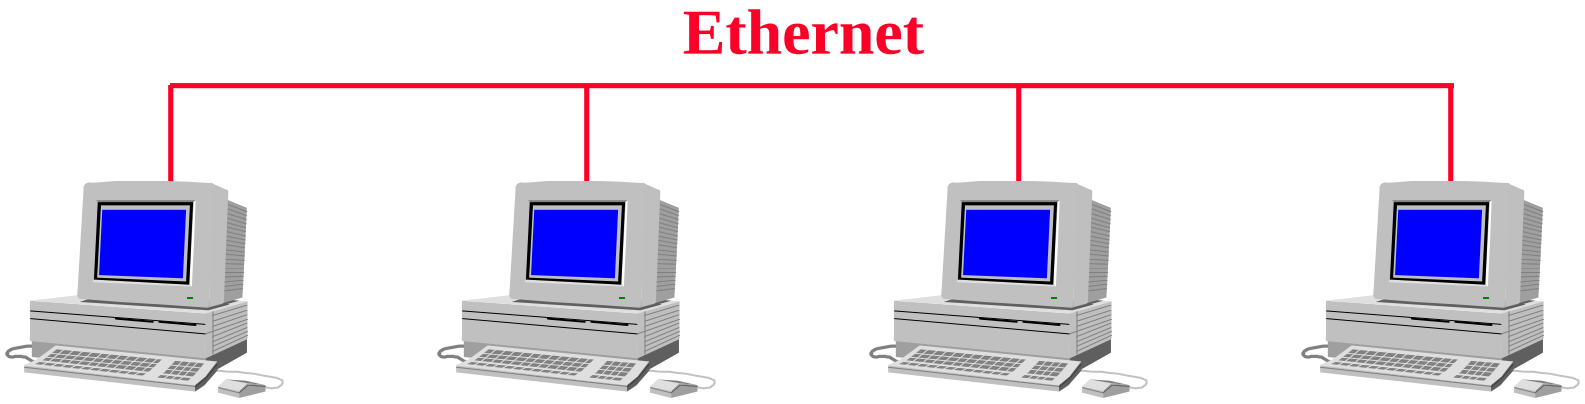

Manager/Workers Model

- **Manager**
  - Initiates computation
  - Tracks progress
  - Handles workers requests
  - Interfaces with user
- **Workers**
  - Spawned and terminated by manager
  - Makes requests to manager
  - Sends results to manager

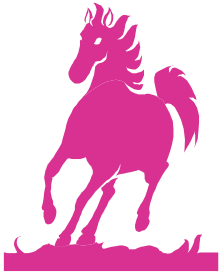

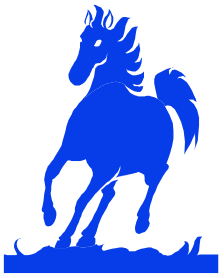

NSF Supercomputing Centers: I

- **Cornell Theory Center (CTC)**
  - IBM SP-2
- **National Center for Supercomputing Applications (NCSA)**
  - HP-Convex Exemplar
  - SGI/CRAY Origin 2000
  - SGI Power Challenge Array
  - TMC CM-5

NSF Supercomputing Centers: II

- **Pittsburgh Supercomputing Center (PSC)**
  - CRAY C90 and J90s
  - CRAY T3E & T3D
  - Alpha Cluster
- **San Diego Supercomputing Center (SDSC)**
  - CRAY C90
  - CRAY T3D
  - Intel Paragon

High-Performance Computing

- **Past Decade:**
  - Flurry of activity in computer development
  - Microprocessors performance astonishing
- **Microprocessor-based supercomputers are now technology of choice**
- **Affordable parallel machines outperform traditional vector supercomputers**
- **New machines are reliable, provide standards and 3rd-party software**

DOE's ASCI Program

- **Department of Energy's Accelerated Strategic Computing Initiative**
  - SNL: 1.8 Tflops Intel Paragon based on *Pentium* processors
  - LLNL: 3.3 Tflops IBM SP (RS6000s) 12/98
  - LANL: 3 Tflops SGI-CRI (alphas) 12/98
  - Planned: 30 & 100 Tflops

Crystallography and Parallel Computers

- **Opportunities for Parallelization**
  - Fast Fourier Transform (FFT)
  - Structure Factors
  - Visualization (Ray Tracing)
- **High Throughput**
  - Use NOW model
- **Examples**
  - *SnB:* NOW model
  - X-PLOR: NOW model

Software Development Considerations

- **Using parallel machines:**
  - Inexpensive access
  - Good access
  - In-house expertise
- **Develop new theories**
- **Large production runs**
- **Production Codes: *Portability!!!***
  - Machines are Cheap
  - People are Expensive

Cost-Effective Solutions

- **NOWs**
  - Network of Pentiums vs.
  - Big Iron (parallel machines)
- **Manager/Workers Model**
- **PVM/MPI/Condor**

Comparative Example

- **Given US $1M**
  - Purchase approx 500 Pentium boxes
  - Purchase approx 32 PE SMP
- **SPECfp95**
  - Pentiums are approx 3-5X slower

Comparative Example

- **Result**
  - Factor of 3-5X improvement in available cycles
- **Considerations**
  - Maintenance
  - Usability
  - Networking

Summary

- **Introduced models, machines, & terminology**
- **Discussed potential uses**
- **Recommendations for using multiprocessor machines**
  - Manager/Workers
  - PVM/MPI
- **Recommendations for *new purchases***
  - NOWs
  - Condor/MPI/PVM