

# Architecture of UNIX Systems

Dirk Bartz  
University of Tübingen

WZ  
GR 15

IEEE Visualization 2000  
Tutorial 7


## Outline

- Parallel Approaches
- Taxonomy
- Memory Models
- Programming Models
- Example Architectures

IEEE Visualization 2000  
Tutorial 7

## Parallel Approaches (1)


Cluster of Workstations/PCs (loosely coupled)



IEEE Visualization 2000  
Tutorial 7

## Parallel Approaches (2)


Multi-processor computers (tightly coupled / SMP)



IEEE Visualization 2000  
Tutorial 7

## Parallel Approaches (3)

Mixture of both



IEEE Visualization 2000  
Tutorial 7

## Outline

- Parallel Approaches
- Taxonomy
- Memory Models
- Programming Models
- Example Architectures

IEEE Visualization 2000  
Tutorial 7

## Taxonomy

Flynn's taxonomy:

- SISD - standard Workstation/PC class
- SIMD - massively-parallel computer
- MISD - does not exist
- MIMD - standard parallel computer type

IEEE Visualization 2000  
Tutorial 7

## Outline

- Parallel Approaches
- Taxonomy
- Memory Models
- Programming Models
- Example Architectures

IEEE Visualization 2000  
Tutorial 7


## Memory Models (1)

- Distributed memory
- Shared-memory
  - Uniform Memory Access (UMA)
  - Non-Uniform Memory Access (NUMA)  
(distributed shared-memory)

IEEE Visualization 2000  
Tutorial 7

## Memory Models (2)


- Uniform Memory Access (UMA)



IEEE Visualization 2000  
Tutorial 7

## Memory Models (3)

- Non Uniform Memory Access (NUMA)



IEEE Visualization 2000  
Tutorial 7

## Memory Models (4)

- Why NUMA architecture?
    - UMA system bus gets saturated (if too much traffic)
    - UMA crossbar gets too complex (too expensive)
- ➔ UMA architecture does not scale beyond a certain level

IEEE Visualization 2000  
Tutorial 7

## Memory Models (5)

- **Typical NUMA problems**
  - High synchronization costs (of subsystem interconnect)
  - High memory access latencies (not in my experience)
  - Might need memory sensitive strategies
    - ➔ loose shared-memory advantage

IEEE Visualization 2000  
Tutorial 7

## Outline

- **Parallel Approaches**
- **Taxonomy**
- **Memory Models**
- **Programming Models**
- **Example Architectures**

IEEE Visualization 2000  
Tutorial 7

## Programming Models (1)

### Message-passing (PVM, MPI)

- Individual processes exchange messages
- Works on clusters and on parallel computers (topology transparent to user)
- Manual parallelization

IEEE Visualization 2000  
Tutorial 7

## Programming Models (2)

### Threading (OpenMP/threads)

- Efficient only on shared memory systems
- One process (environment), multiple threads
- Cheap, implicit communication
- Different scheduling approaches
- Limited (semi-)automatic parallelization

IEEE Visualization 2000  
Tutorial 7

## Outline

- **Parallel Approaches**
- **Taxonomy**
- **Memory Models**
- **Programming Models**
- **Example Architectures**

IEEE Visualization 2000  
Tutorial 7

## Examples Architectures (1)

### Terminology

- A NUMA subsystem/hypernode is the lower hierarchy element of the total system
- A Node is a processing entity; a whole computer in a cluster, or a CPU/processor in a parallel computer

IEEE Visualization 2000  
Tutorial 7

## Examples Architectures (2)


### Terminology, cont'd

- A process is an execution environment. On message-passing systems, individual processes are the smallest processing entity; on thread systems, a process is providing the environment for the threads.

IEEE Visualization 2000  
Tutorial 7

## Examples Architectures (3)

### Sun Enterprise x500/E10000




Interconnect: crossbar/bus  
@ 12.8/3.2 GB/s  
Max #CPUs: 30/64 @ 400 MHz  
Max memory: 8/64 GB/NUMA  
Pthread sched': Mixed  
OpenMP: Third Party  
Node boards: 8/16  
2/4 CPUs  
-/ I/O on board

IEEE Visualization 2000  
Tutorial 7

## Examples Architectures (4)

### Hewlett-Packard N-class




Interconnect: 2 x bus  
@ 1.9 GB/s  
Max #CPUs: 8 @ 440 MHz  
Max memory: 16 GB/UMA  
Pthread sched': User/Kernel  
OpenMP: Third Party

IEEE Visualization 2000  
Tutorial 7

## Examples Architectures (5)

### SGI 2x00 series (Origin 2000)



Interconnect: crossbar/hypercube  
@ 49.9 GB/s  
Max #CPUs: 64 @ 300 MHz  
Max memory: 128 GB/NUMA  
Pthread sched': Mixed  
OpenMP: C/C++  
(MIPSpro 7.3)

Node boards: 32  
2 CPUs

IEEE Visualization 2000  
Tutorial 7