Multiprocessor Architectures

Wolfgang Schreiner
Research Institute for Symbolic Computation (RISC-Linz)
Johannes Kepler University, A-4040 Linz, Austria

Wolfgang.Schreiner@risc.uni-linz.ac.at
http://www.risc.uni-linz.ac.at/people/schreine

Shared Memory Machines

- Multiprocessors
- Closely coupled processors
- UMA architectures

  Uniform Memory Access
- SMP architectures

  Symmetrical Multiprocessing

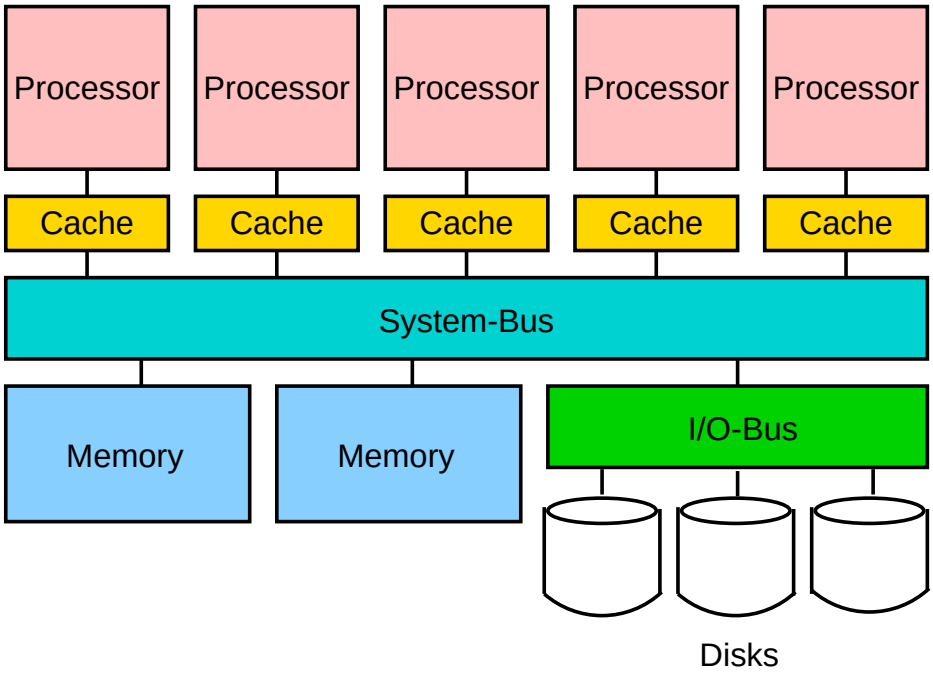

Symmetrical Multiprocessing

- Unix operating system
- Parallelism transparent to users

  Operated like any workstation.

- Queue of executable processes
- Automatic mapping to processors

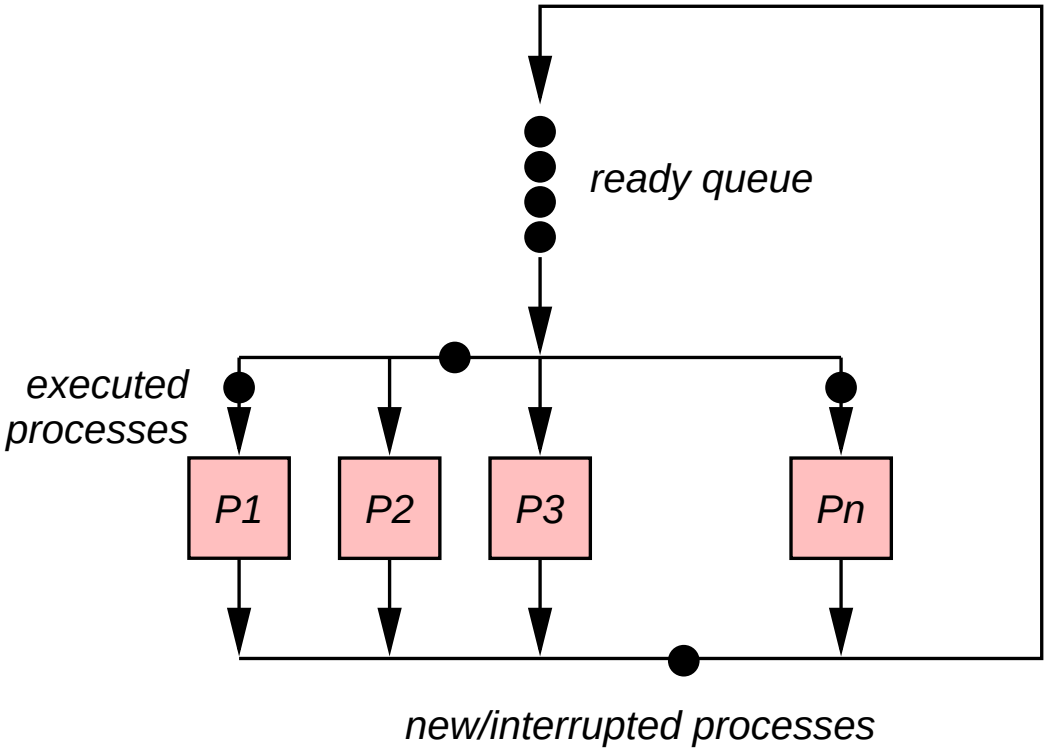

Symmetrical Multiprocessing

Scheduling algorithm

```
foreach P[i] do in parallel
  loop forever
    get job J with higest priority
      from ready queue
    execute at most one timeslice
      of J
    if J not yet terminated then
      put J back into queue
```

Load optimally balanced among processors.

Cache Protocols

- Cache = fast buffer memory.
- Holds recently accessed data.
- Organized in lines of n memory words.

  e.g. 8192 lines $*$ 4 words $*$ 4 bytes.

Read access:

- Cache hit: word is read from cache.
- Cache miss: bus transfer takes place.

  1. Cache line read from memory.
  2. Replaces some (e.g. least recently used) line.

Write access:

- Problem: problem cache consistency.
- Processor modifies content of cache line.
- Other processors hold invalid copies.

Snoopy Bus Protocols

Cache consistency by bus watching mechanism

- Each cache watches bus for write operations
- Write through: forward write immediately to memory
  - Write-invalidate: cache copies are invalidated.
  - Write-update: cache copies are updated.
  - Disadvantage: high bus traffic.
- Copy back: write modified cache line only back if
  1. Line needs to be replaced, or
  2. Another processor reads data.

See additional slides.

Distributed Memory Architectures

- Multicomputers,
- Loosely coupled architectures,
- Message passing architectures,
- MPP Massively parallel processors.

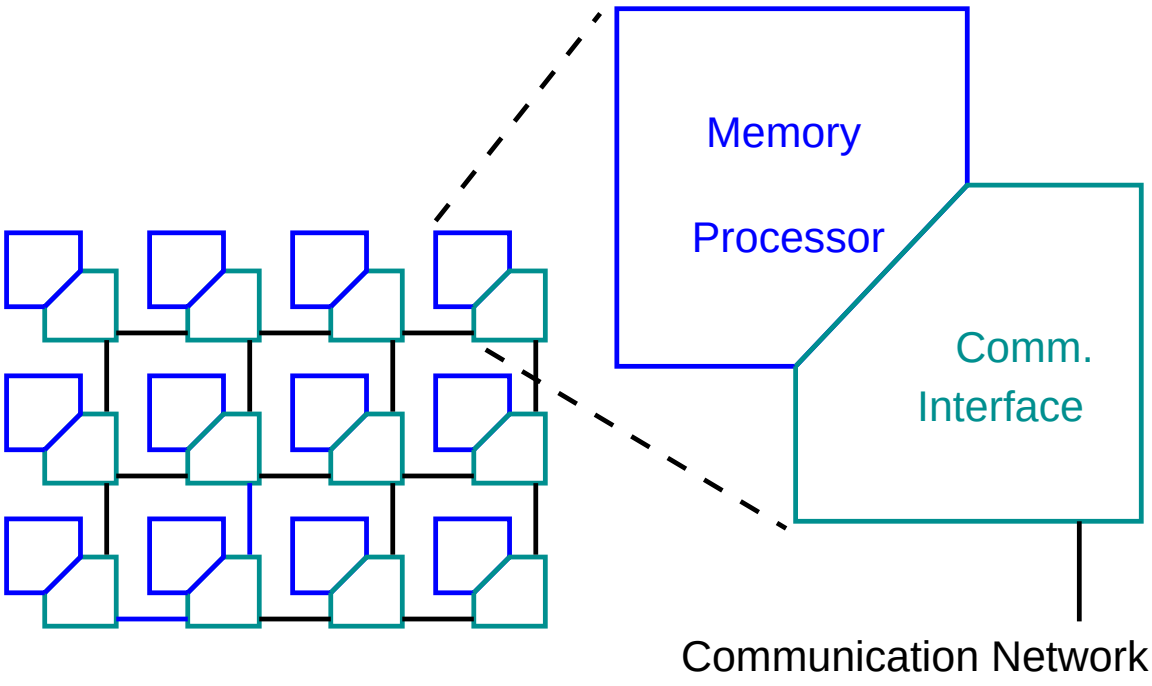

- Local memories private to processors,
- Communication via network.

System Interconnect Architectures

Characteristics

- Degree

  Maximum number of links per node.

- Diameter

  Maximum distance between two nodes

- Bisection width

  Minimum number of edges to be removed to divide network into halves.

- Edge length

  Physical representation in 2D/3D.

Network topology.

System Interconnect Architectures

Performance

- Functionality

  Data routing, global synchronization, ...
- Latency

  Maximum time delay for message transfer.
- Bandwidth

  Maximum data transfer rate.
- Hardware complexity.

  Implementation costs.
- Scalability

  Expandable with more nodes.

Network technology.

Static Connection Networks

See additional Figures.

- k-ary n-cube networks: dimension k, $N = n^k$ nodes.
  - Line/ring (1D), mesh/torus (2D), cube (3D).
  - Degree: $2^k = \text{const}$.
  - Diameter: $k(n-1) = k\sqrt[k]{N}$.
  - Bisection width: n^{k-1}.
  - Edge length: constant.
- Hypercube: binary n-cube, $N = 2^n$.
  - Degree: n *(not constant).*
  - Diameter: $n = \log_2 N$.
  - Bisection width: $2k^{k-1}$.
  - Edge length: not constant.
  - Scalability only by factors of 2!

Today, low-dimensional (2D, 3D) meshes are preferred over hypercubes.

Static Connection Networks

- Trees: degree k, depth d, $N = k^d - 1$ nodes.
  - Degree: $K + 1$.
  - Diameter: $2(d-1) \approx 2 \log_k N$
  - Bisection width: 1! *(hot spot!)*
  - Edge length: not constant.
  - Root becomes communication bottleneck.
- Fat Tree
  - "Branches get thicker toward the root".
  - Channel width increases toward the root.
  - Communication bottleneck is avoided.
  - Example: Connection Machine CM-5.

Fat trees are also popular topologies.

Dynamic Connection Networks

Switches provide dynamic connectivity.

- Digital buses.

  Only one transaction at a time.

- Crossbar network.

  All-to-all connection by n^2 switches.

- Multistage networks.

  Multible levels of $a \times b$ switches; several kinds of inter-stage connection patterns possible.

- Multi-dimensional networks.

  Multiple dimensions of buses or crossbar switches.

Communication patterns implemented on demand.

Message Formats

Message = variable-sized block of data = ⟨destination, address, size⟩.

- Messages = application-level concept.

  Computation processor inititates communication by passing necessary information to communication processor.

- Packages = system-level concept.

  Communication processor splits message into fixed size packages and transfers them over the network.

- Flits (flow control digits).

  Packages are split into fixed-length flits that may flow through the network in pipelined fashion.

Message Routing Schemes

- Store-and-Forward Routing
  - Packages are units of information flow.
  - Each package is stored in intermediate node before it is forwarded to the next node.
  - Latency proportional to communication distance.
- Wormhole Routing
  - Flits are units of information flow.
  - Each received flit of a package is forwarded to next node before the next flit is received.
  - Latency almost independent of communication distance.

Modern multicomputers use wormhole routing!

Wormhole Routing

Transmit a 100 byte message through a 100 MB/s network.

- 1 μs transmission time.
- Store and forward:
  - 1 nodes: 1 μs.
  - 8 nodes: 8 μs.
- Wormhole: 4 bytes/flit.
  - 1 node: 1 μs.
  - 8 node: $1 + \frac{7}{25} = 1.28$ μs.

Distance of communication does not matter any more!

Routing Strategies Characterization

- Deterministic routing.

  Communication path solely determined by source and destination address.

- Adaptive routing.

  Communication path depends on network conditions; alternate paths are possible.

- Non-deterministic routing.

  Random path is selected to avoid hot spots.

Minimization of communication distance is less important than avoiding hot spots in the network!

Dimension-order Routing

- Deterministic strategy for multidimensional networks.
- Choose specific dimension order and route along pathes in that order.
- Meshes: X-Y routing.

  First route along X-dimension, then route along Y-dimension.
- Hypercubes: E routing.

  Route along dimensions in the order of different bit positions of source and destination address.

Dead-lock free routing schemes.

Broadcasting and Multicasting

Communication Patterns

- Broadcast pattern

  One-to-all communication; one source sends the same message to all destinations.

- Multicast pattern

  One-to-many communication; one source sends same message to multiple destinations.

Avoid multiple unicasts by utilizing multicast routing schemes!

Computer Clusters

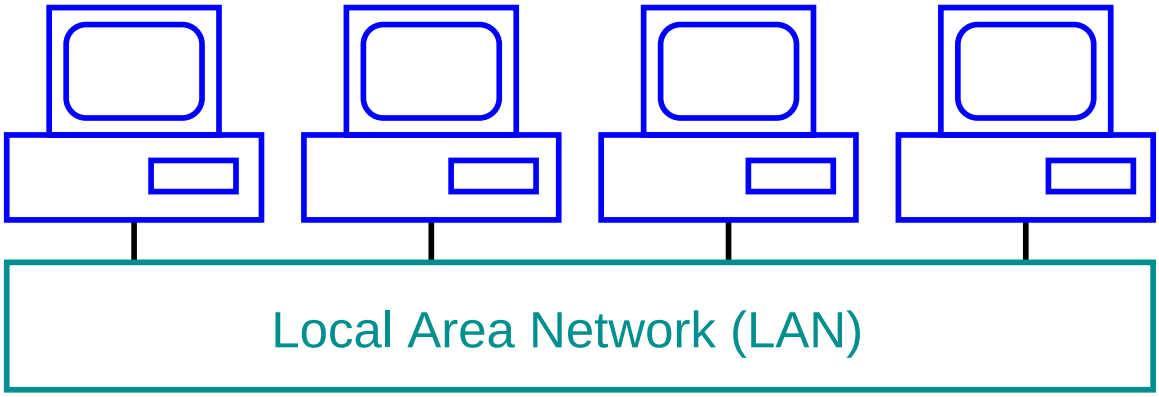

Multiple computers connected via network.

- Nodes are conventional workstations.
- Network is conventional LAN.
- Advantage: low costs.
- Disadvantage: high latencies.
- PVM: Parallel Virtual Machine

  Message passing package for clusters.

Basic operation like MPP; fault-tolerance must be considered.

Distributed Shared Memory

- Logically, a global address space is provided.
- Physically, the memory is distributed among the processors.
- Remote memory accesses via a scalable network.
- NUMA architectures
  - Non-Uniform Memory Access
  - Local memory can be accessed faster.
- COMA architectures
  - Cache Only Memory Achitecture.
  - All memory transformed into caches.

Scalable shared memory multiprocessors.

NUMA Architectures

- Each processor has memory and cache.
- All memories form a global address space.
- *Directory-based* protocols for cache consistency (no snooping possible).
  - Each cache copy is registered in directory.
  - If all cache copies of a data item are shared, memory requests yield new copies.
  - If cache copy is owned, memory request is forwarded to owner.
- Example: Stanford DASH multiprocessor

  Directory Architecture for Shared Memory.

Problem: Number of cache copies may grow arbitrarily.

COMA Architectures

- All distributed memories are transformed into caches.
- No memory hierarchy at processor node.
- All caches form a global address space.
- Example: KSR-1 computer with ALL-CACHE memory.
  - Each address becomes a name without direct physical relevance.
  - Processors form a two-level ring structure; a directory scheme associates names to processors.
  - A search engine transfers memory cells to processors as needed.
  - Prefetch mechanisms are supported to get required data ahead of time.

Problem: Processor are blocked during wait for memory cell.

Multithreaded Architectures

- Hardware-support for many fine-grained processes (threads).
  - Multiple program counters.
  - Multiple register files.
- Very efficient context switch.
  - Possibly after each instruction.
- Memory latency can be hidden.
  - Thread blocked on read of remote memory cell.
  - Other threads executed in the meantime.
- Sufficient parallelism required:
  - Instruction-level parallelism.
  - Application-level parallelism.
- Example: Tera MTA (Multi-Threaded Architecture)

Future of processor architectures.

Commercial Architectures

- IBM SP http://www.ibm.com
  - Message passing architecture based on cluster concept.
  - Paper
- Cray T3E http://www.cray.com
  - Distributed shared memory architecture without cache coherency.
  - Paper
- SGI Origin http://www.sgi.com
  - Distributed shared memory architecture with cache coherency.
  - Paper
- Tera MTA http://www.tera.com
  - Multi-thraded architecture.
  - Paper