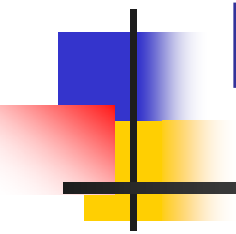

History of Distributed Systems

Joseph Cordina
joseph.cordina@um.edu.mt

Motivation

- Computation demands were always higher than technological status quo
- Obvious answer
  - Several computing elements working in harmony to solve a single problem
  - The need to have shared resources between the computing elements.
- Simplifying computing as 'instruction operating on data', all computing architectures fall under one of the following titles:

SISD	SIMD
MISD	MIMD

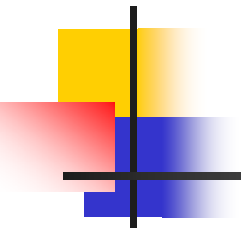

Interconnection Networks

- The computing elements need to communicate using interconnection networks.
- Criteria for such networks are:
  - Latency
  - Bandwidth
  - Connectivity
  - Hardware Cost
  - Reliability
  - Functionality
- End result is *speedup*
- Topologies used include:
  - Linear, Ring, Star, Tree, Nearest-Neighbor Mesh, 3D structures

Static vs. Dynamic Networks

- Static Networks came first and are fixed in topology
- Dynamic Networks contain switches (switching networks)
- Evolution from Static to Dynamic Networks:

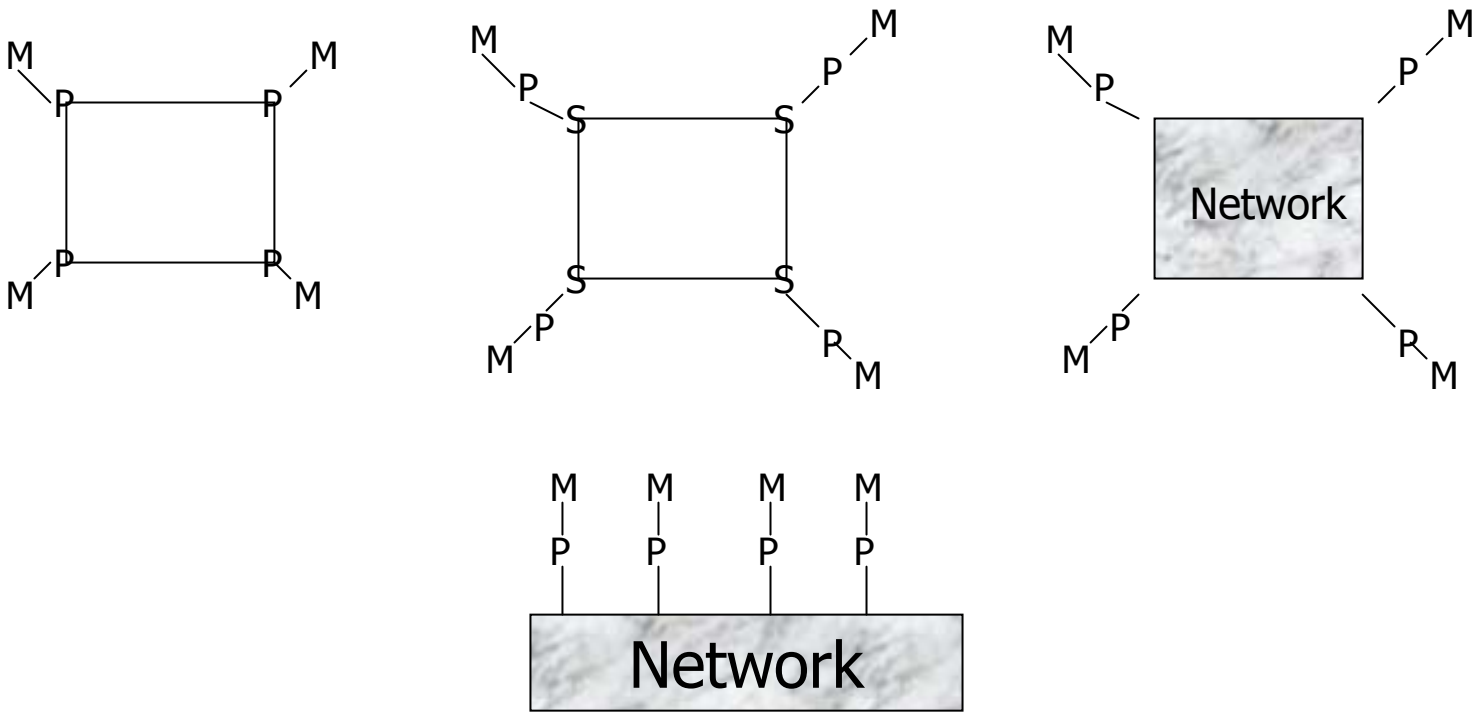

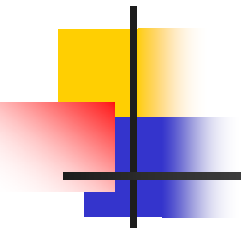

SIMD Architectures

- Originally it was observed that if we have

  For I:=1 to 100

  A(I)=B(I)*C(I);

  we are merely repeating the same instruction on different data
- Thus vector processors were created
- Vector processors take one single *vector instruction* that can simultaneously operate on a series of data arranged in array format.
- First successful computer was the CRAY-1 architecture (1976) that used 11-stage pipelines that could execute concurrently on data
- Other follow-ups include CDC Cyber 205 (4 identical FP units, IBM 3090, ILLIAC IV, GF11 and the Connection Machine

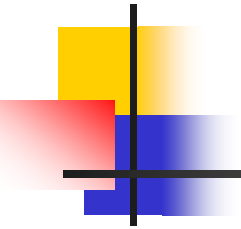

Pros and Cons of SIMD

- Cost is prohibitively expensive
- Special compilers need to be used
- Programming has to be done using specific languages
- Speedup is effective only for specific kind of problems.
- Very difficult to make use of total computing power all the time.
- For specific problems, results can be obtained very quickly.

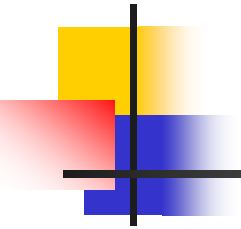

MIMD architectures

- To reduce price one could connect multiple processors, each marching on its own drum beat.
- Most real-life problems can be split up into a large number of small problems that can be solved individually.
- Solution can be gathered through communication between the different computing elements.
  - Connection topology now becomes very effective on speedup
  - Specific connection protocols were used: system bus, Myrinet, etc.
- For most problems, still requires some form of shared resource

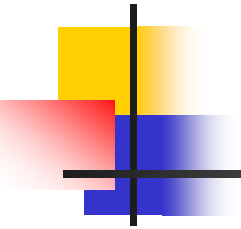

Private Memory

- This is a model where different computing elements with private memory communicate through some form of connection network
- Communication occurs through message passing
- The earliest were clusters of workstations (farms) communicating through a LAN.
  - Slow communication
  - Infrequent communication leads to good speedup
  - Scalability is inherent in LAN architecture
  - Very cheap commodities
- The first real private memory distributed computer was the Cosmic Cube with 64 computing nodes, each node having a direct, point-to-point connection to six others like it.
- Later developments were further hypercubes, meshes and data flow machines.

Shared Memory

- Whenever some data needs to be very frequently accessed, a common alternative is shared memory architectures.
- These are still MIMD architectures, yet the processing elements share the same memory bus.
- These are known as **S**ymmetric **M**ulti-**P**rocessor computers
- Nowadays these are readily available at a fraction of the cost of parallel architectures.

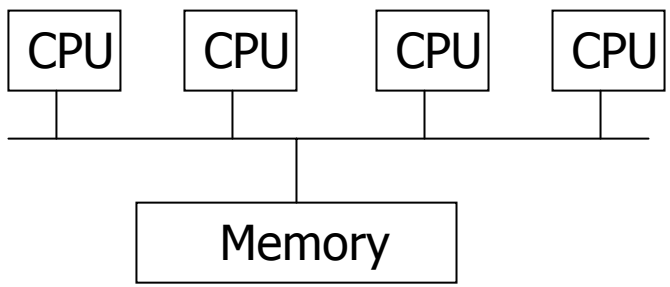

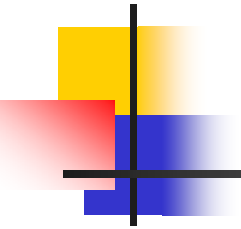

MIMD Characteristics

- Pros:
  - Shareability
  - Expandibility
  - Local Autonomy
  - Improved Performance
  - Improved reliability and availibility
  - Potential cost reductions
- Cons:
  - Network reliance
  - Complexities
  - Security
  - Specific OS features required

Current MIMD

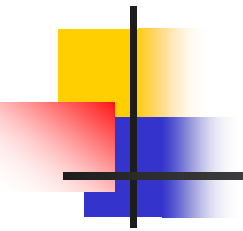

- Beowulf Clusters started in 1993 using commodity platforms and traditional Ethernet.
- LANs run at Gigabit speeds (10Gb expected soon)
- SMP machines are relatively cheap
- Millions of MIMD elements exist connected together for shared resources
  - Internet
- Computing can also be performed on internet connected elements
  - Example stories: SETI
- Two main types of distributed systems:
  - Client-Server (Farming)
  - Peer-to-Peer networks
- Large number of APIs available (PVM, MPI, RMI,…)

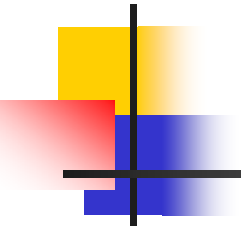

Future

- High-speed networking with external motherboard to motherboard buses (NGIO)
- Super broadband connectivity on the internet
- Huge farms of devices (GRID) offering resources including computing resources.
- High integration of devices at varying levels of computing power with high-bandwidth cross-talk.
- Massively parallel computing elements through the advent of quantum computing.
- Multi-processor machines becoming the norm.
- The development of new computing and programming tools to allow parallel and distributed computation.