Networks

- Networking has become ubiquitous (cf. WWW)
- Two major types of networks depending on geographical distance
  - *Local Area Networks* (LAN)
    - Connect 100's of computers within a maximum distance of a few kilometers
    - Often limited to a single building
    - LANs can be connected together via bridges (also called switched networks)
  - *Wide Area Networks* (WAN or long haul networks)
    - Connect 1000's of computers all over the world

Network architectural issues

- How are "nodes" connected?
- What is a message?
  - Length
  - How are messages routed from one node to another etc.
- What is the "quality of service" of the network?

Some sample network topologies

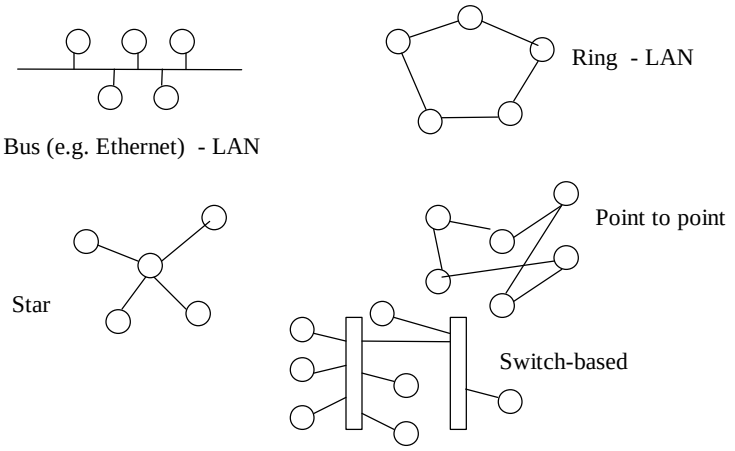

Typical network bandwidths

- Ethernet (LAN)
  - originally 10 Mbit/sec; now 100 M/bit second. Soon Gigabit Ethernet?
- Internet (successor of Arpanet - WAN)
  - Point to point originally with special processors (IMP: Interface message Processor); Currently copper co-axial cables yield bandwidths of 100 Mbits/sec. Fiber optics yield 1 Gbit/sec
  - Internet II is (almost) here
- ATM (Asynchronous transfer modes for switched networks; has been used in WANs and LANs)
  - !50 Mbit/sec to 2.5 Gbit/sec but high latency

Message Format

- Header
  - Request type; where it has to go; sometimes where it comes from; sequence number etc.
- Payload
  - Data (from a few bytes to Kbytes)
- Checksum
  - For checking the data received is correct
- Long messages can be split into *packets*
- Sending/receiving messages require complex protocols with several layers (TCP/IP, ISO standard etc. cf CSE/EE 461)

A little more on LAN's

- Ethernet
  - broadcasts messages; each node recognizes his own address
  - uses a "collision-detection" with back-off method
    - If 2 messages sent at the same time, garbage on the bus. Both senders recognize there is a garbage and resend
    - Since resend could be in sync, senders increase time between resends if there are lots of collisions.
- Rings
  - Allows more than one message on the ring at a given time
- ATM and "special" switches (Myrinet) are becoming popular

A little more on WAN's

- Point to point topology
- Store-and-forward routing
  - virtual circuits (telephone model). Connection is set-up for the duration of the message
  - packet switching
    - long messages split into packets
    - packets do not have to follow the same route. Can arrive in different order they were sent. Protocols reassemble the message
    - More complex hardware for routing , fault recovery etc. but better performance because can avoid congestion

Parallel processors

- Two basic paradigms
  - Shared-memory multiprocessors (communicate through "load-store" instructions)
  - Message-passing (communicate through "send-receive" primitives
- Current trend
  - SMP's (small number of processors) are shared-memory
  - Clusters of SMP's connected either mode. Depends often on the type of *interconnection network*

Interconnection networks for parallel processors

- Shared-bus
  - Least amount of concurrency
  - Good only until the bus gets saturated (used to be 32 processors; now more like 4 or 8)
  - Cache coherence well understood
- Cross-bar
  - Greatest amount of concurrency
  - Expensive in number of switches (grows like $O(n^2)$)
- Multi-stage interconnection networks
  - In between in terms of concurrency and complexity (O(nlogn))