LogP: Towards a Realistic Model of Parallel Computation

David Culler et al

Presentation By: Shoaib Kamil

Motivation

- Low cost of commodity workstations/ processors
- High cost of processor development
- Convergence of parallel machine designs
- Previous models too simplified

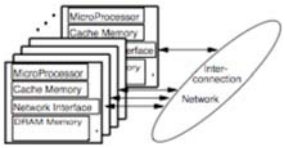

Assumptions

- Algorithms will require large numbers of data elements per processor
- High latency and overhead of communication, and limited bandwidth
- Ignore routing topology
- Cannot be tied to a particular programming model

LogP Model

- L: upper bound on *latency* for communicating a small message from source to destination
- o: *overhead*, the time the processor is engaged in processing the message
- g: *gap*, the minimum interval between **consecutive** message transmissions/ receptions
- P: the number of *processors*
- The network has a capacity of $\lceil L/g \rceil$ messages

Notes on Parameters

- L is an upper bound, so model assumes all messages incur the worst-case latency
- Large logical messages must be divided into a series of small actual messages
- o and g can overlap

Example: Broadcast

- Send a message from processor 0 to all other processors

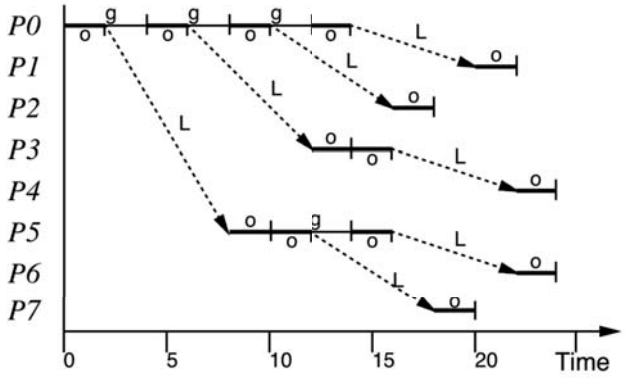

Matching to Real Machines: Does Topology Matter?

Network	Average Distance, P=1k	Average Distance, P=100k
Hypercube	5	8
Butterfly	10	17
Fat tree	9	33
3d torus	7.5	35
3d mesh	10	47
2d torus	16	160
2d mesh	21	213

But send/receive overhead at each hop is more expensive than routing delays…

Comparison to Other Models

- PRAM
  - Unlimited processors & memory with unit access time
  - Synchronous
  - In each cycle, can read any memory, do computation locally, and write any memory
  - Akin to g=o=L=0
- BSP
  - Bounded P, each looks like LogP processor
  - Synchronous
  - In each superstep:
    - Send *h* messages
    - Receive *h* messages
    - Do local computation
- LogP is closest to real machines

Limitations

- Large messages
  - Most NI do not require breaking up messages into smaller ones and involving processor at each
  - LogGP better captures real-world large message performance
- Special barrier networks
- Overlap
  - Gap value can change if multiple communications are queued back-to-back
    - See Welcome's GAS evaluation @ http://upc.lbl.gov