

Lecture 11

Midterm review

Announcements

- Midterm: Monday 10/30, 7:00 to 8:20 PM
- Room: SSB 106
- Open book, but no notes
- No lecture on 11/14
 - But note time changes
 - Makeup on 11/17 9:30AM to 10:50 AM
 - Section will be held on Thurs 11/16 in our class room at usual meeting time

Terms and concepts

- Know the definition and significance of
- Parallel speedup and efficiency, super-linear speedup, scaled speedup,
- Scalability, cost, cost-efficient, isoefficiency
- Amdahl's law, serial bottlenecks
- Granularity, surface-to-volume effect
- Loop carried dependence

Terms and concepts

- CRCW, CREW PRAM
- SPMD, MIMD, SIMD
- Multiprocessors and multicomputers; shared memory, message passing
- Interconnection networks: hypercube, ring, mesh, k-ary d-cube, broadcast and reduction algorithms; diameter and bisection bandwidth
- Message startup, half power point $n_{1/2}$, peak bandwidth
- Processor geometry and blocked data decompositions

Implementation techniques

- Message passing with MPI
 - asynchronous, blocking and non-blocking communication
 - MPI sending modes (buffered, standard, etc)
 - collective communication: reduction, broadcast
 - communicators
 - Eager vs. rendezvous communication
- Performance modeling and performance measurement techniques

Algorithms

- Know the purpose of the following algorithms, and the significant implementation issues affecting performance.
- Be familiar with performance models for each and be prepared to analyze performance and scalability
- **Sorting**
 - Enumeration sort, odd-even sort, odd-even transposition sort, bitonic sort, bucket sort, sample sort
- **Stencil methods**
 - **ODE solver:** 1D case
 - Multidimensional cases : 2D and 3D
 - “Curse of dimensionality:” surface to volume ratios
 - Convergence and how to check it
 - Ghost cells, partitioning, performance models
 - Dealing with Loop carried dependences
- **Matrix multiplication**
 - Canon’s algorithm, SUMMA
- **LU Decomposition (Gaussian elimination)**
 - Cyclic decompositions

Practice problems

- You have 2 machines
 - UMA.net—a Uniform Memory Access (UMA) multiprocessor with 64 processors and 32 GB of shared memory
 - Distributed.net loose connection of 145,000 processors with 9 Terabytes of RAM tied together over the internet with 5Mbit/s network connections with an average latency of 20ms between processors
- Match to the two workloads
 - Analyzes Japanese Haiku poetry. The algorithm almost 100% serial—one can only find parallelism between instructions issued a few cycles apart from each other. Each poem takes 30 seconds on a single processor to compute. Enormous amount of haiku to process, and each poem can be processed separately from the other.
 - LU decomposition on 70,000 x 70,000 matrices

Networks

- Diameter, bisection bandwidth
- Map a 1D ring onto a hypercube, mesh, k-ary d-cube
- How many parallel/unique paths between any two nodes in a hypercube?
- Sort on a 2D network
- Sort on a hypercube

