Session 5 - Introductory Course in Parallel and Distributed Computing

Yves Robert & Behrooz Shirazi

TCPP Workshop
February 2010

Outline

Outline

Approaches/High-level Topics

- Global overview vs in-depth coverage of limited material
- Practical approach vs theory-oriented approach (applies to everything: algos/prog/distributed)
- Algorithms vs programming
- Parallel vs distributed computing
- Fine-grain (threads) vs coarse-grain (MPI) parallelism
- Parallel architectures (shared-memory, distributed memory, multicores) and their impact
- Applications and tools (CS vs non-CS major)

Outline

Overview oriented course

- Introduction to parallel computing
- Multi-core processors
- POSIX and Pthreads programming
- Cache memory and Cache Coherence in MP
- Interconnect Networks for MP
- Multi-computers
- MPI and MPI programming
- Current supercomputers
- Algorithms and programs

Practice-oriented course

- Introduction to parallel processing
- Cray XMT architecture
- Cray XMT programming model
- Cray XMT programming tools
- Reasoning about performance on the Cray XMT
- Data parallelism
- Task parallelism
- Using and optimizing recursive algorithms
- Dataflow
- Graph algorithms
- Other cool algorithms

Theory-oriented course

- Sorting networks
- P-RAM: EREW/CREW/CRCW, simulation, Brent's theorem
- Cole's sorting machine
- Pointer jumping, recursive doubling, list ranking, prefix sums
- Matrix operations
- Connected components, sorted paths
- Graph algorithms

Distributed computing

- Models, broadcast
- Leader election in a ring, in a general graph
- The consensus problem
- Mutual exclusion,
- Atomic snapshots of shared memory
- Peer-to-peer computing
- DHT and join/search algorithms

Parallel architectures

- Shared-memory computers
- SIMD, SPMD and BSP algorithms
- Cache and memory hierarchy, Level3 BLAS
- Multicores and GPUs
- Distributed-memory computers
- Macro-communications (broadcast, scatter, all-to-all, etc)
- Static interconnection networks (rings, grids, hypercubes, fat trees, etc)
- Dynamic interconnection networks (switches, Benes, Clos, etc)
- Routing, store-and-forward, wormhole, bandwidth sharing

Parallel algorithms

- Models :Sorting networks
- Models: P-RAMS
- Interconnection networks
- Case study: ring and hypercube
- Peer-to-peer computing, distributed hash tables
- Algorithms on rings (matrices, stencils)
- Algorithms on 2D grids
- Load balancing for heterogeneous platforms
- Task graph scheduling
- Loop parallelization

Scheduling and automatic parallelization

- Task graphs
- Makespan: polynomial/NP-complete instances
- List scheduling, Graham's bound
- With communications and heterogeneous processors: HEFT
- Realistic models (one-port, bounded multi-port)
- Divisible load scheduling
- Steady-state scheduling
- Multi-criteria (throughput, latency, reliability, energy, cost)
- Dependence levels and Kennedy-Allen algorithm
- Dependence vectors and Lamport hyperplane method
- Unimodular matrices and space-time transformations