


Parallel Programming

Kai Shen

1/24/2011 CSC 258/458 - Spring 2011 1




Parallel Programming Steps

Converting a sequential application to a parallel one

- Decomposition into tasks
- Assign tasks to processors
- Orchestrate data access and synchronization


1/24/2011 CSC 258/458 - Spring 2011 2



Dependences

- Two tasks are dependent if they must follow their order in the sequential program so that the program execution isn't changed.
- Three types of data dependences:
 - Read-after-write
 - Write-after-read
 - Write-after-write
- Only "read-after-write" is called true dependence


1/24/2011 CSC 258/458 - Spring 2011 3



Task Decomposition

- Decomposition:
 - over different functions
 - over different data segments (over loop iterations)
- How good is the decomposition?
 - Assume dependent tasks must run serially. We can build a DAG of task dependencies and critical path length indicates lower bound of parallel execution time.
- Tradeoff on task granularity:
 - smaller tasks may offer more parallelism/concurrency
 - smaller tasks require more management/programming overhead

1/24/2011 CSC 258/458 - Spring 2011 4



Task Assignment


Goals:

- Load balance
- Minimize inter-processor communication/synchronization
⇒ maximize locality

Ways:

- Static assignment
- Dynamic assignment


1/24/2011 CSC 258/458 - Spring 2011 5



Dynamic Task Assignment

- How does it work?
 - Maintain a centralized queue of ready tasks, protected by mutex lock
 - Each thread grabs a task at the beginning; grabs another task after completing the current one
 - New tasks may be generated on the fly and added to queue
- Advantage: good load balancing
- Disadvantages with dynamic task assignment
 - Locality may be lost in the interests of load balancing
 - Synchronization contention on manipulating the task queue


1/24/2011 CSC 258/458 - Spring 2011 6



Dynamic Task Assignment

- Disadvantages with dynamic task assignment
 - Locality may be lost in the interests of load balancing
 - Synchronization contention on manipulating the task queue
- Fixable through distributed queues with work stealing
 - Each thread has an exclusive task queue with good locality and no contention
 - When out of work (load imbalance or synchronization), steal some task from another queue with ready tasks

1/24/2011 CSC 258/458 - Spring 2011 7



Orchestration

- Access shared data
 - Shared-memory parallel platform
 - Distributed-memory parallel platform
- Synchronization
 - mutex lock, condition variables, barrier, ...
 - Enforce dependences
- Fine-grained synchronization between tasks
 - Require tasks to run simultaneously
 - Work better with static task assignment

1/24/2011 CSC 258/458 - Spring 2011 8

Parallel Programming Example: Successive Over Relaxation

- SOR implements a mathematical model for many natural phenomena, e.g., heat dissipation, ocean currents
- Given a 2D grid of data, for some number of iterations:
 - For each internal grid point, compute average of its four neighbors

```

for (i=1; i<n; i++)
  for (j=1; j<n; j++)
    temp[i][j] = 0.25 * (grid[i-1][j]+grid[i+1][j]+grid[i][j-1]+grid[i][j+1]);
for (i=1; i<n; i++)
  for (j=1; j<n; j++)
    grid[i][j] = temp[i][j];
    
```

1/24/2011 CSC 258/458 - Spring 2011 9

Parallel Programming Example: Successive Over Relaxation

```

for (i=1; i<n; i++)
  for (j=1; j<n; j++)
    temp[i][j] = 0.25 * (grid[i-1][j]+grid[i+1][j]+grid[i][j-1]+grid[i][j+1]);
for (i=1; i<n; i++)
  for (j=1; j<n; j++)
    grid[i][j] = temp[i][j];
    
```

- Dependences:
 - First (i,j) loop nest?
 - Second (i,j) loop nest?
 - Between the two loop nests?
 - Between two iterations?

1/24/2011 CSC 258/458 - Spring 2011 10

Parallel Programming Example: Successive Over Relaxation

```

for (i=1; i<n; i++)
  for (j=1; j<n; j++)
    temp[i][j] = 0.25 * (grid[i-1][j]+grid[i+1][j]+grid[i][j-1]+grid[i][j+1]);
for (i=1; i<n; i++)
  for (j=1; j<n; j++)
    grid[i][j] = temp[i][j];
    
```

- Task decomposition:
 - 1D partitioning: each task manages some columns or rows
 - 2D partitioning: each task manages a 2D block of the grid
- Impact on communications in distributed memory parallel computing?

1/24/2011 CSC 258/458 - Spring 2011 11

Parallel Programming Example: Gaussian Elimination

- Solving a system of linear equations
- Reduce an equation matrix into an equivalent upper-diagonal matrix

Diagram illustrating the reduction of matrix A into an upper triangular matrix R. The matrix A is shown as a grid of bars representing elements. The matrix R is shown as a grid of bars representing elements, with a diagonal of bars and zeros below. The transformation is shown as $p \cdot A_1 + A_2$, X , and R with elements r_1 and $r_2 + p \cdot r_1$.

- Partial pivoting to maintain numerical stability

1/24/2011 CSC 258/458 - Spring 2011 12

Assignment #1

Shared Memory Parallel Programming

- All should have accounts in CS research network
 - Familiarize with the parallel machines
- Pre-assignment
 - We provide you sequential/parallel SOR code (a different version, called red-black ordering)
 - You read, understand, and run it
- Main assignment
 - Parallel Gaussian Elimination
 - Different settings (input matrices, machines, proc #'s)
 - Analysis and comparison
 - Grading based on your written report!

1/24/2011

CSC 258/458 - Spring 2011

13

Pthreads Synchronization Primitives

- Mutex lock (mutual exclusion)
 - `pthread_mutex_lock`
 - `pthread_mutex_unlock`
- Condition variable (waiting for a condition)
 - `pthread_cond_wait`
 - `pthread_cond_signal`
 - `pthread_cond_broadcast`

1/24/2011

CSC 258/458 - Spring 2011

14

Barrier Synchronization

- A barrier is set for all threads
- A thread can proceed beyond its barrier if and only if all threads have reached respective barrier points
- Implement the barrier in pthreads?
 - Count the number of arrivals at the barrier
 - When a thread arrives, wait if this is not the last one, otherwise unlock everyone else and proceed
 - Mutex lock to protect the arrival number, condition variable to implement wait and broadcast

1/24/2011

CSC 258/458 - Spring 2011

15