

Parallel Programming

Kai Shen

1/28/2014

CSC 258/458 - Spring 2014

1

Gaussian Elimination: Dependencies

- Simplification – ignore pivoting and final solving step
- Reduce an equation matrix into an equivalent upper-diagonal

$$\begin{matrix} & A & X & R \\ p \cdot A_{:,1} + A_{:,2} & \begin{matrix} A_{:,1} \\ \vdots \\ A_{:,n} \end{matrix} & \begin{matrix} X \\ \vdots \\ X \end{matrix} & = \begin{matrix} r_1 \\ r_2 + p \cdot r_1 \\ \vdots \\ r_n \end{matrix} \end{matrix}$$

for i=1 to N
 for j=i+1 to N
 zero out $A_{i,j}$ by adding $p \cdot A_{:,i}$ to $A_{:,j}$

- Dependencies:
 - Outer loop instances (e.g., i=2 depends on i=1)?
 - Inner loop instances (e.g., j=3 depends on j=2)?
 - Within one inner loop instance?

1/28/2014

CSC 258/458 - Spring 2014

2

Gaussian Elimination: Task Decomposition

$$\begin{matrix} & A & X & R \\ p \cdot A_{:,1} + A_{:,2} & \begin{matrix} A_{:,1} \\ \vdots \\ A_{:,n} \end{matrix} & \begin{matrix} X \\ \vdots \\ X \end{matrix} & = \begin{matrix} r_1 \\ r_2 + p \cdot r_1 \\ \vdots \\ r_n \end{matrix} \end{matrix}$$

for i=1 to N
 for j=i+1 to N
 zero out $A_{i,j}$ by adding $p \cdot A_{:,i}$ to $A_{:,j}$

- Parallelism:
 - Inner loop instances (row-wise)
 - Within one inner loop instance (column-wise)
- Task decomposition:
 - Row vs. column partitioning?
 - 2-dimensional partitioning

1/28/2014

CSC 258/458 - Spring 2014

3


Gaussian Elimination: Task Decomposition

- Row partitioning at different granularities

Row by row,
cyclically



Every proc gets a
contiguous chunk



- Better load balancing under cyclic partitioning

1/28/2014

CSC 258/458 - Spring 2014

4

Pivoting

- Reduce an equation matrix into an equivalent upper-diagonal

$$\begin{matrix} A_{:,i} \\ A_{:,k} \end{matrix} \begin{matrix} A \\ X \end{matrix} = \begin{matrix} r_i \\ r_k \end{matrix}$$

for $i=1$ to N
 swap $A_{:,i}$ (r_i) with $A_{:,k}$ (r_k) where $a_{i,k}$ has largest absolute value
 for $j=i+1$ to N
 zero out $A_{i,j}$ by adding $p \cdot A_{:,i}$ to $A_{:,j}$

- How does it complicate the parallelization (particularly row vs. column partitioning)?

1/28/2014

CSC 258/458 - Spring 2014

5

Dynamic Task Assignment

- How does it work?
 - Maintain a centralized queue of ready tasks, protected by synchronization primitives like 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(s) and disadvantage(s)
- Does it help Gaussian Elimination with pivoting?

1/28/2014

CSC 258/458 - Spring 2014

6

Block Computation

- Matrix multiplication (sequential)
 - Large matrices ($n \times n$) – even a single row/column won't fit in cache

$$\begin{matrix} A \\ B \end{matrix} \begin{matrix} X \\ = \end{matrix} \begin{matrix} C \end{matrix}$$

- How many times each element of A,B has to be loaded?

- Block-by-block multiplication
 - Three blocks ($b \times b$) fit into the cache at the same time

$$\begin{matrix} A \\ B \end{matrix} \begin{matrix} X \\ = \end{matrix} \begin{matrix} C \end{matrix}$$

1/28/2014


CSC 258/458 - Spring 2014

7

Gaussian Elimination: Task Decomposition

- Row partitioning at different granularities

Row by row,
cyclically



Every proc gets a
contiguous chunk



Every proc gets a
block of rows



- A good block size to efficiently utilize processor cache, you produce reasonable load balancing

1/28/2014

CSC 258/458 - Spring 2014

8

Irregular Parallelism

- Real problems contain large, sparse matrices
- Solve them as dense matrices waste time on zero-element operations
- Sparse matrix computation
 - Load imbalance (now you can really benefit from dynamic task assignment)
 - Managing nonzero fillins (hard to do block operations and utilize cache effectively)


1/28/2014

CSC 258/458 - Spring 2014

9

Example Speedup Results

- The speed ratio over the best sequential run



1/28/2014

CSC 258/458 - Spring 2014

10

Alternative Methods for Solving Linear Equations

- Faster solutions if we tolerate less precision
 - Less-than-precise pivoting
 - Iterative method

1/28/2014

CSC 258/458 - Spring 2014

11

Assignment #2 Shared Memory Parallel Programming

- 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 (with pivoting, but without block computation or sparse computation)
 - Different settings (input matrices, machines, proc #'s)
 - Analysis and comparison (Row/column/2D partitioning, blocking sizes, static vs. dynamic task assignment)
 - Grading based on your written report!

1/28/2014

CSC 258/458 - Spring 2014

12

Assignment #2

- All should have accounts in CS grad/research network
- Familiarize with the parallel machines
 - Multi-chip/socket
 - Multi-core
 - Hardware threading (or Intel hyperthreads)

1/28/2014

CSC 258/458 - Spring 2014

13

Suggestions on Testing

- Incremental testing:
 - Try an artificial, small, dense input matrix first
 - Try to parallelize without partial pivoting first

1/28/2014

CSC 258/458 - Spring 2014

14

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

1/28/2014

CSC 258/458 - Spring 2014

15

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
- One way to implement the barrier:
 - 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

1/28/2014

CSC 258/458 - Spring 2014

16