ITCS 4145/5145 Parallel Programming
Final exam
Tuesday May 8th, 2012, 5:00 pm - 7:30pm

Name: .....

This is a closed book test. Do not refer to any materials except those supplied for the test.

Supplied: *"Appendix A Basic MPI Routines"*
"Summary of OpenMP 3.0 C/C++ Syntax."
"CUDA C Quick Reference"

Answer questions in space provided below questions. Use additional paper if necessary but make sure your name is on every sheet.

Total /60

Do not refer to any materials for this part

Qu. 1 Answer each of the following briefly:

(a) Why could modeling the motion of astronomical bodies be regarded as a "Grand Challenge" problem? 2

(b) Suppose a computation has a part that must be executed sequentially and a part consisting of s completely independent parallelizable sections. There are p processors and $p < s$. Derive the maximum speed up factor. A processor can execute more than one parallelizable section in sequence. f is the fraction of the code that cannot be parallelized. Clearly explain your derivation and any assumption 4

(c) What is the purpose of a MPI communicator?

2

(d) Which MPI routine is used to receive a message from an MPI_SSend() routine?

2

(e) How could you send each row of the array A[4][6] to different slave process using the MPI scatter routine (4 rows, 4 slave processes)? Note: C stores its arrays in row major order, that is, one row after another in memory.

2

(f) Why would one sometimes compute the same calculation in more than one process rather than compute it once and send its value to the other processes?

2

(g) Explain how to achieve a parallel computational time complexity of $O(n)$ for matrix multiplication with n x n matrices. 2

(h) What is the prefix sum calculation? 2

(i) A simple barrier implementation uses a counter. Explain how this implementation works. How can it be made re-entrant, that is, called multiple times in processes and still work, taking into account that a process might enter the barrier for a second time before previous processes have left the barrier for the first time. 4

(j) What is a critical section? 2

(k) One way to create threads in a Java program is to extend the class Thread, that is, define a class that is a derived class of the class Thread. What is the disadvantage of this method? 2

(l) Explain what the following OpenMP construct does:

```
#pragma omp for
for (i = 0; i < 10, i++) {
   ...                  //for loop body
}
```

What other construct is also necessary for the above to work? 4

(m) What is meant by the term cost-optimal parallel algorithm? 2

(n) Using Berstein's conditions for parallelism, establish whether:

```
int a[100];
for (i = 0; i < 4; i++) {
   a[i] = a[i + 2];
}
```

can be re-written as:

```
int a[100];
forall (i = 0; i < 4; i++) {
   a[i] = a[i + 2];
}
```

and still obtain the same results? 4

(o) Is following a Bitonic sequence:

4 7 9 8 7 4 5

Explain your answer. 2

(p) In CUDA, what type of C routines cannot be called from the device code (kernel)? 2

(q) In CUDA, what does the qualifier `__device__` indicate when used with a variable declaration? 2

(r) What does *memory coalescing* mean in CUDA programming? 2

Qu. 2 Write a *complete* C program that uses both MPI and OpenMP to add together all the numbers in a provided array A[1000]. Use 10 MPI processes and 100 OpenMP threads in each process.

Provide very clear explanation of how the program works, and comments in your code. If I do not understand the code, I will assume it is incorrect.

8

Qu. 3 Write a *complete* CUDA program to find how many negative numbers they are in a list of integers stored in an array that holds 100,000 integers. Use 1-D grid with 1-D blocks and 10 blocks each of 100 threads.

Provide very clear explanation of how the program works, and comments in your code. If I do not understand the code, I will assume it is incorrect.

8