

ITCS 4145/5145 Parallel Programming
Final exam
11:00 am - 1:30pm , Tuesday May 7th, 2013

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

Qu. 1 Answer each of the following briefly:

(a) Suppose 50% of a computation execution must be executed on a single processor but the remainder can be divided into ten parallel parts of equal duration each of which can be executed by a separate processor at the same time. What is the maximum speed up if you have four processors? Clearly explain. No points for simply putting down a numerical answer with an explanation.

2

(b) Why are critical sections not used in MPI programs?

2

(c) In Assignment 1, what was the command “diff” used for?

2

(d) In MPI, what is the term “rank” used for?

2

(e) What does the MPI code below do?

2

```
int main (int argc, char *argv[]) {
    int token, NP, myrank;
    MPI_Status status;
    MPI_Init (&argc, &argv);
    MPI_Comm_size(MPI_COMM_WORLD, &NP);
    MPI_Comm_rank(MPI_COMM_WORLD, &myrank);
    if (myrank != 0) {
        MPI_Recv(&token, 1, MPI_INT, myrank - 1, 0, MPI_COMM_WORLD, &status);
    } else {
        token = -1;
    }
    token += 2;
    MPI_Send(&token, 1, MPI_INT, (myrank + 1) % NP, 0, MPI_COMM_WORLD);
    if (myrank == 0) {
        MPI_Recv(&token, 1, MPI_INT, NP - 1, 0, MPI_COMM_WORLD, &status);
    }
    MPI_Finalize();
}
```

(f) What is meant by a blocking MPI routine?

2

(g) The MPI collective routines are said to have the same semantics as using point-point `MPI_send()` and `MPI_recv()` routines separately in respect to when they return. If so, when does the `MPI_Gather()` routine return? Explain fully.

2

(h) Describe the 3-message protocol for implementing `MPI_SSend()` synchronism?

2

(i) What is meant by re-entrant code? Why should a barrier routine be re-entrant?

2

(j) In a Jacobi iteration problem, why stopping the iterations when all the values have changed by less than ϵ from the previous iteration does not ensure the solution is accurate to $\pm\epsilon$?

2

(k) What is a thread in the context of shared memory programming?

2

(l) What is a mutually exclusive lock variable (“mutex”)?

2

(m) In OpenMP, how can an integer variable be declared that is accessible by all threads within a parallel region

2

(n) Why does false sharing reduce performance of shared memory programs? (This question is not asking what is false sharing.)

2

(o) Write an OpenMP code sequence using the sections directive to implement two threads where one thread adds one to a shared variable x and another thread sets a private variable y to the same value as x , not taking into account thread contention for the shared variable x . If x is initialized to 0, what are the possible values that x could have after the execution of the threads? Clearly explain your answer.

4

(p) How many processors are needed to achieve a parallel time complexity of $O(\log N)$ for matrix multiplication with $N \times N$ matrices? Explain.

2

(q) What is the parallel time complexity of Odd-Even Transposition Sort with N numbers and N processors?

2

What is the parallel time complexity with P processors where each processor handles N/P numbers in the compare and exchange operation?

2

(r) Which is CUDA? 2

(s) What is nvcc? 2

(t) What is meant by flattening an array in GPU programming? Why is it used? 2

Qu. 2 Write a MPI program that implements a workpool consisting of a master process and 5 slave processes. The master sends a task consisting of two random integers to be added together to each slave. The slaves execute the task by adding the two integers and returns the result to the master. The master then sends another task to the slave, again consisting of two random integers to be added together. This continues until 100 tasks are completed.

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.

Qu 3 Given a sequence of letters consisting of 10,000 groups of 4 characters where each character can be A, C, G or T, write a CUDA program to find it how many groups consist of the sequence GCAT.

The number of threads in a block and number of blocks in the grid are to be given as #define constants, T and B respectively.

Some CUDA predefined variables:

The thread ID within a block in the x direction is given by threadIdx.x

The block ID in the x direction is given by blockIdx.x

The number of bocks in the x direction is given by blockDim.x

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

