

Midterm 2 - ECE 563 Spring 2012

This is a take-home test. You must do your own work, if found cheating you will be failed in the course and you will be turned in to the Dean of Students. To make it easy not to cheat, this exam is open book, open notes, and you may use passive online resources to solve it. By passive I mean you may look at existing online resources, you may not post queries or requests to information to any online resource.

For simplicity, some of the loop nests shown will lead to out-of-bounds array accesses. Ignore these when doing your test.

Please print and sign your name below. By doing so you signify that you have not received or given any prohibited help on this exam.

Name:

Question 1 (14 pts): Consider the N by N arrays A and B shown below.

Each array has its rows distributed as indicated, i.e. array A is row distributed and array B is column distributed. The gray rows and columns contain non-zero values, and the white rows and columns contain zeros. In particular, column k of A is non-zero, and row k of B contains non-zero values, and all other rows and columns contain zero values.

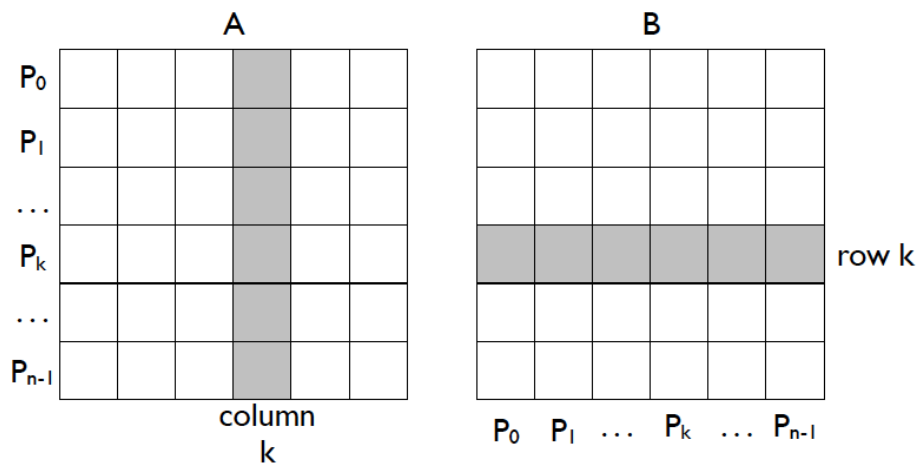
The computation performed on these arrays is to compute $Z = A \times B$, i.e. matrices A and B are multiplied and the result is stored into Z , and then the sum of all of the elements of Z are sent to each process. Z is distributed by rows, just like A .

$$Z = A \times B$$

```
MPI_AllReduce(Z(pid*N), sum, N*N/num_procs, MPI_FLOAT, MPI_PLUS, MPI_COMM_WORLD);
```

It is the intention of the AllReduce for each process to send the rows of Z that they own to the AllReduce.

- How many operations are performed to compute $Z = A \times B$? You can count $a*b+c$ as a single operation. You should not multiply the zero elements.
- Derive the isoefficiency function for the matrix multiply and AllReduce. Assume that A is distributed by rows, and B is distributed by columns, as shown in the figure below.



Question 2 (6 points)

```
for (i = 0; i < n; i++) {  
S1:  a[i] = a[i-2]  
S2:  b[i] = a[i+3]  
}
```

- a. What kind of dependences, if any, exists between the read and write of a in S1? What is its direction and distance?
- b. What kind of dependences, if any, exists between the write of a in S1 and the read of a in S2? What is its direction and distance?

Question 4 (14 points): Consider the loops below:

```
DO I = 1, N
  A(I) = B(I-1) + C(I+1)
END DO
```

```
DO I = 1, N
  D(I) = A(I-1) + B(I-2)
END DO
```

- a. Give statements that distribute A, B, C, and D to minimize communication. Use at least one TEMPLATE and at least one ALIGN statement. The arrays can be distributed onto the default 1-D processor grid.
- b. Is any communication necessary after performing your distribution? If so, say what kind is needed and where it is located (you can either say in words or specify what the MPI commands are.)

Question 5 (14 points). Consider the loops below:

```
float a[n], b[n], c[n+1]
for (i = 1; i < n-1; i++) {
    a[i] = b[i] + c[i+1]
}
```

```
float *pa = (a-1);
for (i = 1; i < n-1; i++) {
    *pa++ = a[i-1] + b[i];
}
```

- a.** Convert the program above to a UPC program, where a , b and c are in shared memory. To the extent possible make pa point to the section of the arrays that is in the threads shared memory.
- b.** Is any communication necessary in your program? If there is, describe the communication.

Question 6 (14 points). Consider the loops below:

```
DO I = 1, N  
  A(I) = B(I-1) + C(I+1)  
END DO
```

```
DO I = 1, N  
  D(I) = A(I-1) + B(I-2)  
END DO
```

Write this as a Co-Array Fortran program.

Question 7 (14 points). Consider the C program below for this question.

```
Before either thread executes
for (i=0; i < n; i++) {
    a[i] = 0;
    flag[i] = 0;
}
```

```
Thread 0
for (i=0; i < n; i++) {
    a[i] = i;
    flag[i] = 1;
}
```

```
Thread 1
for (i=0; i < n; i++) {
    while (!flag[i]);
    printf("%d\n", a[i]);
}
```

- a. Describe the output if this program only has sequentially consistent executions. If multiple outputs can exist (and they may, or may not) describe these multiple outputs. You can do this verbally.
- b. Describe at least one output that can occur if the program does not have sequentially consistent executions. If it is possible for this program to have a different output in a non-sequentially consistent execution than in a sequentially consistent execution, the output you give should be different than what you gave for a.

Question 8 (14 points). A programmer compiles a UPC program in two ways:

1. The program is compiled to execute using a strict consistency model (SC)
2. The program is compiled to execute using a weak (relaxed) consistency model that gives SC results for race free programs.

The programmer executes the program such that all possible outcomes are generated, and the results from the two versions are identical. Circle all that might be true.

- a. The program is sequential
- b. The program is race free (i.e. all shared variable accesses are synchronized)
- c. The program must have races (i.e. there are shared variable accesses that are not synchronized)
- d. The program *must* be race free (i.e. all shared variable accesses are synchronized), since different access orders on variables involved in races always affect the outcome of a program.
- e. The hardware is sequentially consistent, and the compiler options are irrelevant in enforcing sequential consistency.
- f. The hardware is sequentially consistent and both versions the compiler performed no optimizations that might lead to a sequentially inconsistent execution.