

FACULTY OF SCIENCE

ACADEMY OF COMPUTER SCIENCE AND SOFTWARE ENGINEERING

MODULE	IT18X97/IT00197 PARALLEL PROGRAMMING
CAMPUS	AUCKLAND PARK CAMPUS (APK)
EXAM	EXAMINATION SSA OCTOBER 2021

DATE: 28 October 2021

SESSION: 08:30 - 10:30

ASSESSOR(S):

MR R. MALULEKA

EXTERNAL MODERATOR:

DR A. EZUGWU (UKZN)

DURATION: 120 MINUTES

MARKS: 100

Please read the following instructions carefully:

1. An **additional 30 minutes** submission time will be allowed.
 2. Answer all questions.
 3. Answer each question in its entirety before moving on to the next.
 4. **Submit** your answers in a **single PDF document**.
 5. This paper consists of **4** pages, excluding the cover page.
-

QUESTION 1

Suppose you have joined Habitat for Humanity as a volunteer to help build house for families in need. You and the other volunteers will be separated into teams, which will be assigned to different tasks.

- (a) Identify at least four (4) tasks involved in the construction of a house. (2)
- (b) Which tasks from (a) exhibit task parallelism? (2)
- (c) Which tasks from (a) exhibit data-dependency? (2)
- (d) Once construction is finished, the house will need to be painted. How can we use data-parallelism to partition the work of painting the house? (2)
- (e) What are the three sources of overhead in parallel programs? (3)
- (f) How could the sources of overhead identified in (e) present themselves in the construction of a house? (6)
- (g) Give the following metrics for ring network topology (aka linear array with wraparound link) with n nodes: (3)
 - i. diameter
 - ii. bisection width
 - iii. number of links

Total: 20

QUESTION 2

- (a) You have been hired by a company to acquire a multi-processor machine for their resource intensive application. You determine that 90% of the tasks in the application can run in parallel, and that the work can be uniformly divided. You also determine that application is able to handle an increasing amount of work as more cores are added. You can purchase a 4, 8, or 16 core machine. Which machine would give you the best value for money? Motivate your answer. (9)
- (b) If an application running on a single processor has a fixed size problem, and spends 10% of its runtime on serial work and 90% on parallelizable work, what speedups can you expect to see on:
 - (i) 6 processors?
 - (ii) an unlimited amount of processors?(8)
- (c) Suppose that MPI COMM WORLD consists of the three processes 0, 1 and 2, and suppose the following code is executed: (3)

```
int x, y, z;
switch(my_rank) {
case 0:
    x=1; y=2; z=3;
    MPI_Bcast(&x, 1, MPI_INT, 0, MPI_COMM_WORLD);
    MPI_Send(&y, 1, MPI_INT, 2, 12, MPI_COMM_WORLD);
    MPI_Bcast(&z, 1, MPI_INT, 1, MPI_COMM_WORLD);
    break;
```

```

case 1:
    x=4; y=5; z=6;
    MPI_Bcast (&x, 1, MPI_INT, 0, MPI_COMM_WORLD);
    MPI_Bcast (&y, 1, MPI_INT, 1, MPI_COMM_WORLD);
    break;
case 2:
    x=7; y=8; z=9;
    MPI_Bcast (&z, 1, MPI_INT, 0, MPI_COMM_WORLD);
    MPI_Recv (&x, 1, MPI_INT, 0, 12, MPI_COMM_WORLD, &status);
    MPI_Bcast (&y, 1, MPI_INT, 1, MPI_COMM_WORLD);
    break;
}

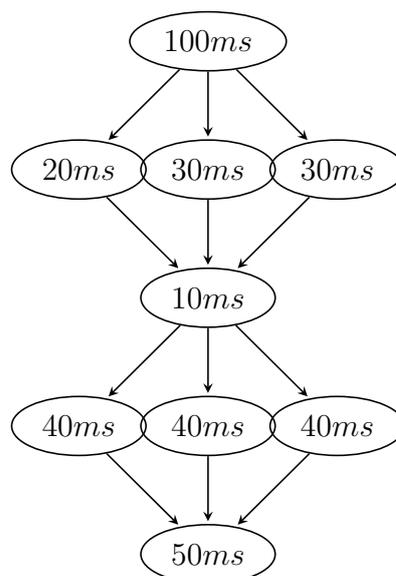
```

What are the values of x , y , and z on each process after the code has been executed?

Total: 20

QUESTION 3

Given the following task graph for a program running on a multiprocessor machine.



- Assuming a single (1) worker thread what is the runtime of this program? (2)
- What is the speedup when two (2) threads are used? (4)
- What is the maximum degree of concurrency of the graph? (1)
- What is the average degree of concurrency of the graph? (3)
- Show an efficient mapping of the tasks onto 3 processes. (5)
- Discuss what data dependency is, how it may hinder parallelism, and how it can be resolved. (4)

- (g) The number of tasks into which a problem is decomposed determines its granularity. Discuss granularity of task decompositions, as well its effect on the performance of parallel programs. Use the example of dense matrix-vector multiplication to support your discussion. (6)

Total: 25

QUESTION 4

Rather than simply finding the sum of n values,

$$x_0 + x_1 + \dots + x_{n-1},$$

the prefix sums are the n partial sums

$$x_0, x_0 + x_1, \dots, x_0 + x_1 + \dots + x_{n-1}$$

- (a) Devise a serial algorithm for computing the n prefix sums of an array with n elements. (4)
- (b) Parallelize your serial algorithm for a system with n processes, each of which stores one of the x_i s. Use only point-to-point communication. Assume that the array is referenced by a variable x at process zero and begin by having p_0 send the other processing their respective values. Each process should store its respective x_i in a variable x_i , and end by storing its corresponding prefix sum in a variable $prefix_x_i$. I.e. (9)

```

1   from mpi4py import MPI
2
3   comm = MPI.COMM_WORLD
4   size = comm.Get_size()
5   rank = comm.Getrank()
6
7   x_i = ...
8   .
9   .
10  .
11  sys.stdout.write("Process %d calculated prefix_x_i = %d.\n"
                    % (rank, prefix_x_i))

```

- (c) Suppose that we are working with a communicator of size 4 and that X is a 4×4 matrix. (12)

$$\begin{bmatrix} x_0 & x_1 & x_2 & x_3 \\ x_4 & x_5 & x_6 & x_7 \\ x_8 & x_9 & x_{10} & x_{11} \\ x_{12} & x_{13} & x_{14} & x_{15} \end{bmatrix}$$

- (i) How would the components of x be distributed among the processes in a program that used a block distribution?
- (ii) How would the components of x be distributed among the processes in a program that used a cyclic distribution?

- (iii) How would the components of x be distributed among the processes in a program that used a block-cyclic distribution with block-size $b = 2 \times 2$?

Total: 25

QUESTION 5

- (a) Briefly reflect on your experiences working on a research project for this module, drawing on the theory you have learnt. Include the following in your discussion: (10)
- A brief background of your research topic;
 - The results of your research;
 - The parallel algorithm design model you used.

Total: 10

— End of paper —