

UNIVERSITI SAINS MALAYSIA

First Semester Examination
Academic Session 2000/2001

September/October 2000

CSC503 - Foundations of Parallel Computing

Duration : [3 hours]

INSTRUCTION TO CANDIDATE:

- Please ensure that this examination paper contains **FOUR** questions in **FIVE** printed pages before you start the examination.
 - Answer **ALL** questions.
 - You can choose to answer either in Bahasa Malaysia or English.
-

ENGLISH VERSION OF THE QUESTION PAPER

1. (a) Describe a situation where you would prefer to use:
- (i) A distributed shared memory system as opposed to a message passing library and vice versa.
 - (ii) A functional decomposition method as opposed to data decomposition method and vice versa

Please give an example for each of the above.

(10 marks)

- (b) Given the following scenarios:

- (i) A student managed to parallelise an application X. The parallelised version takes a longer time to execute compared to the sequential version. Suggest three possible causes of the slow down and provide some solutions.
- (ii) A new research officer is puzzled by the performance of his parallel version of application Y. He is experiencing a superlinear speed up. What could be the causes – give the possible reasons?

(15 marks)

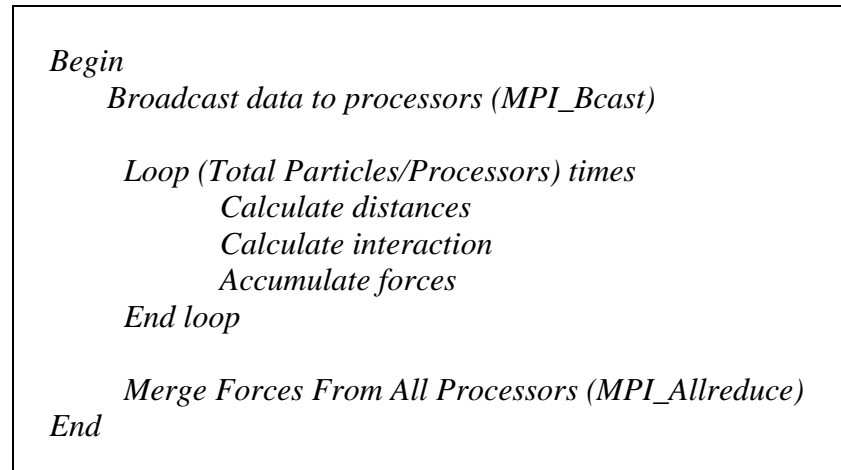
2. (a) Assuming you were given a task to parallelise an atmosphere modeling problem (as discussed in class). Describe the design stages that you may use prior to the code development.

(10 marks)

- (b) (i) Estimate the total execution time for a particle simulation problem with 1000 particles on a 4-node PC cluster. The most expensive computation occurs in the main loop which involves distance calculation (t_d), velocities (t_v) and accumulation of forces (t_f). Hence T_{comp} is the total computation time of all these three values.

Each processor makes two communications; that is sending of data during the configuration set up and also gathering of the final force results. Total communication time (T_{comm}) includes these total collective communications time. Provide the parallel cost model, T_{par} , for the above application.

Below is the general algorithm:



- (ii) Assuming that there exists a large discrepancy between the time obtained from the cost model compared to the actual time measured when running of the actual codes. Suggest the possible causes of discrepancies?

(15 marks)

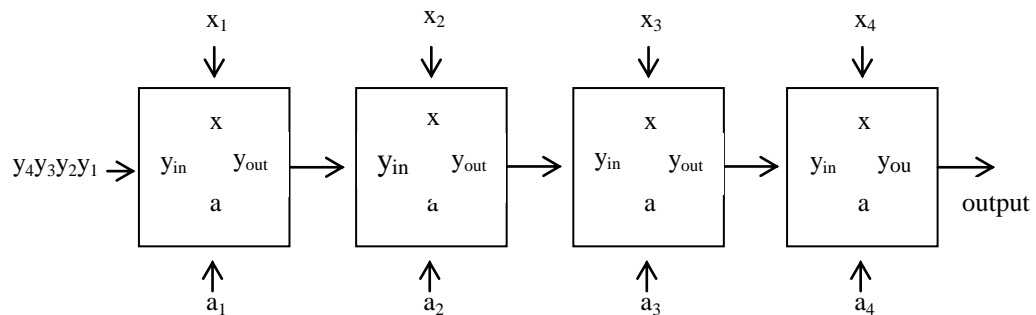
3. (a) Parallel algorithms might be coded in an old sequential language which has some add-on parallel feature, or they might be coded in a completely new parallel language.

Indicate the advantages and disadvantages of these two approaches.

(5 marks)

- (b) A pipeline consists of four stages as shown below. Each stage performs the operation $y_{out} = y_{in} + a * x$.

Determine the overall computation performed.



(5 marks)

- (c) (i) Differentiate between SIMD programming and SPMD programming.
- (ii) A simple example of a SIMD computation is to add the same constant to each element of an array; i.e.

```
for (i = 0; i < n; i++)
    a[i] = a[i] + k
```

A special parallel construct in parallel programming languages to specify the above computation is the *forall* statement. The above loop is written as:

```
forall (i = 0; i < n; i++)
    a [i] = a[i] + k
```

Draw a diagram to illustrate the execution of the *forall* loop above.

- (iii) Simulate the *forall* loop in (ii) on a message-passing computer where the *forall* statement is not available. Use the SPMD computation. (15 marks)

4. (a) History shows that parallelism has been used to improve the effectiveness of computers since the earliest design, and that it has been applied at several distinct levels which can be classified as:

- (i) Job level
- (ii) Program level
- (iii) Instruction level
- (iv) Arithmetic bit level

Explain briefly how parallelism has been achieved at each of the above levels.

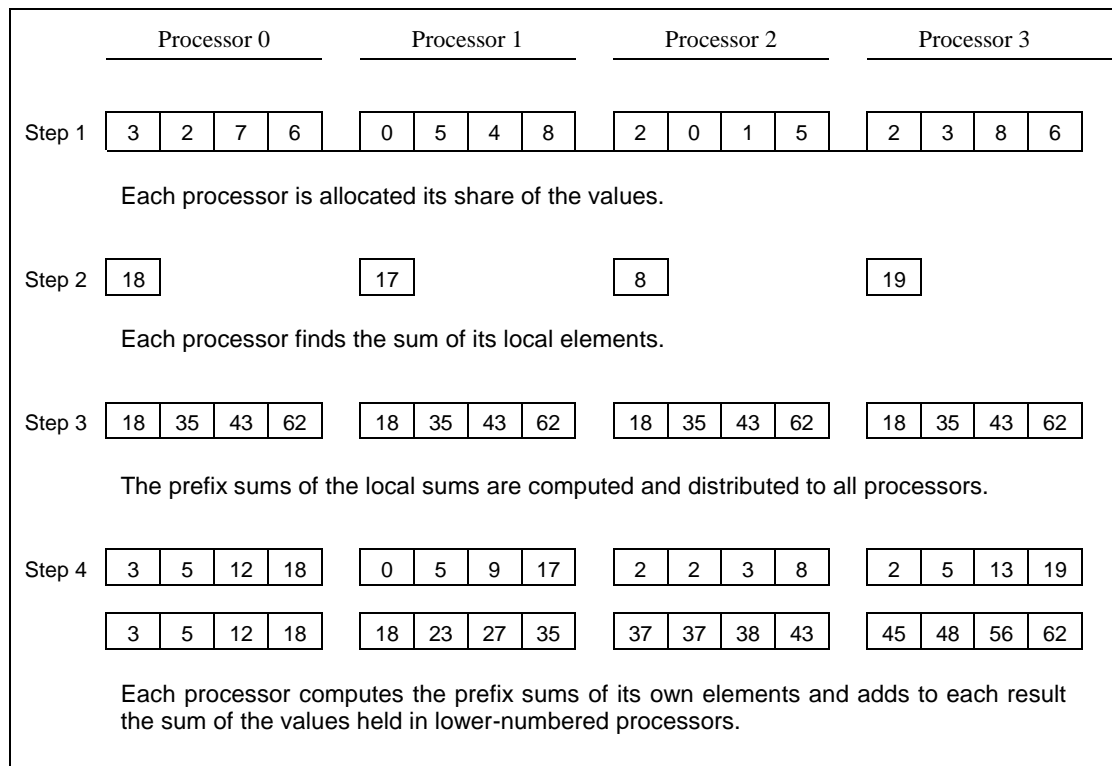
(10 marks)

- (b) Given a set of n values a_1, a_2, \dots, a_n and an associative operation \oplus , the **prefix sums** problem is to compute the n quantities:

$$\begin{aligned} &a_1 \\ &a_1 \oplus a_2 \\ &a_1 \oplus a_2 \oplus a_3 \\ &\dots \\ &a_1 \oplus a_2 \oplus a_3 \oplus \dots \oplus a_n \end{aligned}$$

For example, given the operation $+$ and the array of integers $\{3, 1, 0, 4, 2\}$, the prefix sums of the array are $\{3, 4, 4, 8, 10\}$.

Below is an algorithm to compute the prefix sums of 16 values on a 4-processor multicomputer.



Outline the PVM master/slave program that implements the above algorithm on a network of Sun workstation consisting of 4 nodes.

(You need not produce the exact code with the exact syntax. What is important is the algorithmic structure of the solution.)

(15 marks)