COMSATS University Islamabad
Abbottabad Campus

Department of Computer Science

FINAL TERM EXAMINATION SPRING 2025

Class: BCS-8B, BCS-8C
Subject: Parallel and Distributed Computing
Total Time Allowed: 03-hour
Student Name: _____

Date: 09-July-2024
Instructor: Ehtisham Ul Haque
Max Marks: 50
Registration #_____

Q.1. **Explain** the following questions. (CLO-1, C2) [Marks 8]

a) Briefly explain the following memory types *in parallel computing:*
- Global memory
- Shared memory
- Local/private memory
- Registers

b) Discuss the thread hierarchy in a GPU system (e.g., CUDA), including the terms:
- Thread
- Block
- Grid

Q.2. **Correlate** the following questions. (CLO-2, C3) [Marks 18]

a) Assume that 2D matrices are stored in a 1D array format in row-major order and that only 1D indexing is supported. Rewrite the following matrix multiplication expression using 1D indexing only:
P[i][j] += M[i][k] * N[k][j];
- Write the equivalent 1D-indexed expression.
- Briefly explain how 2D to 1D conversion is done.

b) Analyze efficiency of thread-memory mapping.
You are given a system where:
Global memory access is 100x slower than shared memory.
Each block has 256 threads.
A matrix A[1024][1024] must be multiplied with B[1024][1024] on GPU.
- Suggest a way to use thread *blocks and shared memory to improve* performance.
- What challenges may occur when multiple threads access the same memory location? How can these be avoided?

c) Understand and implement data partitioning for parallel computations
Given an array A[16] = {2, 4, 6, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28, 30, 32} to be processed by p = 4 processors, simulate the following steps:
- What is the partition size?
- What will be the start and end indices of the partition assigned to each processor (P0 to P3)?
- Write down which values of array A will be assigned to each processor.

CamScanner

Also,
If each processor calculates the sum of its assigned partition, write the value of:

- Sum[0] by Processor P0
- Sum[1] by Processor P1
- Sum[2] by Processor P2
- Sum[3] by Processor P3

Finally, compute and show grand_sum at Processor P0.

Q.3. **Code** the following questions. (CLO-3, C3) [Marks 24]

a) Write a CUDA C program to compute the matrix multiplication C = A × B using:
A Rectangular Thread Block of size (4 × 2) (i.e., blockDim.x = 4, blockDim.y = 2).
Given:

- Matrix A of size 16 × 8
- Matrix B of size 8 × 4
- Result matrix C

Your Program Should Include:

- CUDA kernel to compute each element C[row][col] using one thread
- Grid and block configuration using the specified dimensions
- Memory allocation on host and device
- Data transfer between host and device
- Final output of matrix C on the host

b) Write a code for parallel vector addition using MPI, where:

- Two vectors A and B (each of size N) are distributed among p processors.
- Each processor computes a partial result of vector C, where C[i] = A[i] + B[i].

Processor 0 acts as the sever node:

- Initializes vectors A and B
- Distributes data to all other processors
- Gathers the computed segments of vector C
- Prints a portion of the final result

Write a complete MPI program to:

- Accept vector size N as a command-line argument
- Partition vectors equally among p processors
- Perform parallel vector addition across all processes
- Reconstruct the full result in processor 0

CamScanner