BIRLA INSTITUTE OF TECHNOLOGY & SCIENCE, PILANI (RAJ.)
II SEMESTER 2022-2023
Comprehensive Examination

Course No.:	**CS F422**	**Course Title:**	**Parallel Computing**
Date:	**16th May 2023 (9:30-12:30)**	**Maximum Marks:**	**40**
Exam Type:	Part B (Open Book)	Part B Marks:	30

Note:
- Write all parts of a question together.

Q1) Consider a pattern matching algorithm given below.

```
Void search(char* pat, char* txt)
{
    int M = strlen(pat);
    int N = strlen(txt);

    /* A loop to slide pat[] one by one */
    for (inti  = 0; i <= N - M; i++) {
        int j;

        /* For current index i, check for pattern match */
        for (j = 0; j < M; j++)
            if (txt[i + j] != pat[j])
                break;

        if (j == M) // if pat[0...M-1] = txt[i, i+1, ...i+M-1]
            printf("Pattern found at index %d \n", i);
    }
}
```

(a) Using Foster's design methodology, explain how would you parallelize this algorithm. Use task dependency graph to identify parallelism.
(b) Considering a message passing parallel computer with p processors, implement your parallel algorithm in MPI.
(c) Analyze the parallel time complexity and scalability.
(d) Comment on the cost-optimality of the parallel algorithm.

[8M]

Q2) A GPU has 80 streaming multiprocessors each having 64 cores. There are 4 warp schedulers and warpSize is 32. Assuming there are 10^{10} real numbers (8 bytes of storage per number), shared memory per SM is 4KB, registers per SM is 32K and global memory of GPU is 1 GB, answer the following

(a) Write a kernel in CUDA C to find the maximum number.
(b) Identify gird and block dimensions for optimizing GPU performance.
(c) Compute occupancy ratio.
(d) Write the main function to find the maximum number by using the kernel written in (a)
(e) Re-write the main function so that data transfer and kernel execution can happen in parallel.

[8M]

Q3) Consider a large list of numbers, all integers in the csv file *data.csv*. Write a MPI program for the following requirements

(a) each process should read its portion of integers in parallel
(b) the program should find numbers which are divisible by 2, 3, 5, 7, 11 and 13
(c) each process should write the result to a file output.csv in parallel

[4M]

Q4) Given a list of numbers, write a CUDA program to find out if a number is a perfect number. Perfect number is a positive integer that is equal to the sum of its proper divisors. The smallest perfect number is 6, which is the sum of 1, 2, and 3.

[5M]

Q5) For the following, answer briefly.

(a) Using E-cube routing on hypercube of d dimensions, communication between any two nodes can be utmost d hops. State true / false with justification.
(b) If parallel runtime is $n^3/p + p \log p$, what is its per process memory scaling function?
(c) What is the relationship between overhead, work size and cost-optimality?
(d) Explain time taken for all to all broadcast using Ecube routing on Hypercube.
(e) In matrix-vector multiplication, what will be the parallel runtime if vector is distributed among processes and each process is responsible for a single row.

[5M]

--&--