

BIRLA INSTITUTE OF TECHNOLOGY & SCIENCE, PILANI (RAJ.)
II SEMESTER 2019-2020
Comprehensive Examination

Course No.: CS F422	Course Title: Parallel Computing
Date: 12th May 2022 (8-11)	Maximum Marks: 35
Exam Type: Part B (Open Book)	Part B Marks: 20

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.
- (b) Considering a message passing parallel computer with p processors, implement your parallel algorithm in MPI.
- (c) Comment on the cost-optimality of the parallel algorithm.
- (d) Analyze the parallel time complexity and scalability.

[2+4+1+1=8M]

Q2) An array of type float elements is to be processed in a one-element-per-thread fashion by a GPU. Suggest an execution configuration for the following scenarios:

- (a) The array is 1D and of size N. The target GPU has 8 SMs, each with 16 SPs.
- (b) The array is 2D and of size NxN. The target GPU has 5 SMs, each with 48 SPs.

For each of these scenarios calculate the minimum size that N should satisfy to make the GPU computation a desirable alternative to CPU computation.

[2+2=4M]

Q3) Consider the format of data with sample data given in the following table. This is the data of registered courses.

Year	Semester	StudentId	Courseld
2019	1	ld1	cid1
2019	2	ld2	cid2
2020	1	ld3	cid1
2020	1	ld4	cid1
2019	2	ld5	cid2
2019	2	ld2	cid2

For each semester in a year, we need to find out the list of students who have taken the maximum number of courses in a semester. Design map-reduce application with pseudo code for map() and reduce() tasks.

[4M]

Q4) In shared memory programming, locks are used to serialize the access to critical sections. Consider that there is a procedure named Element dequeuer(Queue q) which removes and returns the last element in a given queue.

```

Element dequeuer(Queue q){
    Element e;
    pthread_mutex_lock(&mutex);
    //code
    pthread_mutex_unlock(&mutex);
    return e;
}

```

Considering that the application uses several queue data structures, suggest a locking scheme with better performance.

[2M]

Q5) For the following, justify whether they are true or false.

- (a) Using E-cube routing on hypercube of d dimensions, communication between any two nodes can be utmost d hops.
- (b) Number of registers in a SM determine the block and grid sizes in GPU programming.

[2M]

--&--

BIRLA INSTITUTE OF TECHNOLOGY & SCIENCE, PILANI (RAJ.)
II SEMESTER 2021-2022
Comprehensive Examination

Course No.: CS F422

Course Title: Parallel Computing

Date: 12th May 2022 (8-11)

Maximum Marks: 35

Exam Type: Part A (Closed Book)

Part A Marks: 15

Note:

- Write answers in this sheet itself.
- Overwritten answers will not be accepted for rechecks.
- After submitting part A, collect part B.

IDNO: _____

NAME: _____

Q1) There can be more than one choice correct. Only if all correct options and only correct options are chosen, marks will be awarded. Write answers in the grid. Each carries 0.5 mark.

1	2	3	4	5	6	7	8	9	10		
Recheck:								Q1	Q2	Q3	Total

- | | |
|---|---|
| <ol style="list-style-type: none"> 1. MPI program is compiled using <ol style="list-style-type: none"> (a) mpigcc (b) gcc (c) gcc -mpi (d) mpicc 2. MPI_Reduce function <ol style="list-style-type: none"> (a) requires point to point communication (b) requires collective communication (c) applies given operator associatively on the data held by each process (d) reduces number of dimensions in input data 3. Cost of one-to-all broadcast operation on a hypercube of p nodes is ____. 4. With respect to map reduce framework, <ol style="list-style-type: none"> (a) reduce task can't have several keys in its input (b) a key can be with utmost one reduce task only (c) there can be only one reduce job (d) there can be only one map job 5. MPI follows ____ style of programming <ol style="list-style-type: none"> (a) SIMD (b) SPMD (c) MPMD (d) SISD 6. In Foster's design methodology, the process of grouping tasks into larger tasks in order to | <p>improve performance or simplify programming is called</p> <ol style="list-style-type: none"> (a) Partitioning (b) Communication (c) Agglomeration (d) Mapping <ol style="list-style-type: none"> 7. Mapping strategy used when number of tasks is static, communication is structured and roughly constant time per task is <ol style="list-style-type: none"> (a) Use static load balancing algorithm (b) Use dynamic load balancing algorithm (c) Cyclically map tasks to processors (d) Create one task per processor 8. A collective communication method in MPI <ol style="list-style-type: none"> (a) is called by all processes in the communicator (b) is called only by a few in the communicator (c) uses tags to match messages (d) do not use tags to match messages 9. W.r.t CUDA, choose the correct options <ol style="list-style-type: none"> (e) Host and GPU memories are disjoint (f) Host can access shared memory of GPU (g) Thread index is unique globally (h) Thread index is unique within the block 10. Cost of scatter operation on a linear array of p nodes is ____. |
|---|---|

Q3) Write a CUDA program for adding two vectors and storing result in a third a vector. Syntax errors are not penalized. 2M

Q4) Write an MPI program for adding two vectors and storing result in a third a vector. Syntax errors are not penalized. 2M