

Please fill out the next table with your answers of Part1:

Q. number	Answer	Q. number	Answer
1		11	
2		12	
3		13	
4		14	
5		15	
6		16	
7		17	
8		18	
9		19	
10		20	

Part 1: Select the most correct answer for the following questions.

- [CLO_1.1] Why do we need to write parallel programs?

 - To run multiple instances of a program on a multicore system
 - To exploit the power of multiple processors
 - To run multiple instances of our favourite game program
 - To make use of a single core processor

- [CLO 1.1] Which of the following statements is true?

 - Any object of Thread is not an instance of Runnable
 - Any object of Thread is an instance of Runnable
 - Any object of Runnable is an instance of Thread
 - Thread is a subclass of Runnable

- [CLO 1.1] A 3D-hypercube interconnect is composed of:

 - 2 two-dimensional hypercubes
 - 3 one-dimensional hypercubes
 - A fully connected network
 - A crossbar interconnect

- [CLO 1.1] Which method is used to force one thread to wait for another thread to finish i

finish()
sleep()
join()
wait()

- [CLO 1.1] The bisection width of a fully connected network in which each switch is dire

every other switch is...

 - $p/2$
 - $2p$
 - $2\sqrt{p}$
 - $p^2/4$

6. [CLO 1.1] means one or more cores send their current local sums to the master core.
- Load balancing
 - Cooperation
 - Communication
 - Synchronization
7. [CLO 1.1] In CUDA framework, what does the term "host" mean?
- The main server
 - The GPU
 - The CPU
 - The main memory
8. [CLO 1.1] In MPI implementations, which process is responsible of reading the data (scanf) and send to the other processes?
- Process 1
 - Process 0
 - The process with the highest rank
 - It is up for the programmer to specify that in the code.
9. [CLO 1.1] Select the true statement.
- Shared variables are visible by all running threads
 - Shared variables are visible to a certain thread
 - Private variables are visible by all running threads
 - b and c
10. [CLO 1.1] which of the following is an NP-complete problem?
- vector summation
 - TSP
 - matrix multiplication
 - vectors addition
11. [CLO 1.2] In distributed memory interconnects, which of the following is true regarding direct interconnects?
- Each switch is directly connected only to another switch
 - Each switch is directly connected to a CPU-memory pair
 - Each switch is directly connected only to multiple switches
 - Each CPU-memory pair is directly connected to another CPU-memory pair
12. [CLO 1.2] In OpenMP, which of the following clauses is used to change the distribution of loop's iterations during parallel execution?
- reduction
 - default
 - schedule
 - private

Q6. [CLO 2.1] A. Implement an alternative version of the OpenMP code below where each thread computes its own part to a private variable and then uses a critical directive to compute the global result. [4_marks]

```
int result = 1;
# pragma omp parallel num_threads(4) reduction (*:result)
{
    result *= FunctionA();
}
```

.....

.....

.....

.....

.....

B. Can you replace the critical section in your code with an atomic directive? If the answer is Yes, write the code. If the answer is NO, explain why it cannot be used.

.....

.....

.....

.....

Q7. [CLO 2.1] Find and correct 3 errors in the code below (assume all the required classes have been imported to the project).

(Be specific and mark the error location in the code) (Don't mark more than 3 locations in the code. If so, only the first 3 corrections/suggestions will be considered) [3_marks]

```
public class ExamClass{
    public static void main(String[] args){
        Thread t1 = new Task ();
        Task t2 = new Task ();
        t1.start();
        t2.start();
        t1.join();
    }
}
```

```
class Task extends Thread{
    static Object obj = new Object();
    static Lock something = new Lock();
    void run(){
        something.lock();
        something.release();
        t2.join();
    }
}
```

18. [CLO 2.1] Which of these parallel programming errors are impossible to occur in the given OpenMP construct?

- a) Data conflict in #pragma omp critical
- b) Deadlock in #pragma omp parallel
- c) Data dependency in #pragma omp for
- d) Data race in #pragma omp parallel



19. [CLO 2.1] What does the following MPI statement mean?


```
MPI_Gather(collectedarray, 100, MPI_INT, recv_data, 100, MPI_INT, 0, comm);
```

- a) Broadcast 100 int elements to every process in the communicator comm including the root process 0
- b) Collect 100 int elements from the root process 0 to every process in the communicator comm
- c) Collect 100 int elements from every process in the communicator comm to the root process 0
- d) Obtain 100 int elements from every process in the communicator comm and distribute them to every process in the communicator comm including the root process 0



20. [CLO 2.1] According to the following figure, identify the case that the processes can have simultaneous communication.

- P0 can communicate with P4 and P2 can communicate with P5
- P1 can communicate with P3 and P0 can communicate with P4
- P4 can communicate with P6 and P5 can communicate with P7
- P3 can communicate with P6 and P1 can communicate with P4



Part 2: Answer the following questions


[20 marks]

Q1. [CLO 1.2] Using one of the MPI collective communication functions, write a single MPI statement that will work as an alternative to the following point to point communication code. [1 mark]

```

If (my_rank!=0)
MPI_Send(&var1, 1, MPI_DOUBLE, 0, 0, MPI_COMM_WORLD)
else
{
    For (source=1; source<comm_sz; source++)
    MPI_Recv(&var1, 1, MPI_DOUBLE, source, 0, MPI_COMM_WORLD, MPI_STATUS_IGNORE)
    Total+=var;
}
    
```

13. [CLO 1.2] A suitable programming environment for a parallel system as illustrated in the diagram is...
- a) Java multi-threading
 - b) OpenMP
 - c) Java executor
 - d) MPI



14. [CLO 1.2] Suppose multiple multicore processors with two chips each with two cores. Core 1 in chip 1 needs 1 nanosecond to access variable X, while core 2 in chip 2 needs 1.5 nanoseconds to access the variable X. This system can be described as....
- a) Shared-memory system using UMA
 - b) Distributed-memory system using UMA
 - c) Shared-memory system using NUMA
 - d) Distributed-memory system using NUMA

15. [CLO 1.2] Which of the following methods cannot directly cause a thread to stop executing?
- a) Calling notify() method on an object
 - b) Calling the await() method on an object
 - c) Calling the join() method on a Thread object
 - d) Calling sleep() method on a Thread object

16. [CLO 1.2] Which of the following paradigms is not included in Flynn's taxonomy?
- SIMD
 - SPMD
 - MIMD
 - MISD

17. [CLO 2.1] What is the maximum number of threads that can access the critical section (simultaneously)?

```
Static Semaphore lock = new Semaphore(3);
public void run(){
    lock.acquire(2);
    System.out.println("hello");
    lock.release(2);
}
```

- 3
- 2
- 4
- 1