

# CS6801 2 MARKS

## UNIT –I

### MULTI-CORE PROCESSORS

Single core to Multi-core architectures – SIMD and MIMD systems – Interconnection networks - Symmetric and Distributed Shared Memory Architectures – Cache coherence - Performance Issues –Parallel program design.

Q. No.	Questions	CO	Bloom's Level									
1.	<p><b>What is Single core processors?</b> Single core processors have only one processor in die to process instructions.</p>	C409.1	BTL1									
2.	<p><b>What are the Problems of Single Core Processors:</b> As we try to increase the clock speed of this processor , the amount of heat produced by the chip also increases. It is a big hindrance in the way of single core processors to continue evolving</p>	C409.1	BTL1									
3.	<p><b>What is Multicore processor?</b> A multi-core processor is a single computing component with two or more independent actual processing units (called "cores"), which are units that read and execute program instructions. The multiple cores are embedded in the same die. The multicore processor may looks like a single processor but actually it contains two (dual - core), three (tri - core), four (quad - core), six(hexa-core), eight(octa-core)or ten (deca-core) cores.Some processor even have 22 or 32 cores..</p>	C409.1	BTL1									
4.	<p><b>What are the Problems with multicore processors.</b> According to Amdahl's law , the performance of parallel computing is limited by its serial components. So, increasing the number of cores may not be the best solution .There is need to increase the clock speed of individual cores.</p>	C409.1	BTL1									
5.	<p><b>Comparison Of Single-Core processor And Multi-Core Processor.</b></p> <table border="1"> <thead> <tr> <th>Parameter</th> <th>Single-Core Processor</th> <th>Multi-Core Processor</th> </tr> </thead> <tbody> <tr> <td>Number of cores on a die</td> <td>Single</td> <td>Multiple</td> </tr> <tr> <td>Instruction Execution</td> <td>Can execute Single instruction at a time</td> <td>Can execute multiple instructions by using multiple cores</td> </tr> </tbody> </table>	Parameter	Single-Core Processor	Multi-Core Processor	Number of cores on a die	Single	Multiple	Instruction Execution	Can execute Single instruction at a time	Can execute multiple instructions by using multiple cores	C409.1	BTL1
Parameter	Single-Core Processor	Multi-Core Processor										
Number of cores on a die	Single	Multiple										
Instruction Execution	Can execute Single instruction at a time	Can execute multiple instructions by using multiple cores										

	Gain	Speed up every program or software being executed	Speed up the programs which are designed for multi-core processors		
	Performance	Dependent on the clock frequency of the core	Dependent on the frequency, number of cores and program to be executed		
	Examples	Processor launched before 2005 like 80386,486, AMD 29000, AMD K6, Pentium I,II,III etc.	Processor launched after 2005 like Core - 2-Duo,Athlon 64 X2,13,15 and I7 etc		
6	<p><b>What is meant by Single instruction, multiple data (SIMD)</b></p> <p><b>Single instruction, multiple data (SIMD)</b>, is a class of parallel computers in <u>Flynn's taxonomy</u>. It describes computers with <u>multiple processing elements</u> that perform the same operation on multiple data points simultaneously. Thus, such machines exploit <u>data level parallelism</u>, but not <u>concurrency</u>: there are simultaneous (parallel) computations, but only a single process (instruction) at a given moment</p>			C409.1	BTL1
7	<p><b>What is meant by multiple instruction, multiple data (MIMD)</b></p> <p>In <u>computing</u>, <b>MIMD</b> (multiple instruction, multiple data) is a technique employed to achieve parallelism. Machines using MIMD have a number of <u>processors</u> that function <u>asynchronously</u> and independently. At any time, different processors may be executing different instructions on different pieces of data. MIMD architectures may be used in a number of application areas such as <u>computer-aided design/computer-aided manufacturing</u>, <u>simulation</u>, <u>modeling</u>, and as <u>communication switches</u>. MIMD machines can be of either <u>shared memory</u> or <u>distributed memory</u> categories.</p>			C409.1	BTL1
8	<p><b>Define Multistage interconnection networks.</b></p> <p>Multistage <b>interconnection networks</b> (MINs) are a class of high-speed computer <b>networks</b> usually composed of processing elements (PEs) on one end of the <b>network</b> and memory elements (MEs) on the other end, connected by switching elements (SEs).</p>			C409.1	BTL1
9	<p><b>What is meant by Routing?</b></p> <ul style="list-style-type: none"> <li>-How does a message get from source to destination.</li> <li>-Static or adaptive</li> </ul>			C409.1	BTL1

10	<b>What is Network interface?</b> -Connects endpoints (e.g. cores) to network. -Decouples computation/communication	C409.1	BTL1
11	<b>.What is Centralized shared-memory multiprocessor</b> It share a single centralized memory, interconnect processors and memory by a bus • also known as “uniform memory access” (UMA) or “symmetric (shared-memory) multiprocessor” (SMP) – A symmetric relationship to all processors. – A uniform memory access time from any processor.	C409.1	BTL1
12	<b>What is the concept of Caching in shared-memory machines.</b> • private data: used by a single processor – When a private item is cached, its location is <i>migrated</i> to the cache – Since no other processor uses the data, the program behavior is identical to that in a uniprocessor • shared data: used by multiple processor – When shared data are cached, the shared value may be <i>replicated</i> in multiple caches.	C409.1	BTL1
13	<b>What is Cache Coherence .</b> • migration: a data item can be moved to a local cache and used there in a transparent fashion • replication for shared data that are being simultaneously read both are critical to performance in accessing shared data.	C409.1	BTL1
14	<b>What is meant by Snooping Solution (Snoopy Bus).</b> – Send all requests for data to all processors – Processors snoop to see if they have a copy and respond accordingly – Requires broadcast, since caching information is at processors – Works well with bus (natural broadcast medium) – Dominates for small scale machines (most of the market)	C409.1	BTL1
15	<b>What is meant by Directory-Based Schemes.</b> – Directory keeps track of what is being shared in a centralized place (logically) – Distributed memory => distributed directory for scalability (avoids bottlenecks) – Send point-to-point requests to processors via network – Scales better than Snooping – Actually existed BEFORE Snooping-based schemes	C409.1	BTL1
16	<b>When a memory system is coherent ?</b> A memory system is coherent if: • P writes to X; no other processor writes to X; P reads X and receives the value previously written by P • P1 writes to X; no other processor writes to X; sufficient time lapses; P2 reads X and receives value written by P1 • Two writes to the same location by two processors are seen in the	C409.1	BTL1

	<p>same order by all processors – write serialization</p> <ul style="list-style-type: none"> <li>• The memory consistency model defines “time elapsed” before the effect of a processor is seen by others.</li> </ul>												
17	<p><b>What is meant by a distributed-memory system?</b></p> <p>A distributed-memory system (often called a multicomputer) consist of multiple independent processing nodes with local memory modules which is connected by a general interconnection network. Software DSM systems can be implemented in an operating system, or as a programming library and can be thought of as extensions of the underlying virtual memory architecture.</p>	C409.1	BTL1										
18	<p><b>What the difference between Message passing vs. DSM</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 50%; text-align: left;">Message passing</th> <th style="width: 50%; text-align: left;">Distributed shared memory</th> </tr> </thead> <tbody> <tr> <td>Variables have to be marshalled</td> <td>Variables are shared directly</td> </tr> <tr> <td>Cost of communication is obvious</td> <td>Cost of communication is invisible</td> </tr> <tr> <td>Processes are protected by having private address space</td> <td>Processes could cause error by altering data</td> </tr> <tr> <td>Processes should execute at the same time</td> <td>Executing the processes may happen with non-overlapping lifetimes</td> </tr> </tbody> </table>	Message passing	Distributed shared memory	Variables have to be marshalled	Variables are shared directly	Cost of communication is obvious	Cost of communication is invisible	Processes are protected by having private address space	Processes could cause error by altering data	Processes should execute at the same time	Executing the processes may happen with non-overlapping lifetimes	C409.1	BTL1
Message passing	Distributed shared memory												
Variables have to be marshalled	Variables are shared directly												
Cost of communication is obvious	Cost of communication is invisible												
Processes are protected by having private address space	Processes could cause error by altering data												
Processes should execute at the same time	Executing the processes may happen with non-overlapping lifetimes												
19	<p><b>What are the Advantages of DSM. (Apr/May 2018)</b></p> <ul style="list-style-type: none"> <li>• System scalable</li> <li>• Hides the message passing</li> <li>• Can handle complex and large data bases without replication or sending the data to processes</li> <li>• DSM is usually cheaper than using multiprocessor system</li> <li>• No memory access bottleneck, as no single bus</li> <li>• DSM provides large virtual memory space</li> <li>• DSM programs portable as they use common DSM programming interface</li> <li>• Shields programmer from sending or receive primitives</li> </ul> <p>DSM can (possibly) improve performance by speeding up data access</p>	C409.1	BTL1										

20	<p><b>What the Issues in implementing DSM software</b></p> <ul style="list-style-type: none"> <li>• Data is replicated or cached</li> <li>• Reduce delays</li> <li>• Semantics for concurrent access must be clearly specified</li> <li>• DSM is controlled by memory management software, operating system, language run-time system</li> <li>• Locating remote data</li> <li>• Granularity: how much data is transferred in a single operation</li> </ul>	C409.1	BTL1						
21	<p><b>What are the Disadvantages of DSM (Apr/May 2018)</b></p> <ul style="list-style-type: none"> <li>• Could cause a performance penalty</li> <li>• Should provide for protection against simultaneous access to shared data such as lock</li> <li>• Performance of irregular problems could be difficult</li> </ul>	C409.1	BTL1						
22	<p><b>What are the Methods of achieving DSM.</b>                  There are usually two methods of achieving distributed shared memory:</p> <ul style="list-style-type: none"> <li>• hardware, such as cache coherence circuits and network interfaces;</li> <li>• software. We can use this method in different ways such as modifying the operating system kernel.</li> </ul>	C409.1	BTL1						
23	<p><b>What is meant by Consistency models..</b>                  Memory system tries to behave based on certain rules in the system, which is called system's <i>consistency model</i>.</p>	C409.1	BTL1						
24	<p><b>Define Vector Instruction?(Apr/May2017)</b></p> <p>A vector processor or array processor is a central processing unit (CPU) that implements an instruction set containing instructions that operate on one-dimensional arrays of data called vectors, compared to scalar processors, whose instructions operate on single data items.</p>	C409.1	BTL1						
25	<p><b>What is meant by Snooping cache coherence? (Apr/May 2017)</b></p> <p>Also referred to as a bus-snooping protocol, a protocol for maintaining cache coherency in symmetric multiprocessing environments. In a snooping system, all caches on the bus monitor (or snoop) the bus to determine if they have a copy of the block of data that is requested on the bus.</p>	C409.1	BTL1						
26	<p><b>Compare Symmetric memory architecture and distributed memory architecture. (Nov/Dec 2017)</b></p> <table border="1" data-bbox="289 1724 1260 1862"> <thead> <tr> <th data-bbox="289 1724 407 1766">Sno</th> <th data-bbox="407 1724 748 1862">Symmetric memory architecture</th> <th data-bbox="748 1724 1260 1862">distributed memory architecture.</th> </tr> </thead> <tbody> <tr> <td data-bbox="289 1766 407 1862"></td> <td data-bbox="407 1766 748 1862"></td> <td data-bbox="748 1766 1260 1862"></td> </tr> </tbody> </table>	Sno	Symmetric memory architecture	distributed memory architecture.				C409.1	BTL1
Sno	Symmetric memory architecture	distributed memory architecture.							

	<p><b>1</b> sharedmemory cache-coherent multiprocessor systems. The systems communicated with each other and with shared main memory over a shared bus.</p> <p><b>2</b> any access from any processor to main memory would have equal latency</p>	<p><b>distributed memory</b> refers to a multiprocessor computer system in which each processor has its own private memory. Computational tasks can only operate on local data, and if remote data is required, the computational task must communicate with one or more remote processors</p> <p>any access from any processor to main memory would have different latency</p>		
27	<p><b>What are multiprocessor systems and give their advantages? (Nov/Dec 2017)</b></p> <p>Multiprocessor systems also known as parallel systems or tightly coupled systems are systems that have more than one processor in close communication, sharing the computer bus, the clock and sometimes memory &amp; peripheral devices.</p> <p>Their main advantages are</p> <ol style="list-style-type: none"> <li>1 Increased throughput</li> <li>2 Economy of scale</li> <li>3 Increased reliability</li> </ol>		C409.1	BTL1
28	<p><b>Define Channel?</b></p> <p>A single logical connection between routers/switches</p>		C409.1	BTL1
29	<p><b>List the pros and cons of distributed system (Apr/May 2018)</b></p> <ul style="list-style-type: none"> <li>• System scalable</li> <li>• Hides the message passing</li> <li>• Can handle complex and large data bases without replication or sending the data to processes</li> <li>• DSM is usually cheaper than using multiprocessor system</li> <li>• No memory access bottleneck, as no single bus</li> <li>• DSM provides large virtual memory space</li> <li>• DSM programs portable as they use common DSM programming interface</li> <li>• Shields programmer from sending or receive primitives</li> </ul> <p>DSM can (possibly) improve performance by speeding up data access</p> <ul style="list-style-type: none"> <li>• Could cause a performance penalty</li> <li>• Should provide for protection against simultaneous access to shared data such as lock</li> </ul>		C409.1	BTL1

	Performance of irregular problems could be difficult		
30	<p><b>Define the symmetric shared memory (Apr/May 2018, Nov/Dec 2018)</b></p> <p><b>Symmetric Shared Memory Architecture</b> consists of several processors with a single physical <b>memory shared</b> by all processors through a <b>shared bus</b></p>	C409.1	BTL1
31	<p><b>List out the advantages of multicore CPU</b></p> <ul style="list-style-type: none"> <li>• The largest boost in performance will likely be noticed in improved response time while running CPU intensive processes, like anti-virus scans, ripping/burning media.</li> <li>• Assuming that the die can fit into the package, physically, the multi-core CPU designs require much less printed Circuit Board(PCB) space than multichip SMP designs.</li> <li>• Also, a dual core processor uses slightly less power than two coupled single core processors, principally because of the decreased power required to drive signals external to the chip</li> </ul>	C409.1	BTL1
32	<p><b>Define Vector Registers</b></p> <p>These are registers capable of storing a vector of operands and operating simultaneously on their contents. The vector length is fixed by the system, and can range from 4 to 128 64-bit elements. Vectorized and pipelined functional units.</p>	C409.1	BTL1
33	<p><b>Define Latency and Bandwidth</b></p> <ul style="list-style-type: none"> <li>• The latency is the time that elapses between the source's beginning to transmit the data and the destination's starting to receive the first byte.</li> <li>• The bandwidth is the rate at which the destination receives data after it has started to receive the first byte. So if the latency of an interconnect is 1 seconds and the bandwidth is b bytes per second, then the time it takes to transmit a message of n bytes is</li> <li>• <b>Message transmission time= <math>l+n/b</math></b></li> </ul>	C409.1	BTL1
34	<p><b>List out the approaches in cache coherence</b></p> <ul style="list-style-type: none"> <li>• Snooping cache coherence</li> <li>• Directory-based cache coherence.</li> </ul>	C409.1	BTL1
35	<p><b>List the steps involved in Parallel Program design</b></p> <p>1.Partitioning</p>	C409.1	BTL1

	<p>2. Aggregating</p> <p>3. Communication</p> <p>4. Mapping</p>		
36	<p><b>Define Directory based cache coherence</b></p> <p>It is protocols attempt to solve this problem through the use of a data structure called a <b>directory</b>. The directory stores the status of each cache line. Typically, this data structure is distributed; in our example, each core/memory pair might be responsible for storing the part of the structure that specifies the status of the cache lines in its local memory</p>	C409.1	BTL1
37	<p><b>Define Parallel Overhead</b></p> <p><b><math>T_{parallel} = T_{serial}/p + T_{overhead}</math>.</b></p>	C409.1	BTL1
38	<p><b>Define Scalability</b></p> <p>The number of processes/threads that are used by the program. If we can find a corresponding rate of increase in the problem size so that the program always has efficiency E, then the program is scalable.</p>	C409.1	BTL1
39	<p><b>Define Amdahl's Law (Nov/Dec 2018)</b></p> <p>Amdahl made an observation that's become known as <i>Amdahl's law</i>. It says, roughly, that unless virtually all of a serial program is parallelized, the possible speedup is going to be very limited—regardless of the number of cores available.</p>	C409.1	BTL1
40	<p><b>Define Speedup and Efficiency</b></p> <p>The Serial run-time <math>T_{serial}</math> and our parallel run-time <math>T_{parallel}</math>, then the best we can hope for is <math>T_{parallel} = T_{serial}/p</math>.</p> <p>Parallel program has <b>linear speedup</b>. So if we define the <b>speedup</b> of a parallel program to be linear speedup has <math>S = p</math>, which is unusual. Furthermore, as <math>p</math> increases, we expect <math>S</math> to become a smaller and smaller fraction of the ideal, linear speedup <math>p</math>.</p>	C409.1	BTL1

41	<p><b>How to parallelize the serial program</b></p> <ul style="list-style-type: none"> <li>• For the first step we might identify two types of tasks: finding the bin to which an element of data belongs and incrementing the appropriate entry in bin counts.</li> <li>• For the second step, there must be a communication between the computation of the appropriate bin and incrementing an element of bin counts.</li> </ul>	C409.1	BTL1
42	<p><b>List out the different distributed memory interconnects</b></p> <p>Distributed-memory interconnects are often divided into two groups:  <b>Direct interconnects</b> and <b>Indirect interconnects</b></p>	C409.1	BTL1
43	<p><b>Define Direct Interconnects</b></p> <p>In a direct interconnect each switch is directly connected to processor memory pair, and the switches are connected to each other.  As before, the <i>circles</i> are <i>switches</i>, the <i>squares</i> are <i>processors</i>, and the <i>lines</i> are <i>bidirectional links</i>.  A <i>ring</i> is superior to a simple bus since it allows multiple simultaneous communications.</p>	C409.1	BTL1
44	<p><b>Define Indirect Interconnects</b></p> <ul style="list-style-type: none"> <li>• They provide an alternative to direct interconnects. In an indirect interconnect the switches may not be directly connected to a processor.</li> <li>• They're often shown with unidirectional links and a collection of processors, each of which has an outgoing and an incoming link, and a switching network.</li> </ul>	C409.1	BTL1
45	<p><b>Define Ideal Direct interconnect</b></p> <p>The ideal direct interconnect is a <i>fully connected network</i> in which each switch is directly connected to every other switch</p>	C409.1	BTL1
46	<p><b>Define Hypercube</b></p> <ul style="list-style-type: none"> <li>• It is a highly connected direct interconnect that has been used in actual systems. Hypercubes are built inductively:</li> <li>• A one-dimensional hypercube is a fullyconnected system with two processors.</li> <li>• A two-dimensional hypercube is built from two one-dimensional hypercubes by joining "corresponding" switches</li> </ul>	C409.1	BTL1
47	<p><b>Define Interleaved memory</b></p> <p>The memory system consists of multiple "banks" of memory, which can be accessed more or less independently. After accessing one bank, there will be</p>	C409.1	BTL1

	a delay before it can be reassessed, but a different bank can be accessed much sooner. So if the elements of a vector are distributed across multiple banks, there can be little to no delay in loading/storing successive elements.		
48	<b>Define Strided memory</b>  In strided memory access, the program accesses elements of a vector located at fixed intervals. For example, accessing the first element, the fifth element, the ninth element, and so on, would be strided access with a stride of four	C409.1	BTL1
49	<b>Define Graphics Processor Pipeline</b>  Real-time graphics application programming interfaces, or APIs, use points, lines, and triangles to internally represent the surface of an object. They use a <b>graphics processing pipeline</b> to convert the internal representation into an array of pixels that can be sent to a computer screen	C409.1	BTL1
50	<b>List out the two principal types of MIMD system</b>  <ul style="list-style-type: none"> <li>• Shared Memory System</li> <li>• Distributed Memory system</li> </ul>	C409.1	BTL1

### PART B

Q. No.	Questions	CO	Bloom's Level
1	<b>Explain Single core to Multi-core Architectures .</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:15-20	C409.1	BTL5
2	<b>Explain SIMD and MIMD systems (Apr/May 2017, Nov/Dec 2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:29-34	C409.1	BTL5
3	<b>Explain about Interconnection networks? (Apr/May 2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:35-44	C409.1	BTL5
4	<b>Explain with neat diagram Symmetric Shared Memory Architectures</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:35-42	C409.1	BTL5

5	<p><b>Explain with neat diagram Distributed Shared Memory Architectures (Nov/Dec 2018)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:43-45</p>	C409.1	BTL5
6	<p><b>Explain Cache coherence in Symmetric Shared and Distributed Shared Memory Architectures</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No: 50-55</p>	C409.1	BTL5
7	<p><b>Explain the performance issues of multicore processor. .(Nov/Dec 2017)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:58-64</p>	C409.1	BTL5
8	<p><b>Define cache coherence problem. What are the 2 main approaches to cache coherence? Describe working of snooping cache coherence and explain describe directory based coherence. (Nov/Dec 2017)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:43-45</p>	C409.1	BTL5
9	<p><b>Explain parallel program design (Apr/May 2017,Nov/Dec 2018)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:65-69</p>	C409.1	BTL5
10	<p><b>State and explain Amdahl’s law Outline the steps in designing and building parallel program. Give example (Apr/May 2018)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>	C409.1	BTL5
11	<p><b>Elaborate the classification of computer architecture in parallel computing system (Apr/May 2018)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>	C409.1	BTL5
12	<p><b>Explain Directory Based cache coherence protocol</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>	C409.1	BTL4

13	<b>Generalize the snooping protocol briefly</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.1	BTL6
14	<b>Summarize the Parallelizing the serial program</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.1	BTL5
15	<b>Explain the Shared memory interconnect</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.1	BTL3
16	<b>Highlight the limitations of single core processors and outline how multicore architecture overcome the limitations</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.1	BTL3

## UNIT II PARALLEL PROGRAM CHALLENGES

Performance – Scalability – Synchronization and data sharing – Data races – Synchronization primitives (mutexes, locks, semaphores, barriers) – deadlocks and livelocks – communication between threads (condition variables, signals, message queues and pipes).

### PART A

Q. No.	Questions	CO	Bloom's Level
1	<b>What is data race?</b> A data race occurs when multiple threads use the same data item and one or more of those threads are updating it.	C409.2	BTL1
2	<b>How to avoid data races .</b> One way to avoid data race is by utilizing proper synchronization between threads.	C409.2	BTL1
3	<b>Hoe to Avoid Data Races.</b> Although it can be hard to identify data races, avoiding them can be very simple: Make sure that only one thread can update the variable at a time. The easiest way to do this is to place a <i>synchronization lock</i> around all accesses to that variable and ensure that before referencing the variable, the thread must acquire the lock.	C409.2	BTL1
4	<b>What is the use of Synchronization Primitives? List out the various synchronization primitive in parallel programming (Nov/Dec 2018)</b> Synchronization is used to coordinate the activity of multiple threads. There are various situations where it is necessary; this might be to ensure that shared resources are not accessed by multiple threads simultaneously or that all work on those resources is complete before new work starts. Process Synchronization Memory Synchronization Thread Synchronization	C409.2	BTL1
5	<b>What is the simplest form of synchronization?</b> The simplest form of synchronization is a mutually exclusive ( <i>mutex</i> ) lock. Only one thread at a time can acquire a mutex lock, so they can be placed around a data structure to ensure that the data structure is modified by only one thread at a time.	C409.2	BTL1
6	<b>How to Place Mutex Locks Around Accesses to Variables.</b> int counter; mutex_lock mutex;	C409.2	BTL1

	<pre> void Increment() { <b>acquire( &amp;mutex );</b> counter++; <b>release( &amp;mutex );</b>} void Decrement() { <b>acquire( &amp;mutex );</b> counter--; <b>release( &amp;mutex );</b>} </pre>		
7	<p><b>What is meant by <i>contended</i> mutex.</b></p> <p>If multiple threads are attempting to acquire the same mutex at the same time, then only one thread will succeed, and the other threads will have to wait. This situation is known as a <i>contended</i> mutex.</p>	C409.2	BTL1
8	<p><b>What is critical region.</b></p> <p>The region of code between the acquisition and release of a mutex lock is called a <i>critical section</i>, or <i>critical region</i>. Code in this region will be executed by only one thread at a time.</p>	C409.2	BTL1
9	<p><b>Develop the code for Placing a Mutex Lock Around a Region of Code</b></p> <pre> void * threadSafeMalloc( size_t size ) { <b>acquire( &amp;mallocMutex );</b> void * memory = malloc( size ); <b>release( &amp;mallocMutex );</b> return memory; } </pre>	C409.2	BTL6
10	<p><b>What is meant by Spin Locks.</b></p> <p><i>Spin locks</i> are essentially mutex locks. The thread waiting to acquire a spin lock will keep trying to acquire the lock without sleeping</p>	C409.2	BTL1
11	<p><b>Compare spin lock and mutex lock. (Apr/May 2018)</b></p> <p>The difference between a mutex lock and a spin lock is that a thread waiting to acquire a spin lock will keep trying to acquire the lock without sleeping .In comparison; a mutex lock may sleep if it is unable to acquire the lock.</p>	C409.2	BTL2
12	<p><b>Write the advantage of spin locks.</b></p> <p>The advantage of using spin locks is that they will acquire the lock as soon as it is released, whereas a mutex lock will need to be woken by the operating system before it can get the lock</p>	C409.2	BTL1
13	<p><b>What is the disadvantage of spin locks.</b></p> <p>The disadvantage is that a spin lock will spin on a virtual CPU monopolizing that resource. In comparison, a mutex lock will sleep and free the virtual CPU for another thread to use.</p>	C409.2	BTL1

14	<p><b>What is Semaphores?</b></p> <p><i>Semaphores</i> are counters that can be either incremented or decremented. They can be used in situations where there is a finite limit to a resource and a mechanism is needed to impose that limit. Semaphores will also signal or wake up threads that are waiting on them to use available resources</p>	C409.2	BTL1
15	<p><b>What is an example for semaphore.</b></p> <p>An example might be a buffer that has a fixed size. Every time an element is added to a buffer, the number of available positions is decreased. Every time an element is removed, the number available is increased</p>	C409.2	BTL1
16	<p><b>What is wait and release in semaphore.</b></p> <p>the method that acquires a semaphore might be called <i>wait, down, or acquire</i>, and the method to release a semaphore might be called <i>post, up, signal, or release</i>. When the semaphore no longer has resources available, the threads requesting resources will block until resources are available.</p>	C409.2	BTL1
17	<p><b>Define readerswriter lock.</b></p> <p>A <i>readerswriter lock</i> (or <i>multiple-reader lock</i>) allows many threads to read the shared data but can then lock the readers threads out to allow one thread to acquire a writer lock to modify the data.</p>	C409.2	BTL1
18	<p><b>Show an example for Readers-Writer Lock.</b></p> <pre>int readData( int cell1, int cell2 ) {     <b>acquireReaderLock( &amp;lock );</b>     int result = data[cell1] + data[cell2];     <b>releaseReaderLock( &amp;lock );</b>     return result; } void writeData( int cell1, int cell2, int value ) {     <b>acquireWriterLock( &amp;lock );</b>     data[cell1] += value;     data[cell2] -= value;     <b>releaseWriterLock( &amp;lock);</b> }</pre>	C409.2	BTL1
19	<p><b>What are the use of Barriers.</b></p> <p>There are situations where a number of threads have to all complete their work before any of the threads can start on the next task. In these situations, it is useful to have a barrier where the threads will wait until all are present</p>	C409.2	BTL1
20	<p><b>Show One common example of using a barrier.</b></p> <p>One common example of using a barrier arises when there is a dependence between different sections of code. For example, suppose a number of threads compute the values stored in a matrix. The variable total needs to be calculated using the values stored in the matrix. A barrier can be used to ensure that all the threads complete their computation of the matrix</p>	C409.2	BTL1

	<p>before the variable total is calculated .ZThe following example shows a situation using a barrier to separate the calculation of a variable from its use.</p> <pre> Compute_values_held_in_matrix(); Barrier(); total = Calculate_value_from_matrix();                     </pre>				
21	<p><b>What is data sharing. (Apr/May 2017)</b> Sharing data between multiple threads is called data sharing.</p>	C409.2	BTL1		
22	<p><b>What are the difference between deadlock and livelock. (Apr/May 2017)(Nov/Dec 2018)</b> The <i>deadlock occurs</i> where two or more threads cannot make progress because the resources that they need are held by the other threads. Example: Suppose two threads need to acquire mutex locks A and B to complete some task. If thread 1 has already acquired lock A and thread 2 has already acquired lock B, then A cannot make forward progress because it is waiting for lock B, and thread 2 cannot make progress because it is waiting for lock A. The two threads are <i>deadlocked</i>.</p> <p><u>Listing 4.13 Two Threads in a Deadlock</u></p> <table border="0" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; vertical-align: top; border-right: 1px solid black; padding-right: 10px;"> <p><b>Thread 1</b></p> <pre> void update1() {   acquire(A);   acquire(B); &lt;&lt;&lt; Thread 1                 waits here    variable1++;   release(B);   release(A); }                     </pre> </td> <td style="width: 50%; vertical-align: top; padding-left: 10px;"> <p><b>Thread 2</b></p> <pre> void update2() {   acquire(B);   acquire(A); &lt;&lt;&lt; Thread 2                 waits here    variable2++;   release(B);   release(A); }                     </pre> </td> </tr> </table>	<p><b>Thread 1</b></p> <pre> void update1() {   acquire(A);   acquire(B); &lt;&lt;&lt; Thread 1                 waits here    variable1++;   release(B);   release(A); }                     </pre>	<p><b>Thread 2</b></p> <pre> void update2() {   acquire(B);   acquire(A); &lt;&lt;&lt; Thread 2                 waits here    variable2++;   release(B);   release(A); }                     </pre>	C409.2	BTL1
<p><b>Thread 1</b></p> <pre> void update1() {   acquire(A);   acquire(B); &lt;&lt;&lt; Thread 1                 waits here    variable1++;   release(B);   release(A); }                     </pre>	<p><b>Thread 2</b></p> <pre> void update2() {   acquire(B);   acquire(A); &lt;&lt;&lt; Thread 2                 waits here    variable2++;   release(B);   release(A); }                     </pre>				
23	<p><b>What are conditions under which a deadlock situation may arise? (Nov/Dec 2017)</b></p> <p><b>Mutual Exclusion:</b> At least one resource is held in a non-sharable mode that is only one process at a time can use the resource. If another process requests that resource, the requesting process must be delayed until the resource has been released.</p> <p><b>Hold and Wait:</b> There must exist a process that is holding at least one resource and is waiting to acquire additional resources that are currently being held by other processes.</p> <p><b>No Preemption:</b> Resources cannot be preempted; that is, a resource can only be released voluntarily by the process holding it, after the process has completed its task.</p> <p><b>Circular Wait:</b> There must exist a set <math>\{p_0, p_1, \dots, p_n\}</math> of waiting processes such that <math>p_0</math> is waiting for a resource which is held by <math>p_1</math>, <math>p_1</math> is waiting for a resource which is held by <math>p_2, \dots, p_{n-1}</math> is waiting for a resource which is held by <math>p_n</math> and <math>p_n</math> is waiting for a resource which is held by <math>p_0</math>.</p>	C409.2	BTL1		

24	<p><b>Define thread. mention the uses of swapping. (Nov/Dec 2017)</b>  A thread is the smallest unit of processing that can be performed in an OS.  In most modern operating systems, a thread exists within a process - that is, a single process may contain multiple threads</p>	C409.2	BTL1
25	<p><b>Define deadlock.</b>  <i>Deadlock</i> is the situation where two or more threads cannot make progress because the resources that they need are held by the other threads. It is easiest to explain this with an example. Suppose two threads need to acquire mutex locks A and B to complete some task. If thread 1 has already acquired lock A and thread 2 has already acquired lock B, then A cannot make forward progress because it is waiting for lock B, and thread 2 cannot make progress because it is waiting for lock A. The two threads are <i>deadlocked</i></p>	C409.2	BTL1
26	<p><b>How to communicate multiple threads.</b>  The easiest way for multiple threads to communicate is through memory. If two threads can access the same memory location, the cost of that access is little more than the memory latency of the system. Of course, memory accesses still need to be controlled to ensure that only one thread writes to the same memory location at a time. A multithreaded application will share memory between the threads by default, so this can be a very low-cost approach. The only things that are not shared between threads are variables on the stack of each thread (local variables) and thread-local variables.</p>	C409.2	BTL1
27	<p><b>Illustrate an example which Use Multiple Barriers.</b>  Compute_values_held_in_matrix();  <b>Barrier();</b>  total = Calculate_value_from_matrix();  <b>Barrier();</b>  Perform_next_calculation( total );</p>	C409.2	BTL2
28	<p><b>Define live lock.</b>  A <i>livelock</i> traps threads in an unending loop releasing and acquiring locks. Livelocks can be caused by code to back out of deadlocks.</p>	C409.2	BTL1
29.	<p><b>What is An atomic operation</b>  An <i>atomic operation</i> is one that will either successfully complete or fail; it is not possible for the operation to either result in a “bad” value or allow other threads on the system to observe a transient value.</p>	C409.2	BTL1
30	<p><b>Write down the performance metrics (Apr/May 2018)</b></p>	C409.2	BTL1

31	<p><b>Define Message Queue</b></p> <p>A message queue is a structure that can be shared between multiple processes. Messages can be placed into the queue and will be removed in the same order in which they were added. Constructing a message queue looks rather like constructing a shared memory segment.</p>	C409.2	BTL1
32	<p><b>Define Named Pipes</b></p> <p>UNIX uses pipes to pass data from one process to another. For example, the output from the command <code>ls</code>, which lists all the files in a directory, could be piped into the <code>wc</code> command, which counts the number of lines, words, and characters in the input</p>	C409.2	BTL1
33	<p><b>Mention the mechanism associated with Named Pipes</b></p> <p>Setting Up and Writing into a Pipe <code>Make Pipe( Descriptor ); ID = Open Pipe(Descriptor ); Write Pipe( ID, Message, sizeof(Message) ); ... Close Pipe( ID ); Delete Pipe( Descriptor );</code>  Opening an Existing Pipe to Receive Messages <code>ID=Open Pipe( Descriptor ); Read Pipe( ID, buffer, sizeof(buffer) ); Close Pipe( ID );</code></p>	C409.2	BTL1
34	<p><b>How to create and Place Message Queues</b></p> <p>Creating and Placing Messages into a Queue <code>ID = Open Message Queue Queue(Descriptor ); Put Message in Queue( ID, Message ); Close Message Queue( ID );</code>  Delete Message Queue( Description );  Using the descriptor for an existing message queue enables two processes to communicate by sending and receiving messages through the queue.  Opening a Queue and Receiving Messages <code>ID=Open Message Queue ID(Descriptor);</code>  Message=<code>Remove Message from Queue(ID); ... Close Message Queue(ID);</code></p>	C409.2	BTL1
35	<p><b>What is the fundamental way to share access to resources between threads</b></p> <ul style="list-style-type: none"> <li>• Deadlock</li> <li>• Livelock</li> </ul>	C409.2	BTL1
36	<p><b>Give an example of critical regions</b></p> <p>An operating system does not have an implementation of <code>malloc()</code> that is thread-safe, or safe for multiple threads to call at the same time. One way to fix this is to place the call to <code>malloc()</code> in a critical section by surrounding it with a mutex lock</p>	C409.2	BTL1
37	<p><b>List out the issues in shared caches</b></p> <ul style="list-style-type: none"> <li>• Capacity misses</li> <li>• Conflict misses</li> </ul>	C409.2	BTL1
38	<p><b>Define False sharing</b></p> <p><i>False sharing</i> is the situation where multiple threads are accessing items of data held on a single cache line.  Although the threads are all using separate items of data, the cache line itself is shared between them so only a single thread can write to it at any one time.</p>	C409.2	BTL1
39	<p><b>List out the Memory Bandwidth Measured on a System with Four Virtual CPUs</b></p>	C409.2	BTL1

	<p>Threads Time 7.437563 s Bandwidth 2.63 GB/s                  Threads Time 15.238317 s Bandwidth 2.57 GB/s                  Threads Time 24.580981 s Bandwidth 2.39 GB/s                  Threads Time 37.457352 s Bandwidth 2.09 GB/s</p>		
40	<p><b>What are the measures to be taken when bandwidth size reduces</b>                  The threads are interfering on the processor.</p> <ul style="list-style-type: none"> <li>• A second interaction effect is if the threads start interfering in the caches, such as multiple threads attempting to load data to the same set of cache lines.</li> <li>• One other effect is the behavior of memory chips when they become saturated. The chips start experiencing queuing latencies where the response time for each request increases. Memory chips are arranged in banks.</li> <li>• Accessing a particular address will lead to a request to a particular bank of memory. Each bank needs a gap between returning two responses. If multiple threads happen to hit the same bank, then the response time becomes governed by the rate at which the bank can return memory</li> </ul>	C409.2	BTL1
41	<p><b>Write a code to measure memory bandwidth using Memset</b></p> <pre>#include &lt;stdio.h&gt; #include &lt;stdlib.h&gt; #include &lt;strings.h&gt; #include &lt;pthread.h&gt; #include &lt;sys/time.h&gt; #define BLOCKSIZE 1024*1025 int nthreads = 8; char * memory; double now( ) {     struct timeval time;     gettimeofday( &amp;time, 0 );     return (double)time.tv_sec + (double)time.tv_usec / 1000000.0; } void *experiment( void *id ) {     unsigned int seed = 0;     int count = 20000;     for( int i=0; i&lt;count; i++ )     {         <b>memset( &amp;memory[BLOCKSIZE * (int)id], 0, BLOCKSIZE );</b>     } }</pre>	C409.2	BTL1
42	<p><b>How Bandwidth share between cores</b>                  Bandwidth is another resource shared between threads. The bandwidth capacity of a system depends on the design of the processor and the memory system as well as the memory chips and their location in the system</p>	C409.2	BTL1

43	<p><b>List out the three critical areas to address large difference scaling</b></p> <ul style="list-style-type: none"> <li>• The amount of bandwidth to cache and the memory will be divided among the active threads on the system.</li> <li>• The design of the caches will determine how much time is lost because of capacity and conflict-induced cache misses.</li> <li>• The way that the processor core pipelines are shared between active software threads will determine how instruction issue rates change as the number of active threads increases.</li> </ul>	C409.2	BTL1
44	<p><b>What is the role default malloc()</b> The default malloc() provides better performance than the alternative implementation. The algorithm that provides improved scaling also adds a cost to the single-threaded situation; it can be hard to produce an algorithm that is fast for the single- threaded case and scales well with multiple threads.</p>	C409.2	BTL1
45	<p><b>Define an idea to choose the appropriate data structures</b> Choosing the best structure to hold data, such as choosing an algorithm of the appropriate complexity, can have a major impact on overall performance. Some structures will be efficient when data is accessed in one pattern, while other structures will be more efficient if the access pattern is changed.</p>	C409.2	BTL1
46	<p><b>Define Column major order</b> The opposite ordering is followed, so adjacent elements of the first index are adjacent in memory. This is called <i>column-major</i> order. Accessing elements by a stride is a common error in codes translated from Fortran into C. It shows how memory is addressed in C, where adjacent elements in a row are adjacent in memory.</p>	C409.2	BTL1
47	<p><b>How to select the appropriate array access pattern</b> One common data access pattern is striding through elements of an array. The performance of the application would be better if the array could be arranged so that the selected elements were contiguous.</p>	C409.2	BTL1
48	<p><b>List out the techniques to reduce the latency</b></p> <ul style="list-style-type: none"> <li>• Out of order execution</li> <li>• Hardware prefetching</li> <li>• Software prefetching</li> </ul>	C409.2	BTL1
49	<p><b>List out the non technical reasons why functionality get placed in libraries</b></p> <ul style="list-style-type: none"> <li>• Libraries often represent a convenient product for an organizational unit. One group of developers might be responsible for a particular library of code, but that does not automatically imply that a single library represents the best way for that code to be delivered to the end users.</li> <li>• Libraries are also used to group related functionality. For example, an application might contain a library of string-handling functions. Such a library might be appropriate if it contains a large body of code. On the other hand, if it contains only a few small routines, it might be more appropriate to combine it with another library.</li> </ul>	C409.2	BTL1

50	<b>Why Algorithm complexity is important</b> Algorithmic complexity represents the expected performance of a section of code as the number of elements being processed increases. In the limit, the code with the greatest algorithmic complexity will dominate the runtime of the application	C409.2	BTL1
----	---	--------	------

## PART B

Q. No.	Questions	CO	Bloom's Level
1.	<b>Explain about Synchronization and data sharing in detail.</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:121-126	C409.2	BTL5
2.	<b>Explain Synchronization primitives mutexes and locks.</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:126-128	C409.2	BTL5
3.	<b>Explain Synchronization primitives in semaphores and barriers in detail.</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:128-129	C409.2	BTL5
4.	<b>Explain the concepts of deadlocks and live locks</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:132-133	C409.2	BTL5
5.	<b>Explain communication between threads using condition variables and signals.</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:133-139	C409.2	BTL5
6.	<b>Explain communication between threads using message queues and pipes.</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:138-139	C409.2	BTL5
7.	<b>Explain data races and scalability in parallel program. (apr/may2017)</b> 2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:121-126	C409.2	BTL5

8.	<b>Explain Synchronization primitives in parallel program challenges. (Apr/may2017)</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:126-130	C409.2	BTL5
9	<b>Explain the various approaches to parallel programming. .(Nov/Dec 2017)</b>  1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:43-47	C409.2	BTL5
10.	<b>What is a data race? What are the tools used for detecting data races? How to avoid races? (Nov/Dec 2017) (Apr/May 2018)</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011 Pg.No:121-125	C409.2	BTL5
11	<b>(i)Discuss in detail about producer consumer synchronization (ii)Write a simple semaphore to send a message (Apr/May 2018)</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011	C409.2	BTL5
12	<b>Write a short notes on deadlocks, livelocks and named pipes</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011	C409.2	BTL5
13	<b>Discuss in detail about the importance of algorithmic complexity</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011	C409.2	BTL2
14	<b>Explain the outline about necessity of structure reflects in performance</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011	C409.2	BTL4
15	<b>Write in detail and summarize about hardware constraints applicable in improving scaling</b>  2. Darryl Gove, "Multicore Application Programming for Windows, Linux, and Oracle Solaris", Pearson, 2011	C409.2	BTL5

## UNIT III

## SHARED MEMORY PROGRAMMING WITH OpenMP

OpenMP Execution Model – Memory Model – OpenMP Directives – Work-sharing Constructs – Library functions – Handling Data and Functional Parallelism – Handling Loops – Performance Considerations.

Q. No.	Questions	CO	Bloom's Level
1.	<p><b>What is OpenMP?</b></p> <p>Like Pthreads, OpenMP is an API for shared-memory parallel programming. The “MP” in OpenMP stands for “multiprocessing,” a term that is synonymous with shared-memory parallel computing. Thus, OpenMP is designed for systems in which each thread or process can potentially have access to all available memory, and, when we’re programming with OpenMP, we view our system as a collection of cores or CPUs, all of which have access to main memory.</p>	C409.3	BTL1
2.	<p><b>What is Pthread.</b></p> <p>Pthreads is lower level and provides us with the power to program virtually any conceivable thread behavior. This power, however, comes with some associated cost—it’s up to us to specify every detail of the behavior of each thread.</p>	C409.3	BTL1
3.	<p><b>What the difference between Pthreads and OpenMP.</b></p> <p>Pthreads requires that the programmer explicitly specify the behavior of each thread. OpenMP, on the other hand, sometimes allows the programmer to simply state that a block of code should be executed in parallel, and the precise determination of the tasks and which thread should execute them is left to the compiler and the run-time system. This suggests a further difference between OpenMP and Pthreads, that is, that Pthreads (like MPI) is a library of functions that can be linked to a C program, so any Pthreads program can be used with any C compiler, provided the system has a Pthreads library. OpenMP, on the other hand, requires compiler support for some operations, and hence it’s entirely possible that you may run across a C compiler that can’t compile OpenMP programs into parallel programs.</p>	C409.3	BTL1
4.	<p><b>What is the Execution Model of OpenMp.</b></p> <p>The OpenMP API uses the fork-join model of parallel execution. Multiple threads of execution perform tasks defined implicitly or explicitly by OpenMP directives. The OpenMP API is intended to support programs that will execute correctly both as parallel programs (multiple threads of</p>	C409.3	BTL1

	execution and a full OpenMP support library) and as sequential programs (directives ignored and a simple OpenMP stubs library).		
5.	<b>What is an initial thread in OpenMP?</b> An OpenMP program begins as a single thread of execution, called an initial thread.	C409.3	BTL1
6.	<b>How to compile and running OpenMP programs</b> To compile this with gcc we need to include the -fopenmp option \$ gcc -g -Wall -fopenmp -o omp_hello omp_hello.c To run the program, we specify the number of threads on the command line. For example, we might run the program with four threads and type \$ ./omp_hello 4		
7	<b>What is termed as initial task region. .(Nov/Dec 2017)</b> The initial task region, that is defined by an implicit inactive parallel region surrounding the whole program. When any thread encounters a parallel construct, the thread creates a team of itself and zero or more additional threads and becomes the master of the new team. A set of implicit tasks, one per thread, is generated	C409.3	BTL1
8	<b>Define odd even transportation sort? . (Apr/May 2017)</b> <ul style="list-style-type: none"> <li>Parallelizable version of Bubble sort</li> <li>Requires N passes through the array.</li> <li>Each pass through the array analyzes either: <ul style="list-style-type: none"> <li>Every pair of odd indexed elements and the preceding element, or</li> <li>Every pair of even indexed elements and the preceding element.</li> </ul> </li> <li>Within each pass, elements that are not in order are swapped.</li> </ul> <pre>Sort local keys: for (phase = 0; phase &lt; comm_sz; phase++) {     partner = Compute_partner(phase, my_rank);     if (I'm not idle) {         Send my keys to partner;         Receive keys from partner;         if (my_rank &lt; partner)             Keep smaller keys;         else             Keep larger keys;     } }</pre>	C409.3	BTL1
9.	<b>Develop a "hello word" program in that uses open MP. . (Apr/May 2017)</b> Hello world program in C using MPI: <pre>#include &lt;stdio.h&gt;</pre>	C409.3	BTL1

	<pre> #include &lt;mpi.h&gt;  main(int argc, char **argv) {     int ierr;      ierr = MPI_Init(&amp;argc, &amp;argv);     printf("Hello world\n");      ierr = MPI_Finalize(); } </pre>		
10	<p><b>Define Message Queue.(Nov/Dec 2017)</b></p> <p>Message queues provide an asynchronous communications protocol, meaning that the sender and receiver of the message do not need to interact with the message queue at the same time. Messages placed onto the queue are stored until the recipient retrieves them.</p>	C409.3	BTL1
11.	<p><b>What is an initial task region?</b></p> <p>An initial thread executes sequentially, as if enclosed in an implicit task region, called an initial task region, that is defined by the implicit parallel region surrounding the whole program.</p>	C409.3	BTL1
12	<p><b>Discuss the Structure of the OpenMP Memory Model</b></p> <p>The OpenMP API provides a relaxed-consistency, shared-memory model. All OpenMP threads have access to a place to store and to retrieve variables, called the <i>memory</i>. In addition, each thread is allowed to have its own <i>temporary view</i> of the memory.</p>	C409.3	BTL1
13	<p><b>What is <i>threadprivate</i> memory?</b></p> <p>Each thread also has access to another type of memory that must not be accessed by other threads, called <i>threadprivate memory</i>.</p>	C409.3	BTL1
14	<p><b>Show the format of directive in OpenMP.</b></p> <p><b>#pragma omp directive-name [clause[ [, ] clause]...] new-line</b></p> <p>Each directive starts with <b>#pragma omp</b>. The remainder of the directive follows the conventions of the C and C++ standards for compiler directives. In particular, white space can be used before and after the #, and sometimes white space must be used to separate the words in a directive.</p>	C409.3	BTL2
15	<p><b>What is meant by Stand-alone directives?</b></p> <p>Stand-alone directives do not have any associated executable user code. Instead, they represent executable statements that typically do not have succinct equivalent statements in the base languages. There are some restrictions on the placement of a stand-alone directive within a program. A stand-alone directive may be placed only at a point where a base language executable statement is allowed</p>	C409.3	BTL1
16	<p><b>What is the use of parallel construct?</b></p> <p>This fundamental construct starts parallel execution.</p> <p><b>#pragma omp parallel [clause[ [, ] clause] ...] new-line structured-block</b></p> <p>where <i>clause</i> is one of the following:</p>	C409.3	BTL1

	<p><b>#pragma omp parallel</b> [<i>clause</i>[ [, ]<i>clause</i>] ...] <i>new-line</i>  <i>structured-block</i>  <b>if</b>(<i>scalar-expression</i>)  <b>num_threads</b>(<i>integer-expression</i>)  <b>default</b>(<b>shared</b>   <b>none</b>)  <b>private</b>(<i>list</i>)  <b>firstprivate</b>(<i>list</i>)  <b>shared</b>(<i>list</i>)  <b>copyin</b>(<i>list</i>)  <b>reduction</b>(<i>reduction-identifier</i> :<i>list</i>)  <b>proc_bind</b>(<b>master</b>   <b>close</b>   <b>spread</b>)</p>		
17	<p><b>What is Worksharing Constructs?</b>  A worksharing construct distributes the execution of the associated region among the members of the team that encounters it. Threads execute portions of the region in the context of the implicit tasks each one is executing. If the team consists of only one thread then the worksharing region is not executed in parallel.</p>	C409.3	BTL1
18	<p><b>List some worksharing constructs.</b>  The OpenMP API defines the following worksharing constructs.</p> <ul style="list-style-type: none"> <li>• loop construct</li> <li>• <b>sections</b> construct</li> <li>• <b>single</b> construct</li> <li>• <b>workshare</b> construct</li> </ul>	C409.3	BTL4
19	<p><b>List the the Restrictions apply to worksharing constructs.</b>  The following restrictions apply to worksharing constructs:</p> <ul style="list-style-type: none"> <li>• Each worksharing region must be encountered by all threads in a team or by none at all, unless cancellation has been requested for the innermost enclosing parallel region.</li> <li>• The sequence of worksharing regions and <b>barrier</b> regions encountered must be the same for every thread in a team.</li> </ul>	C409.3	BTL4
20	<p><b>Show the syntax of the loop construct.</b></p> <p>The syntax of the loop construct is as follows:  <b>#pragma omp for</b> [<i>clause</i>[[,] <i>clause</i>] ... ] <i>new-line</i>  <i>for-loops</i>  where <i>clause</i> is one of the following:  <b>private</b>(<i>list</i>)  <b>firstprivate</b>(<i>list</i>)  <b>lastprivate</b>(<i>list</i>)  <b>reduction</b>(<i>reduction-identifier</i>: <i>list</i>)  <b>schedule</b>(<i>kind</i>[, <i>chunk_size</i>])  <b>collapse</b>(<i>n</i>)  <b>ordered</b>  <b>nowait</b></p>	C409.3	BTL1
21	<p><b>What is meant by binding?</b>  The binding thread set for a loop region is the current team. A loop region</p>	C409.3	BTL1

	binds to the innermost enclosing <b>parallel</b> region. Only the threads of the team executing the binding <b>parallel</b> region participate in the execution of the loop iterations and the implied barrier of the loop region if the barrier is not eliminated by a <b>nowait</b> clause		
22	<p><b>What is the use of collapse clause.</b></p> <p>The <b>collapse</b> clause may be used to specify how many loops are associated with the loop construct. The parameter of the <b>collapse</b> clause must be a constant positive integer expression. If no <b>collapse</b> clause is present, the only loop that is associated with the loop construct is the one that immediately follows the loop directive.</p>	C409.3	BTL1
23	<p><b>List the Restrictions to the loop construct.</b></p> <p>Restrictions to the loop construct are as follows:</p> <ul style="list-style-type: none"> <li>• All loops associated with the loop construct must be perfectly nested; that is, there must be no intervening code nor any OpenMP directive between any two loops.</li> <li>• The values of the loop control expressions of the loops associated with the loop construct must be the same for all the threads in the team.</li> <li>• Only one <b>schedule</b> clause can appear on a loop directive.</li> <li>• Only one <b>collapse</b> clause can appear on a loop directive.</li> <li>• <i>chunk_size</i> must be a loop invariant integer expression with a positive value.</li> <li>• The value of the <i>chunk_size</i> expression must be the same for all threads in the team.</li> <li>• The value of the <i>run-sched-var</i> ICV must be the same for all threads in the team.</li> <li>• When <b>schedule(runtime)</b> or <b>schedule(auto)</b> is specified, <i>chunk_size</i> must not be specified.</li> <li>• Only one <b>ordered</b> clause can appear on a loop directive.</li> <li>• The <b>ordered</b> clause must be present on the loop construct if any <b>ordered</b> region ever binds to a loop region arising from the loop construct.</li> <li>• The loop iteration variable may not appear in a <b>threadprivate</b> directive</li> </ul>	C409.3	BTL4
24	<p><b>How to Determine the Schedule of a Worksharing Loop?</b></p> <p>When execution encounters a loop directive, the <b>schedule</b> clause (if any) on the directive, and the <i>run-sched-var</i> and <i>def-sched-var</i> ICVs are used to determine how loop iterations are assigned to threads.. If the loop directive does not have a <b>schedule</b> clause then the current value of the <i>def-sched-var</i> ICV determines the schedule. If the loop directive has a <b>schedule</b> clause that specifies the <b>runtime</b> schedule kind then the current value of the <i>run-sched-var</i> ICV determines the schedule. Otherwise, the value of the <b>schedule</b> clause determines the schedule.</p>	C409.3	BTL1
25	<p><b>What is The sections construct.</b></p> <p>The <b>sections</b> construct is a non-iterative work sharing construct that</p>	C409.3	BTL1

	contains a set of structured blocks that are to be distributed among and executed by the threads in a team. Each structured block is executed once by one of the threads in the team in the context of its implicit task.		
26	<p><b>Show the the syntax of the sections construct.</b></p> <p>The syntax of the <b>sections</b> construct is as follows:</p> <pre>#pragma omp sections [clause[[,] clause] ...] new-line {[ #pragma omp section new-line] structured-block [#pragma omp section new-line structured-block ] ... }</pre> <p>where <i>clause</i> is one of the following</p> <p><b>private(list)</b>  <b>firstprivate(list)</b>  <b>lastprivate(list)</b>  <b>reduction(reduction-identifier:list)</b>  <b>nowait</b></p>	C409.3	BTL1
27	<p><b>List the Restrictions to the sections construct .</b></p> <ul style="list-style-type: none"> <li>• Orphaned <b>section</b> directives are prohibited. That is, the <b>section</b> directives must appear within the <b>sections</b> construct and must not be encountered elsewhere in the <b>sections</b> region.</li> <li>• The code enclosed in a <b>sections</b> construct must be a structured block.</li> <li>• Only a single <b>nowait</b> clause can appear on a <b>sections</b> directive.</li> </ul> <p><b>firstprivate(list)</b>  <b>lastprivate(list)</b>  <b>reduction(reduction-identifier:list)</b></p> <ul style="list-style-type: none"> <li>• A throw executed inside a <b>sections</b> region must cause execution to resume within the same section of the <b>sections</b> region, and the same thread that threw the exception must catch it.</li> </ul>	C409.3	BTL4
28	<p><b>Show The syntax of the single construct .</b></p> <pre>!\$omp single [clause[[,] clause] ...] structured-block !\$omp end single [end_clause[[,] end_clause] ...]</pre> <p>where <i>clause</i> is one of the following:</p> <p><b>private(list)</b>  <b>firstprivate(list)</b></p> <p>and <i>end_clause</i> is one of the following:</p>	C409.3	BTL1

	<b>copyprivate(list) nowait</b>		
29	<p><b>Show the syntax of the workshare construct.</b></p> <p><b>!\$omp workshare</b> <i>structured-block</i></p> <p><b>!\$omp end workshare [nowait]</b></p> <p>The enclosed structured block must consist of only the following:</p> <ul style="list-style-type: none"> <li>• array assignments</li> <li>• scalar assignments</li> <li>• <b>FORALL</b> statements</li> <li>• <b>FORALL</b> constructs</li> <li>• <b>WHERE</b> statements</li> <li>• <b>WHERE</b> constructs</li> <li>• <b>atomic</b> constructs</li> </ul>	C409.3	BTL1
30	<p><b>List the restrictions apply to the workshare construct:</b></p> <ul style="list-style-type: none"> <li>• All array assignments, scalar assignments, and masked array assignments must be intrinsic assignments.</li> <li>• The construct must not contain any user defined function calls unless the function is <b>ELEMENTAL</b></li> </ul>	C409.3	BTL4
31	<p><b>What is the use of schedule clause.</b></p> <p>If more than one loop is associated with the loop construct, then the iterations of all associated loops are collapsed into one larger iteration space that is then divided according to the <b>schedule</b> clause. The sequential execution of the iterations in all associated loops determines the order of the iterations in the collapsed iteration space</p>	C409.3	BTL1
32	<p><b>What is the use of single construct.</b></p> <p>The <b>single</b> construct specifies that the associated structured block is executed by only one of the threads in the team (not necessarily the master thread), in the context of its implicit task. The other threads in the team, which do not execute the block, wait at an implicit barrier at the end of the <b>single</b> construct unless a <b>nowait</b> clause is specified</p>	C409.3	BTL1
33	<p><b>List the Restrictions to the single construct are as follows:</b></p> <ul style="list-style-type: none"> <li>• The <b>copyprivate</b> clause must not be used with the <b>nowait</b> clause.</li> <li>• At most one <b>nowait</b> clause can appear on a <b>single</b> construct.</li> </ul>	C409.3	BTL4
34	<p><b>What is workshare construct?</b></p> <p>The <b>workshare</b> construct divides the execution of the enclosed structured block into separate units of work, and causes the threads of the team to share the work such that each unit is executed only once by one thread, in the context of its implicit task.</p>	C409.3	BTL1
35	<b>Explain Scope of a variable (Apr/May 2018)</b>	C409.3	BTL1

36	<b>Define Race Condition (Apr/May 2018)</b>	C409.3	BTL1
37	<b>State the trapezoidal rule in OpenMP (Nov/Dec 2018)</b> <pre>for (i = 1; i &lt;= n-1; i++) { x_i = a + i*dx; approx += f(x_i); } approx = dx*approx;</pre>	C409.3	BTL1
38	<b>Define Coherence and consistency</b> <ul style="list-style-type: none"> <li>• Coherence refers to the behavior of the memory system when a single memory location is accessed by multiple threads.</li> <li>• Consistency refers to the ordering of accesses to different memory locations, observable from various threads in the system</li> </ul>	C409.3	BTL1
39	<b>Define data race</b> <ul style="list-style-type: none"> <li>• A data race is defined to be accesses to a single variable by at least two threads, at least one of which is a write, not separated by a synchronization operation.</li> <li>• OpenMP does guarantee certain consistency behavior, however. That behavior is based on the OpenMP flush operation.</li> </ul>	C409.3	BTL1
40	<b>Define OpenMP flush operation</b> The OpenMP flush operation is applied to a set of variables called the flush set. Memory operations for variables in the flush set that precede the flush in program execution order must be firmly lodged in memory and available to all threads before the flush completes, and memory operations for variables in the flush set, that follow a flush in program order cannot start until the flush completes.		
41	<b>Mention the actions to be taken to move the value from one thread to another thread</b> <ul style="list-style-type: none"> <li>• The first thread writes the value to the shared variable,</li> <li>• The first thread flushes the variable</li> <li>• The second thread flushes the variable and</li> <li>• The second thread reads the variable</li> </ul>	C409.3	BTL1
42	<b>List out the two methods available for enabling nested parallel regions</b> 1.The <b>omp_set_nested()</b> library routine 2. Setting of the <b>OMP_NESTED</b> environment variable to TRUE	C409.3	BTL1
43	<b>What is data dependences</b> 1. OpenMP compilers don't check for dependences among iterations in a loop that's being parallelized with a parallel <b>for</b> directive. It's up to us, the programmers, to identify these dependences. 2. A loop in which the results of one or more iterations depend on other	C409.3	BTL1

	iterations <i>cannot</i> , in general, be correctly parallelized by OpenMP.		
44	<p><b>Define Atomic Directive</b></p> <p>The atomic directive ensures that a specific memory location is updated atomically, rather than exposing it to the possibility of multiple, simultaneous writing threads.</p> <p>The atomic directive supports no OpenMP clauses.</p> <p><b>Syntax</b></p> <p><b>#pragma omp atomic expression</b></p>	C409.3	BTL1
45	<p><b>Define Critical Directive</b></p> <p>Specifies that code is only executed on one thread at a time. OpenMP <i>does</i> provide the option of adding a name to a critical directive:</p> <p><b>Syntax</b></p> <p><b># pragma omp critical(name)</b></p>	C409.3	BTL1
46	<p><b>List out the two types of locks in OpenMP to destroy lock data structures</b></p> <ul style="list-style-type: none"> <li>• Simple locks</li> <li>• Nested Locks</li> </ul>	C409.3	BTL1
47	<p><b>Define Barrier Directive and Master Directive</b></p> <p>A barrier directive will cause the threads in a team to block until all the threads have reached the directive.</p> <p><b>Syntax</b></p> <p><b>#pragma omp barrier</b></p> <p>Specifies that only the master thread should execute a section of the program.</p> <p><b>Syntax</b></p> <p><b># pragma omp master</b></p>	C409.3	BTL1
48	<p><b>Which are the OpenMP clauses supported by Single directive</b></p> <ul style="list-style-type: none"> <li>• copyprivate</li> <li>• firstprivate</li> <li>• nowait</li> <li>• private</li> </ul>	C409.3	BTL1
49	<p><b>Define Synchronization Clauses</b></p> <ul style="list-style-type: none"> <li>• Critical</li> <li>• Atomic</li> <li>• Ordered</li> <li>• Barrier</li> <li>• Nowait</li> </ul>	C409.3	BTL1
50	<p><b>List out the three types of Scheduling?</b></p> <ul style="list-style-type: none"> <li>• Static</li> <li>• Dynamic</li> <li>• Guided</li> </ul>	C409.3	BTL1
51	<p><b>Define Data parallelism</b></p> <p><b>Data parallelism</b> is a form of parallelization across multiple</p>	C409.3	BTL1

	processors in parallel computing environments. It focuses on distributing the data across different nodes, which operate on the data in parallel. It can be applied on regular data structures like arrays and matrices by working on each element in parallel.												
52	<p><b>List out the four discrete steps to parallelization</b></p> <table border="1"> <thead> <tr> <th>Type</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Decomposition</td> <td>The program is broken down into tasks, the smallest exploitable unit of concurrence.</td> </tr> <tr> <td>Assignment</td> <td>Tasks are assigned to processes.</td> </tr> <tr> <td>Orchestration</td> <td>Data access, communication, and synchronization of processes.</td> </tr> <tr> <td>Mapping</td> <td>Processes are bound to processors.</td> </tr> </tbody> </table>	Type	Description	Decomposition	The program is broken down into tasks, the smallest exploitable unit of concurrence.	Assignment	Tasks are assigned to processes.	Orchestration	Data access, communication, and synchronization of processes.	Mapping	Processes are bound to processors.	C409.3	BTL1
Type	Description												
Decomposition	The program is broken down into tasks, the smallest exploitable unit of concurrence.												
Assignment	Tasks are assigned to processes.												
Orchestration	Data access, communication, and synchronization of processes.												
Mapping	Processes are bound to processors.												
53	<p><b>What are Loop Carried Dependencies</b>  <b>Loop-Carried Dependence</b> Verification in <b>OpenMP</b>. Data <b>dependence</b> analysis is a very difficult task, mainly due to the limitations imposed by pointer aliasing, and by the overhead of dynamic data <b>dependence</b> analysis.</p>	C409.3	BTL1										

### PART B

Q. No.	Questions	CO	Bloom's Level
1.	<p><b>Explain OpenMP Execution Model in detail with example.(Nov/Dec 2017) (Apr/May 2018)</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:210-215</p>	C409.3	BTL5
2.	<p><b>Explain the Memory Model of OpenMP. (Apr/May 2018)</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:213-215</p>	C409.3	BTL5
3.	<p><b>Explain the OpenMP Directives . (Apr/may2017)</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:224-231</p>	C409.3	BTL5
4.	<p><b>Explain the Library functions used in OpenMP.</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:</p>	C409.3	BTL5

5.	<b>Explain in detail how to Handle Loops in OpenMP (Nov/Dec 2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No: 236-240	C409.3	BTL5
6.	<b>Explain OpenMP directives. . (Apr/may 2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:224-231	C409.3	BTL5
7.	<b>How data and functional parallelism are handled in shared memory programming with open MP? (Apr/may2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:212-215	C409.3	BTL5
8.	<b>Explain in detail about the handling loops in parallel operations (Nov/Dec 2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:236-240	C409.3	BTL5
9	<b>Write an example program for shared memory programming with Pthread (Apr/May 2018)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:236-240	C409.3	BTL5
10	<b>Explain in detail about the synchronization primitives in parallel program challenges (Apr/May 2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:236-240	C409.3	BTL5

11	<p><b>Explain the type shared memory model</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>		BTL5
12	<p><b>Collect the all information about internal control variable</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>		BTL1
13	<p><b>Explain briefly about General Data Parallelism</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>		BTL4
14	<p><b>(i)Explain the NoWait Clause</b></p> <p><b>(ii)Explain the single pragma</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>		BTL4
15	<p><b>(i)Illustrate the runtime library definitions</b></p> <p><b>(ii)Illustrate the execution environment routines</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>		BTL3

## UNIT IV

**DISTRIBUTED MEMORY PROGRAMMING WITH MPI**

MPI program execution – MPI constructs – libraries – MPI send and receive – Point-to-point and Collective communication – MPI derived datatypes – Performance evaluation.

## PART A

Q. No.	Questions	CO	Bloom's Level
1.	<p><b>What is MPI?</b> One process calls a <i>send</i> function and the other calls a <i>receive</i> function. The implementation of message-passing that will be using is called <b>MPI</b> (Message-Passing Interface). MPI is not a new programming language. It defines a <i>library</i> of functions that can be called from C, C++, and FORTRAN Programs</p>	C409.4	BTL1
2.	<p><b>How to execute MPI programs?</b> Many systems use the command called mpicc to compile and run MPI programs: <b>mpicc -g -Wall -o mpi hello mpi hello.c</b></p>	C409.4	BTL1
3.	<p><b>What is the use of Wrapper Script. (Apr/May 2017)</b> A wrapper script is a script whose main purpose is to run some program. In this case, the program is the C compiler. The wrapper simplifies the running of the compiler by telling it where to find the necessary header files and which libraries to link with the object file.</p>	C409.4	BTL1
4.	<p><b>Show the syntax of MPI_init and MPI_finalize</b> The syntax of MPI_Init():   <pre>int MPI_Init( int*      argc p /* in/out */, char***   argv p /* in/out */);</pre> The syntax of MPI_Finalize():  <pre>int MPI_Finalize(void);</pre></p>	C409.4	BTL1
5.	<p><b>What is Communicator in MPI?</b> In MPI a <b>communicator</b> is a collection of processes that can send messages to each other. One of the purposes of MPI Init is to define a communicator that consists of all of the processes started by the user when she started the program. This communicator is called MPI_COMM_WORLD.</p>	C409.4	BTL1

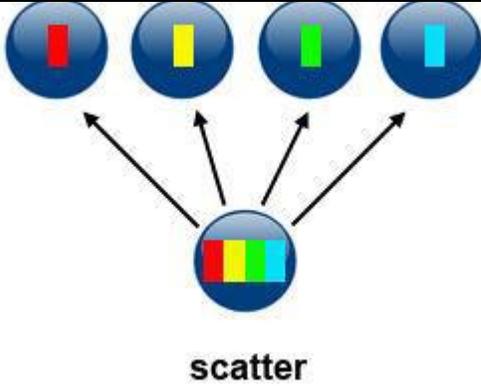
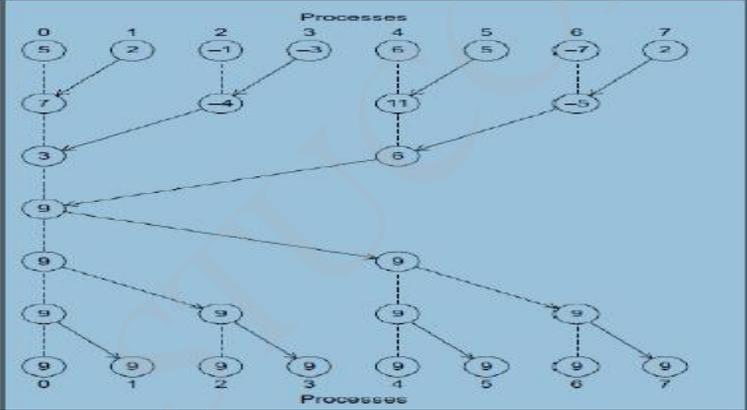
6.	<p><b>What is a SPMD program?</b> A single program is written so that different processes carry out different actions, and this is achieved by simply having the processes branch on the basis of their process rank. This approach to parallel programming is called <b>single program, multiple data, or SPMD</b></p>	C409.4	BTL1
7.	<p><b>Give the syntax of MPI_Get_count</b> The syntax of MPI_Get_count is  <pre>int MPI_Get_count(     MPI_Status* status_p    /* in */,     MPI_Datatype type       /* in */,     int* count p           /*out */);</pre></p>	C409.4	BTL1
8.	<p><b>What is Non-overtaking?</b> If process q sends two messages to process r, then the first message sent by q must be available to r before the second message. There is no restriction on the arrival of messages sent from different processes. ie., if q and t both send messages to r, then even if q sends its message before t sends its message, there is no requirement that q's message become available to r before t's message. This is essentially because MPI can't impose performance on a network.</p>	C409.4	BTL1
9	<p><b>Evaluate the performance evaluation methods in distributed memory programming. .(Nov/Dec 2017,</b>  1.Timing performance  2.performance Results  3.Speedup performance  4.efficiency performance  5.Scalability performance</p>	C409.4	BTL5
10	<p><b>What is wildcard argument?</b> The wildcard arguments:  <ul style="list-style-type: none"> <li>○ Only a receiver can use a wildcard argument. Senders must specify a process rank and a nonnegative tag. Thus, MPI uses a “push” communication mechanism rather than a “pull” mechanism.</li> </ul> There is no wildcard for communicator arguments; both senders and receivers must always specify communicators</p>	C409.4	BTL1
11	<p><b>Show the area of the trapezoid.</b>  Area of one trapezoid = <math>h/2[f(x_i) + (f(x_{i+1}))]</math></p>	C409.4	BTL1
12	<p><b>Define collective communications in MPI.</b>  In MPI parlance, communication functions that involve all the processes in a communicator are called collective communications.</p>	C409.4	BTL1
13	<p><b>What is Local Variables? Give Examples.</b>  Local variables are variables whose contents are significant only on the process that's using them. Some examples are: <b>local a, local b</b> and</p>	C409.4	BTL1

	<b>local n.</b>		
14	<p><b>What is Global Variables? Give Examples.</b> Variables whose contents are significant to all the processes are sometimes called global variables. Some examples are: <b>a, b</b> and <b>n</b>.</p>	C409.4	BTL1
15	<p><b>What is point to point communications in MPI?</b> <b>MPI_Send</b> and <b>MPI_Recv</b> are often called point-to-point communications.</p>	C409.4	BTL1
16	<p><b>List any two points stating that how collective communication differ from point to point communication.</b> ▪ All the processes in the communicator must call the same collective function. Example, a program that attempts to match a call to <b>MPI_Reduce</b> on one process with a call to <b>MPI_Recv</b> on another process is erroneous, and, in all likelihood, the program will hang or crash.</p>	C409.4	BTL4
17	<p><b>What is a butterfly-structured global sum?</b> The processes exchange partial results instead of using one-way communications. Such a communication pattern is sometimes called a butterfly-structured global sum.</p>	C409.4	BTL1
18	<p><b>What is broadcast?</b> A collective communication in which data belonging to a single process is sent to all of the processes in the communicator is called a <b>broadcast</b></p>	C409.4	BTL1
19	<p><b>Show a broadcast function.</b> A broadcast function:  <pre>int MPI_Bcast ( void*      data p      /* in/out */ , int        count       /* in      */ , MPI_Datatype datatype  /* in      */ , int        source proc /* in      */ , MPI_Comm   comm        /* in      */ );</pre> </p>	C409.4	BTL1
20	<p><b>Define block partition of the vector.</b> The work consists of adding the individual components of the vectors, so we might specify that the tasks are just the additions of corresponding components. Then there is no communication between the tasks, and the problem of parallelizing vector addition boils down to aggregating the tasks and assigning them to the cores. If the number of components is <math>n</math> and we have <math>comm\_sz</math> cores or processes, let's assume that <math>n</math> evenly divides <math>comm\_sz</math> and define <math>local\ n = n / comm\_sz</math>. Then we can simply assign blocks of <math>local\ n</math> consecutive components to each process. This is often called a <b>block partition</b> of the vector.</p>	C409.4	BTL1
21	<p><b>Show MPI Scatter function.</b> The communication MPI provides a function:  <pre>int MPI_Scatter( void*      send_buf_p /* in */ , int        send_count /* in */ , MPI_Datatype send_type /* in */ ,</pre> </p>	C409.4	BTL1

	<pre> void*      recv_buf_p      /* out */, int        recv_count     /* in   */, MPI_Datatype recv_type    /* in   */, int        src_proc       /* in   */, MPI_Comm   comm          /* in   */);         </pre>		
22	<p><b>Give MPI Gather function.</b>  The communication in this function can be carried out by MPI Gather,</p> <pre> int MPI_Gather( void*      send_buf_p      /* in   */, int        send_count     /* in   */, MPI_Datatype send_type    /* in   */, void*      recv_buf_p     /* out  */, int        recv_count     /* in   */, MPI_Datatype recv_type    /* in   */, int        dest_proc      /* in   */, MPI_Comm   comm          /* in   */);         </pre>	C409.4	BTL1
23	<p><b>What is Derived data type in MPI?</b>  In MPI, a derived data type can be used to represent any collection of data items in memory by storing both the types of the items and their relative locations in memory.</p>	C409.4	BTL1
24	<p><b>Show the use MPI_Type_create_struct.</b>  We can use MPI Type create struct to build a derived datatype that consists of individual elements that have different basic types:</p> <pre> int MPI_Type_create_struct( int        count          /* in        */, int        array_of_blocklengths[] /* in        */, MPI_Aint   array_of_displacements[] /* in        */, MPI_Datatype array_of_types[] /* in        */, MPI_Datatype* new_type_p /* out       */);         </pre>	C409.4	BTL1
25	<p><b>Show the ratio of the serial and parallel runtime.</b>  The most widely used measure of the relation between the serial and the parallel run-times is the speedup. It's just the ratio of the serial run-time to the parallel run-time:</p> $S(n, p) = \frac{T_{serial}(n)}{T_{parallel}(n, p)}$	C409.4	BTL1
26	<p><b>Show "per process" speedup?</b>  The "Per Process speedup is</p>	C409.4	BTL1

	$E(n, p) = \frac{S(n, p)}{p} = \frac{T_{serial}(n)}{p * T_{parallel}(n, p)}$		
27	<p><b>what is a wrapper script? . (apr/may2017)</b>                      A wrapper script is a script whose main purpose is to run the wrapper the running of the compiler by telling it where to find the necessary which libraries to link with the object file .</p>	C409.4	BTL1
28	<p><b>How to design a parallel program?</b>                      A parallel program can be designed using four basic steps:                      1. Partition the problem solution into tasks.                      2. Identify the communication channels between the tasks.                      3. Aggregate the tasks into composite tasks.                      4. Map the composite tasks to cores.</p>	C409.4	BTL1
29	<p><b>What is linear speedup?</b>                      The ideal value for <math>S(n,p)</math> is <math>p</math>. If <math>S(n,p) = p</math>, then our parallel program with <math>comm\_sz = p</math> processes is running <math>p</math> times faster than the serial program. This speedup, sometimes called <b>linear speedup</b>.</p>	C409.4	BTL1
30	<p><b>What is block-cyclic partition?</b>                      Instead of using a cyclic distribution of individual components, use a cyclic distribution of blocks of components, so a block-cyclic distribution isn't fully specified until we decide how large the blocks are</p>	C409.4	BTL1
31	<p><b>What are the possibilities for choosing a destination when sending requests for work with MPI</b>                      MPI is designed to allow users to create programs that can run efficiently on most parallel architectures. The design process included vendors (such as IBM, Intel, TMC, Cray, Convex, etc.), parallel library authors (involved in the development of PVM, Linda, etc.), and applications specialists</p>	C409.4	BTL1
32	<p><b>List the restrictions work sharing constructs</b>                      The sequence of <i>work-sharing constructs</i> and barrier directives encountered must be the same for every thread in a team</p>	C409.4	BTL1
33	<p><b>Write the evaluation methods is distributed memory programming</b></p>	C409.4	BTL1
34	<p><b>Give the commands for MPI</b></p>	C409.4	BTL1
35	<p><b>Define and broadcast and butterfly MPI</b></p>	C409.4	BTL1
36	<p><b>What is MPI W_Time</b>                      MPI provides a function, MPI_Wtime, that returns the number of seconds that have elapsed since some time in the past:  <b>Double MPI_Wtime(void);</b></p>	C409.4	BTL1
37	<p><b>Define MPI_Barrier</b>                      The MPI collective communication function MPI_Barrier insures that no process will return from calling it until every process in the communicator has started calling it.</p>	C409.4	BTL1

38	<p><b>Define Speed-Up and Efficiency</b> The most widely used measure of the relation between the serial and the parallel run-times is the <b>speedup</b>. It's just the ratio of the serial run-time to the parallel run-time:</p> $S(n,p) = T_{\text{Serial}(n)} / T_{\text{Serial}(n,p)}$	C409.4	BTL1
39	<p><b>How the <math>T_{\text{overhead}}</math> is represented</b> The parallel run-time is denoted by <math>T_{\text{parallel}}</math>. Since it depends on both the input size, <math>n</math>, and the number of processes, <math>\text{comm\_sz} = p</math>, we'll frequently denote it as <math>T_{\text{parallel}}(n,p)</math>. The parallel program will divide the work of the serial program among the processes, and add in some overhead time, which we denoted <math>T_{\text{overhead}}</math>:</p> $T_{\text{parallel}}(n,p) = T_{\text{serial}}(n/p) + T_{\text{overhead}}$	C409.4	BTL1
40	<p><b>Define Speed up</b> The most widely used measure of the relation between the serial and the parallel run-times is the <b>speedup</b>. It's just the ratio of the serial run-time to the parallel run-time:</p> $S(n,p) = T_{\text{serial}}(n) / T_{\text{parallel}}(n,p)$	C409.4	BTL1
41	<p><b>Define Efficiency</b> This speedup, sometimes called <b>linear speedup</b>. Another widely used measure of parallel performance is parallel <b>efficiency</b>.</p> $E(n,p) = S(n,p) / p = T_{\text{serial}}(n) / (P * T_{\text{parallel}}(n,p))$	C409.4	BTL1
42	<p><b>What is MPI derived data types</b> In MPI, a <b>derived datatype</b> can be used to represent any collection of data items in memory by storing both the types of the items and their relative locations in memory.</p>	C409.4	BTL1
43	<p><b>Define Gather</b> Gathers distinct messages from each task in the group to a single destination task. This routine is the reverse operation of MPI_Scatter. The data stored in the memory referred to by <code>send_buf_p</code> on process 0 is stored in the first block in <code>recv_buf_p</code>, the data stored in the memory referred to by <code>send_buf_p</code> on process 1 is stored in the second block referred to by <code>recv_buf_p</code>, and so on.</p>	C409.4	BTL1
44	<p><b>Define Scatter</b> Distributes distinct messages from a single source task to each task in the group.</p>	C409.4	BTL1

	 <p style="text-align: center;"><b>scatter</b></p>		
45	<p><b>List out the types of collective operations</b></p> <p><b>Synchronization</b> - processes wait until all members of the group have reached the synchronization point.</p> <ul style="list-style-type: none"> <li>• <b>Data Movement</b> - broadcast, scatter/gather, all to all</li> </ul> <p><b>Collective Computation (reductions)</b> - one member of the group collects data from the other members and performs an operation (min, max, add, multiply, etc.) on that data.</p>	C409.4	BTL1
46	<p><b>Define MPI_Allreduce</b></p> <p>Consider a situation in which <i>all</i> of the processes need the result of a global sum in order to complete some larger computation. For example, if we use a tree to compute a global sum, we might —reverse the branches to distribute the global sum.</p> 	C409.4	BTL1
47	<p><b>List out the two possibilities when the message are assembled</b></p> <ul style="list-style-type: none"> <li>• the sending process can <b>buffer</b> the message or</li> <li>• it can <b>block</b></li> </ul>	C409.4	BTL1
48	<p><b>What are the types of MPI type</b></p> <p>The MPI type MPI_Status is a struct with at least the three members <b>MPI_SOURCE</b>, <b>MPI_TAG</b>, and <b>MPI_ERROR</b>.</p> <p>Suppose our program contains the definition MPI_Status status; Then, after a call to MPI Recv in which &amp;status is passed as the last argument, we can determine the sender and tag by examining the two members</p> <p style="text-align: center;"><b>status.MPI_SOURCE</b> <b>status.MPI_TAG</b></p>	C409.4	BTL1
49	<p><b>What is status_p argument</b></p> <ul style="list-style-type: none"> <li>• The Amount of Data In The Message</li> </ul>	C409.4	BTL1

	<ul style="list-style-type: none"> <li>The Sender of The Message, Or</li> <li>The Tag of The Message.</li> </ul>		
50	<p><b>Mention the syntax for MPI_send</b></p> <pre>int MPI_Send(     void*      msg_buf_p    /* in */,     int        msg_size     /* in */,     MPI_Datatype msg_type   /* in */,     int        dest         /* in */,     int        tag          /* in */,     MPI_Comm   communicator /* in */);</pre>	C409.4	BTL1
51	<p><b>Write a note on Distributed memory machines (Nov/Dec 2018)</b>          Programming on a distributed memory machine is a matter of organizing a program as a set of independent tasks that communicate with each other via messages. In addition, programmers must be aware of where data is stored, which introduces the concept of locality in parallel algorithm design. An algorithm that allows data to be partitioned into discrete units and then runs with minimal communication between units will be more efficient than an algorithm that requires random access to global structures.</p>	C409.4	BTL1
52	<p><b>How to compile an MPI Program (Nov/Dec 2018)</b>          Many systems use a command called mpicc for compilation mpicc is a script that's a <b>wrapper</b> for the C compiler.</p>	C409.4	BTL1

### PART B

Q. No.	Questions	CO	Bloom's Level
1.	<p><b>What is MPI? Write a program "hello, world" that makes some use of MPI. how to compile and execute MPI programs. (Nov/Dec 2017).</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:86-90</p>	C409.4	BTL1
2.	<p><b>Explain about Trapezoidal rule in MPI in detail</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:94-96</p>	C409.4	BTL1
3.	<p><b>Give in detail about Collective Communication.</b>            1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:101-115</p>	C409.4	BTL1

4.	<p><b>Explain the following MPI functions:</b></p> <ul style="list-style-type: none"> <li>▪ <b>MPI_Reduce</b></li> <li>▪ <b>MPI_Allreduce</b></li> <li>▪ <b>MPI_Scatter</b></li> <li>▪ <b>MPI_Gather</b></li> <li>▪ <b>MPI_Allgather</b></li> </ul> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:88-91</p> <ul style="list-style-type: none"> <li>▪</li> </ul>	C409.4	BTL1
5.	<p><b>Compare and contrast Collective communication Vs Point to point communication.(16) (Nov/Dec 2017).</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:105-106</p>	C409.4	BTL2
6.	<p><b>Discuss about MPI Derived data types with example programs. (10)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:116-118</p>	C409.4	BTL6
7.	<p><b>i.Explain about Performance Evaluation of MPI Programs in detail.(Apr/May 2017)</b></p> <p><b>ii.What are the performance issues in multicore processor?</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:119-126</p>	C409.4	BTL5
8.	<p><b>i.Explain tree structured communication</b></p> <p><b>ii.What are the differences between point to point and collective communication? (Apr/May 2017)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011 Page No:102-103</p>	C409.4	BTL5
9	<p><b>(i)Explain Loop Handling in detail</b></p> <p><b>(ii)Describe about MPI programs execution with example (Apr/May 2018)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>	C409.4	BTL5
10	<p><b>Explain the virtual memory in detail (Apr/May 2018)</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>	C409.4	BTL5
11	<p><b>(i)Describe the Attribute Caching</b></p> <p><b>(ii)Discuss about the communicators</b></p> <p>1. Peter S. Pacheco, “An Introduction to Parallel Programming”, Morgan-Kauffman/Elsevier, 2011</p>	C409.4	BTL2

12	<b>(i) Describe the Distributed array datatype constructor</b> <b>(ii) Explain the Cartesian constructor</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.4	BTL5
13	<b>(i) Generalize the group of process</b> <b>(ii) Explain the virtual topology</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.4	BTL6
14	<b>(i) Describe the MPI program execution</b> <b>(ii) Describe about MPI Init and Finalize</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.4	BTL1
15	<b>(i) Describe about the Datatype constructor</b> <b>(ii) Discuss about the subarray datatype constructor</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.4	BTL5

## UNIT V

### PARALLEL PROGRAM DEVELOPMENT

Case studies - n-Body solvers – Tree Search – OpenMP and MPI implementations and comparison.

#### PART A

Q. No.	Questions	CO	Bloom's Level
1.	<b>What is an n-body problem?</b> In an n-body problem, it needs to find the positions and velocities of a collection of interacting particles over a period of time. An n-body solver is a program that finds the solution to an n-body problem by simulating the behavior of the particles. The input to the problem is the mass, position, and velocity of each particle at the start of the simulation, and the output is the position and velocity of each particle at a sequence of user-specified times, or simply the position and velocity of each particle at the end of a user-specified time period.	C409.5	BTL1

2.	<p><b>Show the pseudocode of a serial n-body solver.</b></p> <pre> 1      Get input data; 2      for each timestep { 3          if (timestep output) Print positions and velocities of particles; 4          for each particle q 5              Compute total force on q; 6          for each particle q 7              Compute position and velocity of q; 8      } 9      Print positions and velocities of particles;</pre>	C409.5	BTL1
3.	<p><b>List the two algorithms in the n-body solver.</b></p> <p>The n-body solver with the original force calculation, the <b>basic</b> algorithm, and the solver with the number of calculations reduced the <b>reduced</b> algorithm</p>	C409.5	BTL1
4.	<p><b>List the two algorithms in the n-body solver.</b></p> <p>The n-body solver with the original force calculation, the <b>basic</b> algorithm, and the solver with the number of calculations reduced the <b>reduced</b> algorithm</p>	C409.5	BTL1
5.	<p><b>Differentiate between two algorithms in n-body solver</b></p> <p>The <i>n</i>-body solver with the original force calculation, the <i>basic</i> algorithm, and the solver with the number of calculations reduced, the <i>reduced</i> algorithm.</p>		
6.	<p><b>Define Graph.</b></p> <p>A <b>graph</b> is a collection of vertices and edges or line segments joining pairs of vertices.</p>	C409.5	BTL1
7.	<p><b>List the use of Recursive depth-first search algorithm.</b></p> <p>The algorithm makes use of several global variables:</p> <ul style="list-style-type: none"> <li>• <b>n</b>: the total number of cities in the problem</li> <li>• <b>digraph</b>: a data structure representing the input digraph</li> <li>• <b>hometown</b>: a data structure representing vertex or city 0, the salesperson's hometown</li> <li>• <b>best tour</b>: a data structure representing the best tour so far</li> </ul>	C409.5	BTL1
8	<p><b>What are the disadvantages of recursive depth first stage</b></p> <p>It also has the disadvantage that at any given instant of time only This could be a problem when we try to parallelize tree search by threads or processes</p>		
8.	<p><b>List the similarities parallelizing the solvers using pthreads.</b></p> <p>The more important similarities are:</p> <ul style="list-style-type: none"> <li>➤ By default local variables in Pthreads are private, so all shared variables are global in the Pthreads version.</li> <li>➤ The principal data structures in the Pthreads version are identical to those in the OpenMP version: vectors are two-dimensional arrays of doubles, and the mass, position, and velocity of a single particle are stored in a struct. The forces are stored in an array of vectors.</li> <li>➤ Startup for Pthreads is basically the same as the startup for OpenMP: the main thread gets the command-line arguments, and allocates and initializes</li> </ul>	C409.5	BTL1

	<p>the principal data structures.</p> <ul style="list-style-type: none"> <li>➤ The main difference between the Pthreads and the OpenMP implementations is in the details of parallelizing the inner loops. Since Pthreads has nothing analogous to a parallel for directive, we must explicitly determine which values of the loop variables correspond to each thread’s calculations. To facilitate this, a function Loop schedule, which determines             <ul style="list-style-type: none"> <li>• the initial value of the loop variable,</li> <li>• the final value of the loop variable, and</li> <li>• the increment for the loop variable.</li> </ul> </li> <li>➤ The input to the function is             <ul style="list-style-type: none"> <li>• the calling thread’s rank,</li> <li>• the number of threads,</li> <li>• the total number of iterations, and</li> </ul>             an argument indicating whether the partitioning should be block or cyclic           </li> </ul>		
9.	<p><b>Define communication structure</b></p> <p>A communication structure is called a <b>ring pass</b>. In a ring pass, imagine the processes as being interconnected in a ring. Process 0 communicates directly with processes 1 and comm._sz-1, process 1 communicates with processes 0 and 2, and so on. The communication in a ring pass takes place in phases, and during each phase each process sends data to its “lower-ranked” neighbor, and receives data from its “higher-ranked” neighbor. Thus, 0 will send to comm._sz -1 and receive from 1. 1 will send to 0 and receive from 2, and so on. In general, process q will send to process (q-1+comm._sz)%comm_sz and receive from process (q+1)%comm_sz.</p>	C409.5	BTL1
9	<p><b>List the properties of function First_index.</b></p> <p>The function First index should determine a global index glb part2 with the following properties:</p> <ol style="list-style-type: none"> <li>1. The particle glb_part2 is assigned to the process with rank owner.</li> <li>2. <math>glb\_part1 &lt; glb\_part2 &lt; glb\_part1 + comm\_sz</math>.</li> </ol>	C409.5	BTL1
10	<p><b>What is a word about I/O.</b></p> <p>The basic I/O was designed for use by single-process, single-threaded programs, and when multiple processes or multiple threads attempt to access the I/O buffers; the system makes no attempt to schedule their access.</p>	C409.5	BTL1
11	<p><b>What is race condition? .(Nov/Dec 2017)</b></p> <p>A race condition is an undesirable situation that occurs when a device or system attempts to perform two or more operations at the same time, but because of the nature of the device or system, the operations must be done in the proper</p>	C409.5	BTL1

	sequence to be done correctly.		
12	<p><b>Show the Pseudocode for the MPI implementation of the reduced n-body solver.</b></p> <pre> 1   source = (my_rank + 1) % comm._sz; 2   dest = (my_rank - 1 + comm._sz) % comm._sz; 3   Copy loc_pos into tmp_pos; 4   loc_forces = tmp_forces = 0; 5 6   Compute forces due to interactions among local particles; 7   for (phase = 1; phase &lt; comm._sz; phase++) { 8   Send current tmp_pos and tmp_forces to dest; 9   Receive new tmp_pos and tmp_forces from source; 10  /* Owner of the positions and forces we're receiving */ 11  owner = (my_rank + phase) % comm._sz; 12  Compute forces due to interactions among my particles 13  and owner's particles; 14  } 15  Send current tmp_pos and tmp_forces to dest; 16  Receive new tmp_pos and tmp_forces from source;</pre>	C409.5	BTL1
13	<p><b>What is NP-complete problem? .(Apr/May 2017)</b></p> <p>In computational complexity theory, a decision problem is NP-complete when it is both in NP and NP-hard. The set of NP-complete problems is often denoted by NP-C or NPC. The abbreviation NP refers to "nondeterministic polynomial time".</p>	C409.5	BTL1
14	<p><b>What is Depth-first search?</b></p> <p>Depth-first search (DFS) is an algorithm .for traversing or searching tree or graph data structures. One starts at the root (selecting some arbitrary node as the root in the case of a graph) and explores as far as possible along each branch before backtracking.</p>	C409.5	BTL1
15	<p><b>What is Recursive depth-first search?</b></p> <p>Depth-first search (DFS) is an algorithm that traverses a graph in search of one or more goal nodes. The defining characteristic of this search is that, whenever DFS visits a maze cell c, it recursively searches the sub-maze whose origin is c. This recursive behaviour can be simulated by an iterative algorithm using a stack. A cell can have three states:</p> <ul style="list-style-type: none"> <li>• <b>Unvisited.</b> The cell has not yet been visited by DFS.</li> <li>• <b>Visit In Progress.</b> The cell has been discovered, but not yet finished. Ie, the recursive search which begins at this node has not yet terminated.</li> <li>• <b>Visited.</b> The cell has been discovered, and the submazes which start at this node have been completely visited also.</li> </ul>	C409.5	BTL1
16	<p><b>What problem occurs when test lock condition and update lock condition combined?</b></p> <p>The combination of “test lock condition” and “update lock condition” can</p>	C409.5	BTL1

	<p>cause a problem: the lock condition (e.g. the cost of the best tour) can change between the time of the first test and the time that the lock is acquired. Thus, the threads also need to check the lock condition after they acquire the lock.</p>		
17	<p><b>Show the pseudocode for updating the best tour.</b>  The pseudocode for updating the best tour should look something like this:</p> <pre> if (new tour cost &lt; best tour cost) {     Acquire lock protecting best tour;     if (new tour cost &lt; best tour cost)         Update best tour;     Relinquish lock; } </pre>	C409.5	BTL1
18	<p><b>Show the syntax for mutex checking.</b>  Pthreads has a <i>nonblocking</i> version of <i>pthread_mutex_lock</i> called <i>pthread_mutex_trylock</i>. This function checks to see if the mutex is available. If it is, it acquires the mutex and returns the value 0. If the mutex isn't available, instead of waiting for it to become available, it will return a nonzero value.</p>	C409.5	BTL1
19	<p><b>Define MPI.</b>  The <i>message passing interface (MPI)</i> is a standardized means of exchanging messages between multiple computers running a parallel program across distributed memory</p>	C409.5	BTL1
20	<p><b>Show a pseudocode for a recursive solution to TSP using depth first search.(Apr/May2017)</b></p> <pre> 1 void Depth_first_search(tour_t tour) { 2     city_t city; 3 4     if (City_count(tour) == n) { 5         if (Best_tour(tour)) 6             Update_best_tour(tour); 7     } else { 8         for each neighboring city 9             if (Feasible(tour, city)) { 10                Add_city(tour, city); 11                Depth_first_search(tour); 12                Remove_last_city(tour, city); 13            } 14        } 15    } /* Depth_first_search */ </pre>	C409.5	BTL1
21	<p><b>What are the features of distributed memory?(Nov/Dec 2017)</b></p> <p>distributed memory refers to a multiprocessor computer system in which each processor has its own private memory. Computational tasks can only operate on local data, and if remote data is required, the computational task must communicate</p>	C409.5	BTL1

	with one or more remote processors.		
22	<p><b>Show the syntax of MPI_Pack and MPI_Unpack.</b></p> <pre> int MPI_Pack( void*          data to be packed          /* in */ int          to be packed count / _ in _/, MPI_Datatype  datatype /_ in _/, void_ contig buf /_ out _/, int contig buf size /_ in _/, int_ position p /_ in/out _/, MPI Comm comm /_ in _/);  int MPI_Unpack( void_ contig buf /_ in _/, int contig buf size /_ in _/, int_ position p /_ in/out _/, void_ unpacked data /_ out _/, int unpack count /_ in _/, MPI Datatype datatype /_ in _/, MPI Comm comm /_ in _/);                     </pre>	C409.5	BTL1
23	<p><b>List the use of MPI_IN_PLACE argument.</b></p> <p>The use of argument <i>MPI_IN_PLACE</i> is that the input and output buffers are the same. This can save on memory and the implementation may be able to avoid copying from the input buffer to the output buffer</p>	C409.5	BTL1
24	<p><b>What the use of functions MPI Scatter and MPI Gather.</b></p> <p>The functions <i>MPI_Scatter</i> and <i>MPI_Gather</i> can be use to split an array of data among processes and collect distributed data into a single array, respectively</p>	C409.5	BTL1
25	<p><b>What are the use of functions MPI_Scatterv and MPI_Gatherv.</b></p> <p>When the amount of data going to or coming from each process is the same for each process. If we need to assign different amounts of data to each process, or to collect different amounts of data from each process we can use <i>MPI_Scatterv</i> and <i>MPI_Gatherv</i>, respectively.</p>	C409.5	BTL1
26	<p><b>Show the syntax of MPI_Scatterv.</b></p> <pre> int MPI_Scatterv( void*          sendbuf          /* in */ int*          sendcounts       /* in */ int*          displacements    /* in */ MPI_Datatype  sendtype        /* in */ Void*          recvbuf        /* out */ int          recvcnt          /* in */ MPI_Datatype  recvtype        /* in */ int          root              /* in */ MPI_Comm     comm             /* in */);                     </pre>	C409.5	BTL1
27	<p><b>Show the syntax of MPI_Gatherv.</b></p>	C409.5	BTL1

	<pre> int MPI_Gatherv( void*          sendbuf          /* in */, int           sendcount        /* in */, MPI_Datatype  sendtype         /* in */, void*          recvbuf         /* out */, int*          recvcnts         /* in */, int*          displacements    /* in */, MPI_Datatype  recvtype         /* in */, int           root             /* in */, MPI_Comm      comm            /* in */);         </pre>		
28	<p><b>What are the three modes provided by MPI?</b>  MPI provides three other modes for sending: <b>synchronous, standard, and ready.</b></p>	C409.5	BTL1
29	<p><b>List the use of functions MPI_Pack and MPI_Unpack.</b>  MPI provides the function <i>MPI_Pack</i> for storing a data structure in a single, contiguous buffer before sending. Similarly, the function <i>MPI_Unpack</i> can be used to take data that's been received into a single contiguous buffer and unpack it into a local data structure</p>	C409.5	BTL1
30	<p><b>List the use of ready mode.</b>  Ready sends (<i>MPI_Rsend</i>) are erroneous unless the matching receive has already been started when <i>MPI_Rsend</i> is called. The ordinary send <i>MPI_Send</i> is called the standard mode send.</p>	C409.5	BTL1
31	<p><b>List the use of Synchronous mode.</b>  Synchronous sends won't buffer the data; a call to the synchronous send function MPI_Ssend won't return until the receiver has begun receiving the data.</p>	C409.5	BTL1
32	<p><b>How to compute n-body forces</b></p> <pre> for each particle q forces[q] = 0; for each particle q { for each particle k &gt; q { x_diff = pos[q][X] - pos[k][X]; y_diff = pos[q][Y] - pos[k][Y]; dist = sqrt(x_diff_x_diff + y_diff_y_diff); dist_cubed = dist_dist_dist; force_qk[X] = G_masses[q]_masses[k]/dist_cubed _ x_diff; force_qk[Y] = G_masses[q]_masses[k]/dist_cubed _ y_diff forces[q][X] += force_qk[X]; forces[q][Y] += force_qk[Y]; forces[k][X] -= force_qk[X]; forces[k][Y] -= force_qk[Y]; } }         </pre>	C409.5	BTL1

33	<p><b>List the data structures used for serial implementation</b></p> <p>The data structures are the tour, the digraph, and, in the iterative implementations, the stack and the stack are essentially list structures. For tour instead of array structure with three arrays: the array storing the cities, the number of cities, and the cost of the partial tour. When the digraph are represented using List.</p>	C409.5	BTL1
34	<p><b>Difference between Parallelizing the two <math>n</math>-body solvers using pthread and OpenMP.</b></p> <p>The main difference between the Pthreads and the OpenMP implementations is in the parallelizing the inner loops. Since Pthreads has nothing analogous to a parallel for loop, it must explicitly determine which values of the loop variables correspond to each thread's calculations. To facilitate this a function Loop schedule which contains</p> <ul style="list-style-type: none"> <li>. the initial value of the loop variable,</li> <li>. the final value of the loop variable, and</li> <li>. the increment for the loop variable.</li> </ul>	C409.5	BTL1
35	<p><b>How the performance of the reduced solver is much superior to the performance of the basic solver?</b></p> <p>The efficiency of the basic solver on 16 nodes is about 0.95, while the efficiency of the reduced solver on 16 nodes is only about 0.70. A point to stress here is that the reduced MPI solver makes much more efficient use of memory than the basic MPI solver the basic solver must provide storage for <math>n</math> positions on each process, while the reduced solver only needs extra storage for <math>n = \text{comm\_sz}</math> positions. <math>n = \text{comm\_sz}</math> forces</p>	C409.5	BTL1
36	<p><b>How can we use OpenMP to map tasks/particles to cores in the basic version of the <math>n</math>-body solver?</b></p> <pre> for each timestep { if (timestep output) Print positions and velocities of particles; for each particle q Compute total force on q; for each particle q Compute position and velocity of q; } </pre> <p>The two inner loops are both iterating over particles. So, in principle, parallelizing the two inner for loops will map tasks/particles to cores, and we might try something like this:</p> <pre> for each timestep { if (timestep output) Print positions and velocities of particles; </pre>	C409.5	BTL1

	<pre> # pragma omp parallel for for each particle q   Compute total force on q; # pragma omp parallel for for each particle q   Compute position and velocity of q; } </pre>		
37	<p><b>Write the pseudocode for the MPI version of the basic <math>n</math>-body solver?</b></p> <pre> Get input data; 2 for each timestep { 3   if (timestep output) 4     Print positions and velocities of particles; 5   for each local particle loc q 6     Compute total force on loc q; 7   for each local particle loc q 8     Compute position and velocity of loc q; 9   Allgather local positions into global pos    array; 10 } 11 Print positions and velocities of particles; </pre>	C409.5	BTL1
38	<p><b>Write the pseudocode for the MPI implementation of the reduced <math>n</math>-body solver?</b></p> <pre> 1 source = (my_rank + 1) % comm_sz; 2 dest = (my_rank - 1 + comm_sz) % comm_sz; 3 Copy_loc_pos_into_tmp pos; 4 loc_forces = tmp_forces = 0; 5 6 Compute forces due to interactions among local particles; 7 for (phase = 1; phase &lt; comm_sz; phase++) { 8   Send current tmp_pos and tmp_forces to dest; 9   Receive new tmp_pos and tmp_forces from source; 10  /_ Owner of the positions and forces we're receiving _/ 11  owner = (my_rank + phase) % comm_sz; 12  Compute forces due to interactions among my particles    and owner's particles; 13 } 14 } </pre>	C409.5	BTL1

	15 Send current tmp_pos and tmp_forces to dest; 16 Receive new tmp_pos and tmp_forces from source;		
39	<p><b>What are the two phases for computation of forces?</b></p> <p>The following choices with respect to the data structures: Each process stores the entire global array of particle masses. Each process only uses a single n-element array for the positions. Each process uses a pointer loc_pos that refers to the start of its block of pos.</p>	C409.5	BTL1
40	<p><b>Write the pseudo code for an implementation of a depth-first solution to TSP using recursion?</b></p> <pre> for (city = n-1; city &gt;= 1; city--) Push(stack, city); while (!Empty(stack)) { city = Pop(stack); if (city == NO CITY) // End of child list, back up Remove last city(curr tour); else { Add city(curr tour, city); if (City count(curr tour) == n) { if (Best tour(curr tour)) { Update best tour(curr tour); Remove last city(curr tour); } else { Push(stack, NO CITY); for (nbr = n-1; nbr &gt;= 1; nbr--) if (Feasible(curr tour, nbr)) Push(stack, nbr); } }/_ if Feasible _/ }/_ while !Empty _/ </pre>	C409.5	BTL1
41	<b>How the function Push_copy is used in TSP</b>	C409.5	BTL1

	<p>It is necessary to push onto the stack to create a copy of the tour before actually pushing it on to the stack using the function <code>Push_copy</code>. The extra memory is required to allocating storage for a new tour and copying the existing tour is time-consuming. Reduce the costs by saving freed tours in our own data structure, and when a freed tour is available we can use it in the <code>Push_copy</code> function instead of calling <code>malloc</code>.</p>		
42	<p><b>What are the algorithms for identifying which subtrees we assign to the p threads</b></p> <ul style="list-style-type: none"> <li>➤ depth-first search</li> <li>➤ breadth-first search</li> </ul>	C409.5	BTL1
43	<p><b>Define the term POSIX or PThreads</b></p> <p>Pthreads are libraries of type definitions, functions, and macros that can be used in C program a standard for Unix-like operating systems—for example, Linux and Mac OS X. It specifies facilities that should be available in such systems. In particular, it specifies an application programming interface (API) for <i>multithreaded</i> programming. Pthreads is not a programming language (such as Java). Rather, like MPI, Pthreads specifies a <i>library</i> that can be linked with C programs. Unlike MPI, Pthreads API is only available on POSIX systems—Linux, Mac OS X, Solaris, HPUX, and s</p>	C409.5	BTL1
44	<p><b>What are the reason for parameter <code>threads_in_cond_wait</code> used in Tree search?</b></p> <p>There are also two cases to consider:</p> <ul style="list-style-type: none"> <li>o <code>threads_in_cond_wait &lt; thread_count</code>, it tells us how many threads are wait</li> <li>o <code>threads_in_cond_wait == thread count</code>,all the treads are out of work, and its</li> </ul>	C409.5	BTL1
45	<p><b>What are the global variables for Recursive depth first search?</b></p> <p>n: the total number of cities in the problem .              digraph: a data structure representing the input digraph .              hometown: a data structure representing vertex or city 0, the salesperson’s home              tour: a data structure representing the best tour so far.</p>	C409.5	BTL1
46	<p><b>Mention the performance of MPI solvers</b></p> <p>The performance of the reduced solver is much superior to the performance of the basic solver, although the basic solver achieves higher efficiencies.              A point to stress here is that the reduced MPI solver makes much more efficient use of memory than the basic MPI solver; the basic solver must provide storage for all <i>n</i> positions on each process, while the reduced solver only needs extra storage for <i>n/commsz</i> positions and <i>n/commsz</i> forces.</p>	C409.5	BTL1
47	<p><b>Mention the principal data structures on pthread</b></p> <p>The vectors are two-dimensional arrays of <b>doubles</b>, and the mass, position, and</p>	C409.5	BTL1

	velocity of a single particle are stored in a struct. The forces are stored in an array of vectors.		
48	<b>Name any two OpenMp environment variables (Nov/Dec 2018)</b> <b>omp_set_num_threads(<i>num_threads</i>)</b> <b>omp_get_num_threads()</b> <b>omp_get_max_threads()</b> <b>omp_get_thread_num()</b>	C409.5	BTL1
49	<b>List any two scoping clauses in OpenMP (Nov/Dec 2018)</b> <ul style="list-style-type: none"> <li>• Shared Variables</li> <li>• Private Variables</li> </ul>	C409.5	BTL1
50	<b>What are the reason for parameter threads_in_cond_wait used in Tree search?</b>  there are also two cases to consider: <ul style="list-style-type: none"> <li>• threads_in_cond_wait &lt; thread_count, it tells us how many threads are waiting</li> <li>• threads_in_cond_wait == thread count,all the treads are out of work, and its time to quit.</li> </ul>	C409.5	BTL1

## PART-B

Q. No.	Questions	CO	Bloom's Level
1.	<b>1.Explain n-Body solvers in OpenMP.</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kauffman/Elsevier, 2011 Page No:271-297	C409.5	BTL5
2.	<b>Explain about Tree Search.</b>  1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kauffman/Elsevier, 2011 Page No:229-318	C409.5	BTL5
3.	<b>Explain OpenMP implementations in detail.</b>  1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kauffman/Elsevier, 2011 Page No:15-20	C409.5	BTL5

4.	<b>Explain MPI implementations in detail.</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No: 86-88	C409.5	BTL4
5.	<b>Compare OpenMP and MPI implementations.</b>  <b>Refer Notes</b>	C409.5	BTL4
6.	<b>Explain how to Parallelizing the tree search in detail.</b>  1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No: 306-308	C409.5	BTL5
7.	<b>i.How to parallelize the basic solver using MPI. (Apr/May2017)</b> <b>ii.Explain Non recursive depth first search. (Apr/May2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No: 316-317 & 303-304	C409.5	BTL5
8.	<b>Explain the implementation of tree search using MPI and dynamic partitioning. (Apr/May2017)</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No:229-318	C409.5	BTL5
9	<b>What does the n-body problem do/give the pseudocode for serial n-body solver and for computing n-body forces. (Nov/Dec 2017).</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No: 271-297	C409.5	BTL1
10	<b>How will you parallelize the reduced solver using Open Mp? How will you parallelize the reduced solver using Open MP? (Nov/Dec 2017).</b>  1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011 Page No: 271-297	C409.5	BTL1
11	<b>Generalize about the two Serial program</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.5	BTL6
12	<b>Describe about the Parallelizing the reduced solver using OpenMP</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kaufman/Elsevier, 2011	C409.5	BTL2

13	<b>Describe about the Recursive and non recursive DFS</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kauffman/Elsevier, 2011	C409.5	BTL2
14	<b>Examine the Data structures for the serial implementation</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kauffman/Elsevier, 2011	C409.5	BTL3
15	<b>Express detail about the static parallelizing of tree search using PThreads</b> 1. Peter S. Pacheco, "An Introduction to Parallel Programming", Morgan-Kauffman/Elsevier, 2011	C409.5	BTL1