Masters

Parallel and Distributed Computing (CS621)

Question: 1 (Marks: 1) **Attempted Questions: 0** **Total Questions: 50**

In _____ control strategy, the control function is distributed among different components in the system.

Choices:

- decentralized
- distributed
- centralized
- supervised

Masters

Parallel and Distributed Computing (CS621)

Question: 2 (Marks: 1) Attempted Questions: 1 Total Questions: 50

In _____ switching, a complete path has to be established prior to the start of communication between a source and a destination.

Choices:

- packet
- port
- circuit
- connect

Masters

Parallel and Distributed Computing (CS621)

Question: 3 (Marks: 1) Attempted Questions: 2 Total Questions: 50

Number of bits transmitted in a unit of time is known as_____.

Choices:

- latency
- diameter
- bandwidth
- bisection

Masters

Parallel and Distributed Computing (CS621)

Question: 4 (Marks: 1) **Attempted Questions: 3** **Total Questions: 50**

A _____ topology is a network interconnection scheme in which each node is connected to every other node with a direct link.

Choices:

- bus
- star
- mesh ✓
- switch

Masters

Parallel and Distributed Computing (CS621)

Question: 5 (Marks: 1) Attempted Questions: 4 Total Questions: 50

The possibility to use the same software in different environments is known as _____.

Choices:

- reliability
- portability
- robustness
- availability

Masters

Parallel and Distributed Computing (CS621)

Question: 10 (Marks: 1) **Attempted Questions: 6** **Total Questions: 50**

Decomposition of computation into a large number of tasks results in _____ decomposition.

Choices:

- fine-grained
- course-grained
- medium-grained
- best-grained

Masters

Parallel and Distributed Computing (CS621)

Question: 11 (Marks: 1) Attempted Questions: 7 Total Questions: 50

Discrete optimization problems can be solved by using _____ decomposition technique.

Choices:

- Recursive decomposition
- Data decomposition
- Exploratory decomposition
- Speculative decomposition

Masters

Parallel and Distributed Computing (CS621)

Question: 12 (Marks: 1) Attempted Questions: 8 Total Questions: 50

The problem of finding the minimum in a list is an example of decomposition of data based on _____

Choices:

- output data
- input data
- intermediate data
- both input and output data

Masters

Parallel and Distributed Computing (CS621)

Question: 13 (Marks: 1) Attempted Questions: 9 Total Questions: 50

Which of the following decomposition techniques is/are more desirable for quicksort problem?

Choices:

- Recursive decomposition
- Data decomposition
- Exploratory decomposition
- Both data and recursive decomposition

Masters

Parallel and Distributed Computing (CS621)

Question: 26 **(Marks: 1)** **Attempted Questions:** 10 **Total Questions:** 50

Sequential Consistency ensures that the result of any execution in a distributed system is equivalent to:

Choices:

- Operations being executed in a random order
- All processes executing concurrently without synchronization
- Read and write operations being executed in a sequential order
- Each process executing independently of others

Masters

Parallel and Distributed Computing (CS621)

Question: 27 (Marks: 1) Attempted Questions: 11 Total Questions: 50

What does causal consistency guarantee regarding writes in a distributed system?

Choices:

- All writes are seen by all processes in the same order.
- Concurrent writes are seen in the same order on different machines.
- All writes are executed concurrently to improve performance.
- Writes that are causally related are seen in the same order by all processes.

Parallel and Distributed Computing (CS621)

Question: 29 (Marks: 1) **Attempted Questions: 13** **Total Questions: 50**

GPUs have gained popularity due to their ability to:

Choices:

- Run only real-world applications efficiently.
- Execute single-threaded tasks faster than multi-core CPUs.

correct

- Replace multi-core CPUs in all computing tasks.
- Provide parallel processing capabilities for certain applications.

Parallel and Distributed Computing (CS621)

Question: 30 (Marks: 1) **Attempted Questions: 14** **Total Questions: 50**

CUDA is designed for computation-intensive tasks to be executed on which hardware?

Choices:

- CPU (Central Processing Unit)
- GPU (Graphics Processing Unit) ✓
- RAM (Random Access Memory)
- SSD (Solid State Drive)

Parallel and Distributed Computing (CS621)

Question: 31 (Marks: 1) Attempted Questions: 15 Total Questions: 50

GPUs accelerated libraries are commonly used in which type of applications?

Choices:

- Text editors
- Web browsers
- Data science and machine learning
- File compression utilities

Parallel and Distributed Computing (CS621)

Question: 32 (Marks: 1) Attempted Qu

In message passing, how do processors exchange data in a distributed memory system?

Choices:

- They share a global memory space.
- They use cache synchronization mechanisms.
- They communicate through a communication network. ✓
- They directly access each other's local memory.

Question: 34 (Marks: 1) Attempted Questions: 18 Total Questions: 50

What is the size of the Basic MPI library in terms of the number of functions?

Choices:

- Large (over 50 functions)
- Medium (around 25 functions)
- Small (around 10 functions)
- Basic (6 functions)

Parallel and Distributed Computing (CS621)

Question: 35 (Marks: 1) Attempted Questions: 19 Total Questions: 50

In MPI, which function is commonly used to send a message from one process to another?

Choices:

- MPI_Broadcast
- MPI_Reduce
- MPI_Send
- MPI_Receive

Parallel and Distributed Computing (CS621)

Question: 36 (Marks: 1) Attempted Questions: 20 Total Questions: 50

MapReduce is a software framework which supports _____computing on large data sets.

Choices:

- serial
- parallel
- distributed
- both parallel and distributed

Parallel and Distributed Computing (CS621)

Question: 37 (Marks: 1) Attempted Questions: 21 Total Questions: 50

Which of the following is NOT true regarding Google File System (GFS)?

Choices:

- High security
- Limited security
- Only viable in a specific environment
- Custom designed

Parallel and Distributed Computing (CS621)

Question: 38 (Marks: 1) Attempted Questions: 22 Total Questions: 50

_____ of Hadoop Distributed File System (HDFS) focuses on the individual instances of application data or the data content.

Choices:

- Structural metadata
- Descriptive metadata
- Contextual metadata
- Contextual and structural metadata

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 39 (Marks: 1) **Attempted Questions: 23** **Total Questions: 50**

What is the primary function of Hadoop's HDFS (Hadoop Distributed File System)?

Choices:

- Running complex machine learning algorithms.
- Storing and distributing data across a cluster of machines. ✓
- Managing user access control.
- Creating graphical visualizations

Parallel and Distributed Computing (CS621)

Question: 40 (Marks: 1) Attempted Questions: 24 Total Questions: 50

Which component of Hadoop is responsible for storing and managing data across a cluster?

Choices:

- Hadoop Mapper
- Hadoop Reducer
- Hadoop NameNode
- Hadoop TaskTracker

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 41 (Marks: 3) **Attempted Questions: 25** **Total Questions: 50**

Define the degree of the switch in interconnection networks.

the total number of ports in a switch is called the degree of switch.

.....

Answer:

Parallel and Distributed Computing (CS621)

Question: 42 (Marks: 3) Attempted Questions: 26 Total Questions: 50

Bill Dally, a prominent computer scientist has highlighted power efficiency and massive parallelism as major advantages of GPUs by considering two aspects. You are required to mention these aspects.

1- data access rate capability
2- data processing capability

.....

Answer:

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 43 (Marks: 3) **Attempted Questions: 28** **Total Questions:**

Write down the names of the data-centric consistency models.

1- 1- Continuous Consistency
2- Consistent Ordering of Operations
3- Causal Consistency
4- Grouping Operations

.....

Answer:

Parallel and Distributed Computing (CS621)

Question: 44 (Marks: 3) Attempted Questions: 28 Total Questions: 50

What will happen when multiple threads attempt to manipulate the same data item, if proper care is not taken to synchronize?

can results incoherent

Answer:

Parallel and Distributed Computing (CS621)

Question: 45 (Marks: 3) Attempted Questions: 29 Total Questions: 5(

The I/O tools are the collections of system software and libraries have grownup to address I/O issues.
You are required to list three I/O tools.

Answer:

Here is the names of three I/O tools.

At Parallel file systems
- MPI-IO
- High level libraries
- Relationships between these are not always clear.

Parallel and Distributed Computing (CS621)

Question: 46 (Marks: 3) Attempted Questions: 30 Total Questions:

Do you think, a parallel runtime of a program can be evaluated asymptotically as a function of input size only? Justify your answer in either case.

sequential algorithm can be evaluated asymptotically as a ..... same question aghy

.....

Answer:

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 47 (Marks: 5) **Attempted Questions: 31**

Write down any five interconnection network's properties in parallel and distributed computing.

.....

Answer:

Following are five interconnection network's properties in parallel and distributed computing:

1- Bandwidth
2- Latency
3- bisection
4- diameter
5- degree of node

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 49 (Marks: 5) **Attempted Questions: 34** **Total Questions: 5**

The process of dividing a computation into smaller parts, some or all of which may potentially be executed in parallel, is called decomposition. You are required to list 4 key decomposition techniques.

.....

Answer:

Following are the 4 key decomposition techniques.

1- Recursive Decomposition
2- Data Decomposition
3- Exploratory Decomposition
4- Speculative Decomposition

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 50 (Marks: 5) **Attempted Questions: 34** **Total Questions: 50**

Differentiate the following statement as Independent MPI IO and Collective MPI IO calls.

1. MPI_File_write()
2. MPI_File_write_all()
3. MPI_File_read().
4. MPI_File_read_all()

1ST AND 3RD independent MPI IO
2nd and 4th collective MPI IO

Answer:

Parallel and Distributed Computing (CS621)

Question: 7 (Marks: 1) Attempted Questions: 35 Total Questions: 50

Much of the effort associated with writing correct threaded programs is spent on _____ concurrent threads.

Choices:

- scheduling
- data accesses for
- synchronizing
- both scheduling and data access

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 8 **(Marks: 1)** **Attempted Questions: 36** **Total Questions: 50**

Which of the following can be used to overcome the overhead of polling for availability of locks?

Choices:

- Semaphores
- Mutex-Locks
- Condition variable
- trylock () function

Parallel and Distributed Computing (CS621)

Question: 9 (Marks: 1) Attempted Questions: 37 Total Questions: 50

The length of the _____in a task dependency graph is called the critical path length.

Choices:

- shortest path
- longest path
- average of longest and shortest path
- intersection of longest and shortest path

Parallel and Distributed Computing (CS621)

Question: 13 (Marks: 1) Attempted Questions: 38

Which of the following decomposition techniques is/are more desirable for quicksort problem?

Choices:

- (●) Recursive decomposition
- ( ) Data decomposition
- ( ) Exploratory decomposition
- Both data and recursive decomposition

Parallel and Distributed Computing (CS621)

Question: 14 (Marks: 1) Attempted Questions: 38 Total Questions: 50

Keep in view of Lustre file system, Identify, which of the following provides file service, and network request handling for one or more local OSTs?

Choices:

- Matadata Server (MDS)
- Matadata Target (MDT)
- Object Storage Server (OSS)
- Object Storage Target (OST)

Parallel and Distributed Computing (CS621)

Question: 15 (Marks: 1) Attempted Questions: 39 Total Questions: 5(

Keep in view of Lustre file system, Identify, which of the following stores file data (chunks of files)?

Choices:

- Matadata Server (MDS)
- Matadata Target (MDT)
- Object Storage Server (OSS)
- Object Storage Target (OST)

Parallel and Distributed Computing (CS621)

Question: 16 (Marks: 1) Attempted Questions: 40 Total Questions: 5

Which of the following is not true regarding POSIX I/O interface?

Choices:

- Standard I/O interface across many platforms
- open, read/write, close functions in C/C++/Fortran
- Mechanism almost all serial applications use to perform I/O
- Many ways for describing collective access

Virtual University Overseas Examination System V2.1 (x86)

Parallel and Distributed Computing (CS621)

Question: 17 (Marks: 1)

Which of the following is not true regarding MPI-IO?

Choices:

- Provides a high-level interface
- Has a large number of routines
- Facilitate Collective I/O
- Facilitate Atomicity rules

Parallel and Distributed Computing (CS621)

Question: 18 (Marks: 1) **Attempted Questions: 42** **Total Questions: 50**

Aggregates smaller read/write operations into larger operations is an example of_____.

Choices:

- POSIX IO
- MPI IO
- Buffered IO
- Non-Buffered IO

Parallel and Distributed Computing (CS621)

Question: 19 (Marks: 1) **Attempted Questions: 43** **Total Questions: 50**

The time elapsed between the beginning and the end of its execution on a sequential computer is known as _____.

Choices:

- serial runtime
- parallel runtime
- distributed runtime
- both sequential and parallel runtime

Parallel and Distributed Computing (CS621)

Question: 20 (Marks: 1) Attempted Questions: 44 Total Questions: 50

Speedup S is computed by _____.

Choices:

- $S = \Theta (n / \log n)$
- $S = \Theta (n \log n)$
- $S = \Theta (\log n)$
- $S = \Theta (\log n)^n$

Parallel and Distributed Computing (CS621)

Question: 21 (Marks: 1) Attempted Questions: 44 Total Questions: 50

The phenomenon when the speedup become greater than p (number of processing elements) is known as _____ speedup.

Choices:

- linear
- superlinear
- non-linear
- supernon-linear

Parallel and Distributed Computing (CS621)

Question: 22 (Marks: 1) Attempted Questions: 46 Total Questions: 50

_____ is the product of parallel runtime and the number of processing elements used (p x T_P).

Choices:

Cost

Overhead

Speedup

Efficiency

Parallel and Distributed Computing (CS621)

Question: 23 (Marks: 1) Attempted Questions: 47 Total Questions: 50

Using fewer than the maximum possible number of processing elements to execute a parallel algorithm is called _____.

Choices:

- scaling up
- scaling down
- scaling out
- scaling in

Question: 24 (Marks: 1) Attempted Questions: 48 Total Questions: 50

_____ is a special class of distributed data-store which is characterized by the lack of simultaneous updates. Here, the emphasis is more on maintaining a consistent view of things for the individual client process that is currently operating on the data-store.

Choices:

- Sequential model
- Client centric consistency model
- Data consistency model
- Incremental model

Question: 25 (Marks: 1) Attempted Questions: 49 Total Questions: 50

_____ the data and moving it closer to where it is needed helps to solve the scalability problem.

Choices:

- Replicating
- Aborting
- Terminating
- Paradigm