24PC1010

High Performance Computing

Objectives:

- Understand High Performance Computing (HPC) system architectures and various computational models.
- Learn basics of CUDA programming.
- Apply parallel execution models and methodologies for parallel programming and parallel applications development.
- Design and implement compute intensive applications on HPC platform.

Course Outcomes:

CO1: Understand High Performance Computing (HPC) system architectures and various computational models.
CO2: Learn basics of CUDA programming.
CO3: Apply parallel execution models and methodologies for parallel programming and parallel applications development.
CO4: Design and implement compute intensive applications on HPC platform.

Module-I

Introduction to High Performance Computing: Era of Computing, Scalable Parallel Computer Architectures, towards low-cost computing, Network of Workstations project by Berkeley, Cluster Computing Architecture, Components, Cluster Middleware and SSI, Need of Resource Management and Scheduling, Programming Environments. **Introduction to High Performance Computing:** Scalable Parallel Computer Architectures, towards low-cost computing, Network of Workstations project by Berkeley, Cluster Computing Architecture, Components, Cluster Middleware and SSI, Need of Resource Management and Scheduling, Programming Environments

Module-II

Cluster Computing: Clustering Models, Clustering Architectures, Clustering Architectures key factors, types of clusters, Mission critical Vs Business Critical Applications, Fault Detection and Masking Algorithms, Check pointing, Heartbeats, Watchdog Timers, Fault recovery through Failover and Failback Concepts. **High Speed Networks & Message Passing:** Introduction to High-Speed Networks, Lightweight Messaging Systems, Xpress Transport Protocol, Software RAID and Parallel File systems, Load Balancing Over Networks – Algorithms and Applications, Job Scheduling approaches and Resource Management in Cluster

Module-III

CUDA Programming: Introduction to CUDA architecture for parallel processing, CUDA Parallelism Model, Foundations of Shared Memory, Introduction to CUDA-C, Parallel programming in CUDA-C, Thread Cooperation and Execution Efficiency, Constants memory and events, memory management, CUDA C on multiple GPUs, Hashing and Natural Parallelism, Scheduling and Work Distribution, Atomics, Barriers and Progress, Transactional Memory

Module-IV

Open CL Programming: Introduction to OpenCL, OpenCL Setup, Basic OpenCL, Advanced OpenCL. **Shared-memory programming:** OpenMP: Introduction to OpenMP, Parallel Programming using OpenMP

Books

1. Rajkumar, High Performance Cluster Computing: Architectures and Systems, Vol. 1 Pearson Education.
2. Georg Hager and Gerhard Wellein, Introduction to High Performance Computing for Scientists and Engineers, CRC Press.
3. Kai Hwang, Advanced Computer Architecture: Parallelism, Scalability, Programmability, McGraw Hill International Editions.
4. Course Name: High Performance Computing Link: https://nptel.ac.in/courses/106/108/106108055/