MAP-I doctoral program

Advanced Computing

Summary

This document presents a Ph.D. level course, as a joint FCUP-DIUM proposal in the MAP-I doctoral program, corresponding to a Curriculum Unit - UCT - credited with 5 ECTS.

I. Theme, Justification and Context

Introduction

Traditionally, grid computing only existed in the realms of high performance computing and technical computer farms. Increasing demand for computing power and data sharing, combined with technological advances, particularly related to network bandwidth increases, has extended the scope of the grid beyond its traditional bounds.

Today, large-scale production Grid infrastructures such as EGEE in Europe, OSG in the US, and NAREGI in Japan are offering their services to many scientific and industrial applications, from domains as diverse as Astronomy, Biomedicine, Computational Chemistry, Earth Sciences, Financial Simulations, and High Energy Physics. Grid infrastructures provide these applications a new means for collaborative research by facilitating the sharing of computational and data resources at an unprecedented scale.

Experience has shown that the technologies involved in sharing resources across the Internet (i.e., grid computing) are very different and their design are more challenging than the technologies used in high-performance cluster computing. Cluster computing builds on knowledge of algorithms, and message passing programming in the C or Fortran programming languages. In contrast, grid computing builds on knowledge of network security, the client/server paradigm, XML, web services and network programming in the Java programming language. Grid computing has many open problems, from which we can mention large scale efficient load balancing and scheduling, fault-tolerance mechanisms in a heterogeneous and diverse environment, failure detection and automatic recovery, workflow languages and tools, “gridfication” tools, among others.

Grid protocols and technologies are being adopted in a wide variety of academic, government, and industrial environments, and there is a growing body of research-oriented literature in grid computing. As so, there is a need for training on the fields of Cluster and Grid Computing, as the next generation of scientists and engineers need to be prepared for a technological workplace that is changing the world.

Cluster Computing

Clustering is a powerful concept and technique for deriving extended capabilities from existing classes of components. In the field of computing systems, clustering is being applied to render new systems structures from existing computing elements to deliver capabilities that through other approaches could easily cost ten times as much. In recent years clustering hardware and software have evolved so that today potential user institutions have a plethora of choices in terms of form, scale, environments, cost, and means of implementation to meet their scalable computing requirements. Some of the largest computers in the world are cluster systems.
Beowulf-class clusters systems have become extremely popular, providing exceptional price/performance, flexibility of configuration and upgrade, and scalability. For many applications they replace previous-generation monolithic vector supercomputers and MPPs.

Beowulf integrates widely available and easily accessible low-cost or no-cost system software to provide many of the capabilities required by system environment. As result of these attributes and the opportunities they imply, Beowulf-class clusters have penetrated almost every aspect of computing and are rapidly coming to dominate the medium to high end.
Computing with Beowulf cluster engages four distinct but interrelated areas of consideration: i) hardware system structure, ii) resource administration and management environment, iii) distributed programming libraries and tools and iv) parallel algorithms.

Hardware system structure encompasses all aspects of the hardware node components and their capabilities, the dedicated network controllers and switches, and the interconnection topology that determines the system's global organization.
The resource management environment is the battery of system software and tools that govern all phases of system operation from installation, configuration, and initialization, through administration and task management, to system status monitoring, fault diagnosis, and maintenance.
The distributed programming libraries and tools determine the paradigm by which the end user coordinates the distributed computing resources to execute simultaneously and cooperatively the many concurrent logical components constituting the parallel application program.
Finally, the domain of parallel algorithms provides the models and approaches for organizing a user's application to exploit the intrinsic parallelism of the problem while operating within the practical constraints of effective performance.

Grid Computing

Our day's computational grids have been accepted as one of the most powerful tools to tackle large-scale problems of computational science and engineering, does promoting the transformation of traditional applications into grid-enabled software packages.
Grid computing is an important new approach to distributed computing that uses geographically distributed computers collectively to achieve higher performance computing and resource sharing. Its combination of regular structures and dynamic algorithms will deliver large and sustained data intensive, knowledge intensive and computational resources across heterogeneous contributing sites. It is enabling rapid advances in many disciplines.
On the front of these applications are domains where scientists need the grid to analyze huge amounts of data as provided by sophisticated experiments such as the Large Hadron Collider at CERN.
Grids underpin the rapidly emerging e-Infrastructure and Cyber infrastructure. It will enable global collaborations to make rapid advances addressing challenges in science, economics, design, engineering and medicine, in fact in all walks of life.
Because grid computing infrastructures are so large and heterogeneous, new challenges arise in diverse areas such as data security, confidentiality, authentication and authorization, load balancing and scheduling, failure detection, prevention and diagnosis, among others.

Description of the Course

Cluster and Grid Deployment (CGD) Curriculum Unit is offered as a fully comprehensive discussion of the foundations and practices for the operation and application of commodity clusters and Grids.
The course is divided into several parts that cover several broad topic areas, including high-performance parallel computing, cluster and grid Computing and cluster management (software and hardware deployment, queuing system and accounting).

Part I begins by presenting Parallel Programming Paradigms and the most common languages, libraries, concepts and techniques for writing parallel application programs and exploit cluster computing.
Next it briefly exposes high-performance computing from the perspective of clusters: it introduces the basic terminology, gives a broad overview of the different architectures and presents some of its limitations. It also describes the hardware components that make up a Beowulf system and shows how to assemble such a system.
Topics also include a brief discussion of Linux for cluster and coverage of the basis of installing and configuring the operating system.
Cluster planning determines what we want a cluster to do and gives an overview of the different types of interconnections communication technologies and cluster file systems alternatives as well as take it out for an initial spin using some readily available parallel benchmarks.

Part II starts describing the tools that may be used to replicate software installed on one machine to others. Thus explaining how to duplicate a pre-configured node tuned installation on a number of machines quickly and efficiently. Finally, in this part we also explain how to manage the resources of Beowulf systems, including system administration and task scheduling, monitoring and accounting.
To illustrate most of the concepts presented in the previous course components, we study and use the SeARCH cluster installed at the Dep. Informatics of Universidade do Minho.

Part III of the course begins by providing an introduction to Grid technologies that underpin e-infrastructure and cyber infrastructure offering an opportunity to hear about the latest achievements from Europe, North America and Asia, and to experience a variety of Grid systems. It will present a conceptual framework to provide students with theory, technology and implementation skills that form the foundation of the field of grid computing.

Part IV builds on the EU Grid Initiative and EGEE project and middleware addressing the concepts and topics needed by site Administrators in charge of middleware installation and configuration.
After discussing the overall *gLite* architecture and the security framework the focus goes to gLite services used for Information and Monitoring, Workload Management, and Data Management.

Finally it reports on experiences gained using the EGEE infrastructure to configure sites at University of Minho and connection to several Virtual Organizations. This part is also dedicated to the study of several aspects of porting existing applications to the grid and development of new applications to the grid.

Related Courses

To our knowledge there is no course equivalent in other Education Institutions. However, there are several other courses and training events, built on recent advances on cluster and grid technologies such as the various editions of the International Summer School on Grid Computing. This school provides in-depth introduction to cluster and grid technologies, through lectures on the principles, technologies experience and exploitation of grids.

II. Objectives

The key distinction between clusters and grids mainly lies in the way resources are managed. In case of clusters, the resource allocation is performed by a centralised resource manager and all nodes cooperatively work together as a single unified resource. In case of Grids, each node has its own resource manager and does not aim at providing a single system view.

We identify a two set of component interesting parts that can help to structure course material in cluster and grid deployment. Each set is constituted by several lectures organized in parts that conclude with a case study.

The work will be challenging but rewarding, enhancing each student's ability to work in this rapidly advancing field and giving participants experience to widely used Grid middleware so that students find themselves to:

- be familiar with the fundamental components of Cluster Grid environments, such as: authentication, authorization, resource discovery and resource access;
- be able to use Cluster Grid environments for basic and advanced job submission, and distributed data management;
- be conversant with Grid achievement worldwide;
- be alert to emerging Grid applications; appreciate the potential of e-infrastructure; and be aware of new research opportunities.

The course is focused in the concepts and technologies administration and implementation aspects needed to provide the requirements for a Grid infrastructure using the current and foreseen capabilities of the EGEE/gLite middleware, running under Scientific Linux, to deploy a dedicated Grid test-bed, consisting on the following components: UI (User Interface), RB (Resource Broker), VO (Virtual Organisation server), RLS (Replica Location Server), CEs (Computing Element), SE (Storage Element), WNs (Worker Nodes).

Target Audience

We will assume that students have diverse backgrounds and build on that diversity. They may come from any country, from computer science and computational science from enthusiasts who have recently heard about this emerging technology to ambitious researchers planning to work on cluster and Grid projects.
Students may be planning to engage in fundamental distributed systems research, or be involved in advanced middleware design and development, to develop new methods in any discipline that depends on the emerging capabilities of e-Infrastructure. Other students may plan to use advanced Grid middleware or be engaged in Grid deployment and operations.

III. Learning Outcomes

By the end of this module, participants should be able to design, and deploy Cluster and Grid based systems and software applications using development and execution environment, such as Rocks and latest Grid software, such as Globus and gLite.
Students should leave the course with:

- A good understanding of the principles and foundations of cluster and Grid computing
- An appreciation of the variations in technologies and views of Cluster and Grid computing
- An awareness of good strategies for using these emerging technologies in their research
- Confidence in using several of the available technologies
- An introduction to the research challenges and the future of Grid computing
- Insight into the architectural implications of Grid-scale computation

- Awareness of current research issues in: Grid architecture and infrastructure
- Integration of applications across autonomous organisations
- Practical experience of current Grid technologies and the associated standards
- Skills in utilising current Grid tools and technologies
- Appreciation of the weaknesses of existing tools and technologies, and potential areas for improvement

In addition student will also gain proficiency into the following areas:

- Cluster Administration - advanced scientific parallel computing on the grid requires the installation and configuration of a cluster that is accessible using grid protocols.
- Parallel Computing - developers of advanced scientific parallel computing on the grid need to know how to use tools and parallel programming languages and libraries to develop and run parallel applications.
- Grid Administration - grid computing requires the installation, configuration, and administration of a grid development and execution environment.
- Security - administrators, developers, and grid users all need to understand the use of public key encryption, certificates, and how security is implemented in the Grid.
- Grid Access and Management – administrators and grid users all need to know how to use commands and tools to submit and execute different grid applications, how to locate, access, and manage data repositories and other grid resources, and how to interface to advanced grid computing toolkits.

IV. Course Contents

Part I. An Introduction to Clusters

Programming Paradigms

- Shared memory
- Distributed memory
- Message passing
- Parallel Programming Languages

Overview of Cluster Computing

- The Role of Clusters
- Definition and Taxonomy
- Distributed Computing
- Limitations

Cluster Planning

- Beowulf Cluster
- Architecture and Cluster Software
- Design Decisions
- Virtualization technologies
- Benchmarks

Part II. Building Custom Clusters

Cloning Systems

- Configuring Systems
- Automating Installations

Workload Management Software

- Queuing /Scheduling
- Monitoring
- Resource Management /Accounting

Cluster Case Study

- *SeARCH* at DI/Univ. Minho**:** with Rocks, Ganglia, Torque/Maui,

Part III. Introduction to Grid Computing

Overview of Grid Computing

- Grid architecture and OGSA/OGSI
- Grid architectures and systems:
- Emerging standards: OGF
- Grid Initiatives and Applications

Grid Computing Components

- Operations Infrastructure
- Middleware and Application
- Virtual Organization management
- Resource Discovery and Info Services
- Data Access Integration and Management
- Web Services
- Grid Portals

Open Problems

- Failure Prevention, Detection, Diagnostics and Recovery
- Environments and Tools for application development
- Gridification of applications
- Workflow languages and tools
- Load balancing and scheduling
- Data integration

Part IV. EU Grid Initiative

EGEE project and middleware

- Authentication & Authorisation
- Job submission service
- Data Management
- Logging & Bookkeeping
- Information Services
- Software stack
- Monitoring and fault recovery

Installing and configuring gLite

- gLite Architecture
- Deployment concepts
- Installation

Grid Case Study

- UMINHP-CO: a semi-automatic gLite/site configuration
- Porting Applications to the Grid

V. Methodology

The overall aim is to give students a broad and well balanced understanding of Cluster and Grid computing that will serve well as a foundation for more specific work or research.
The Course will include lectures and discussions on the principles, technologies, experience and exploitation of clusters and Grids, through which students will receive an integrated and well-structured introduction to the fields of Cluster and Grid computing, its applications and deployment.
Lectures will also review the research horizon and report recent significant successes.
The methodology will also favour critical discussions and reasoning about large-scale distributed system architectures, infrastructures and technologies
Students will be exposed to cluster computing platforms and solutions to study first, how to install, configure and manage computer cluster. Secondly, how to apply high-performance parallel computing using cluster architectures and models, based on the latest technologies such as Beowulf. This will expose principles, research challenges, technical issues and leading views.
Two cases of study of widely used middleware produced by projects in the USA and Europe have been selected to provide experience of collaboratively using Cluster and Grid-based resources, and support the hands-on sessions.
The test bed will be connected to SeARCH and EGEE to provide a rich environment for learning and experimentation. Exercises and team work will encourage students to learn by using this test bed.

There is no set text for this module. Books and research papers will be distributed as required and technical manuals and related documentation will be issued as part of the practical activities. Specific lecture notes for each course unit will be written.

VI. Grading

70% Examination
30% Course work

VII. Bibliographic References

1. *Education: Integrating Parallel and Distributed Computing in Computer Science Curricula*, Marcin Paprzycki, IEEE Distributed Systems Online, vol. 7, no. 2, Feb., 2006
2. *Grid computing in the undergraduate classroom: Topics, exercises, and experiences*, Mache, J., Apon, A., *Proceedings of the 5th IEEE/ ACM International Symposium on Cluster Computing and the Grid*, 2005
3. *SeARCH*, Services and Advanced Research Computing with HTC/HPC clusters", 2005, http://www.di.uminho.pt/search
4. *Classroom exercises for grid services*, *Proceedings of the Linux Cluster Institute International Conference on High Performance Computing*, [Apon, A., Mache, J., Yara, Y., Landrus, K.,2004
5. *EGEE Consortium. EGEE: Enabling Grids for E-Science in Europe,* Project funded by the European Union under contract IST-2003-508833, June 2004
6. *High-Performance Linux Clusters*, with Oscar, Rocks, openMosix & MPI, Joseph D. Sloan, O' Reilly, 2004
7. *Beowulf Cluster Computing with Linux*, Second Edition, William Gropp, Ewing Lusk and Thomas Sterling, MIT PRESS, ISBN 0-262-69292-9, 2003
8. *Grid Computing Making the Global Infrastructure a Reality*, F. Berman, G. C. Fox and A. J. G. Hey editors, Wiley, 2003, ISBN 0-470-85319-0