

Tutorial

Title: Current and Emerging Trends in Cluster Computing

Presenter(s):

Mark Baker

Division of Computer Science
University of Portsmouth
Mercantile House, Hampshire Terrace
Portsmouth
Hants, UK, PO1 2EG
Phone: +44 2392 845826
Fax: +44 2392 843106
Email: mark.baker@computer.org
URL: <http://www.dcs.port.ac.uk/~mab/>

Rajkumar Buyya

School of Computer Science and Software Engineering
Monash University,
C5.10, Caulfield Campus
Melbourne, VIC 3145, Australia
Phone: +61-3-9903 1969
Fax: +61-3-9903 2863
Email: rajkumar@csse.monash.edu.au
URL: <http://www.csse.monash.edu.au/~rajkumar>

Tutorial Abstract:

The availability of high-speed networks and increasingly powerful commodity microprocessors are making clusters of computer an ever more pervasive technology. Clusters consisting of commodity-of-the-shelf (COTS) hardware components as well as commonly used and freely available software are being found in all organisations where high-performance or availability computing is required.

The commercial success of clusters has pushed them into mainstream general purpose computing. This necessitates the need for a software tools and techniques that allows one to manage cluster resources effectively in a manner that is acceptable to majority of users. In this tutorial, we discuss current and emerging trends in cluster computing. In particular we cover cluster technologies in the areas of architectures, networking, single system image, management and application tools. We then move onto talk about four deployed cluster systems that are being used in commerce, industry and research environments. Finally, based on our experience, we summarise our findings, drawing a number of conclusions about current clusters and then briefly discuss emerging technology trends and how these will influence clusters of the future.

Detailed Description:

The availability of high-speed networks and increasingly powerful commodity microprocessors are making the usage of clusters, or networks, of computers an appealing vehicle for cost effective parallel computing. Clusters, built using commodity-of-the-shelf (COTS) hardware components as well as free, or commonly used, software, are playing a major role in redefining the concept of supercomputing. In this tutorial, we discuss the motivation for the transition from

using dedicated parallel supercomputers, to COTS-based cluster supercomputers. We also describe the enabling technologies and then present a number of case studies of cluster-based projects to support our discussion. Finally, we summarise our findings and draw a number of conclusions relating to the usefulness and likely future of cluster computing.

The commercial success of clusters has pushed them into mainstream general purpose computing. Today clusters are not only used for HPC, they have been used in the areas of mission-critical, web serving, and database applications. A number of commercial applications are exploiting commercial applications. In this tutorial we will discuss architectural model of such applications with case studies.

This tutorial will start by discussing the motivation for cluster computing and endeavour to make a clear distinction between it and parallel/distributed computing. We then focus on cluster computer architectures, enabling technologies, hardware and software structures with an emphasis on recipes for building ones own high performance cluster system based different OS, networking, and middleware technologies with example systems. The remaining part of the tutorial focuses on architectural issues related to cluster middleware, the provision of environments with a Single System Image and various programming paradigms; these include distributed shared memory models, message passing as well as Java for HPC. During the discussion of each of these topics, we will present latest developments in that area and possible future issues that researchers and developers need to address.

The tutorial concludes with an open discussion of the work being carried out by various international groups on cluster computing. It is hoped that this discussion will help stimulate interest in the widespread take up of COTS-based clusters for high performance computing.

Outline:

1. Introduction and background
2. Cluster Architectures/Components
 - Linux/Solaris/Windows
 - Middleware
 - Programming Environments
 - Applications
3. Cluster Networking
 - Network technologies (NICs and switches)
 - Ethernet, Myrinet, ATM, SCI, etc.
 - Communication APIs VIA, Active Messages, BIP, etc.
 - Cluster topologies
4. Single System Image alternatives
 - Hardware
 - Operating System
 - Subsystems
 - Tools and Applications
5. Cluster Tools
 - Management
 - Administrative
 - Application development
6. Case Studies
 - Web serving (Hot Bot)

- HPC (CPlant)
- GRID (Internet Applications)
- E-Commerce
- ATLAS, PAPI, and NetSolve

7. Summary and Conclusions

- Near and future trends in clusters

Schedule:

Part 1 (5%): Introduction and background

Part 2 (10%): Cluster Architectures

Part 3 (20%): Cluster Networking

Part 4 (10%): SSI

Part 5 (15%): Cluster Tools

Part 6 (30%): Case Studies

Part 7 (10%): Conclusions

Note - % indicates the percent of overall time dedicated to each topic.

Duration:

Half day tutorial.

Level:

50% Introductory, 30% Intermediate, and 20% Advanced.

Required experience:

The purpose of this tutorial is to overview the current trends and options in cluster-based technologies and systems. The purpose of the case studies is to highlight how the choices made influence the services that can be expected from each cluster type.

Expected audience:

Students, academics, application developers, system designers

Presenter's profile:

Dr Mark Baker

University of Portsmouth, UK

Mark Baker started working in the field of High Performance Computing at Edinburgh University (UK) in 1988. In Edinburgh he was involved in the development of parallel linear solvers on a large Transputer-systems using Occam. From 1990 until 1995 Mark was a project leader of a group at the University of Southampton (UK). This group was involved in developing and supporting environments and tools for a range of parallel and distributed systems. It was whilst at Southampton that Mark started to actively investigate and research software for managing and monitoring distributed environments. In 1995 Mark took up a post as Senior Research Scientist at NPAC, Syracuse University (USA). Whilst at NPAC Mark researched and wrote the widely cited critical review of the Cluster Management Systems¹. At Syracuse Mark worked on a range projects involving the major HPC groups and Labs. in the US. It was during this period that he worked closely with Prof. Geoffrey Fox on a variety of cluster and metacomputing related projects.

1. M.A. Baker, G.C. Fox and H.W. Yau, Review of Cluster Management Software, *NHSE Review*, May 1996 – <http://www.nhse.org/>

Since 1996, Mark has been a Senior Lecturer in the Division of Computer Science at the University of Portsmouth. At Portsmouth Mark lectures on network architectures, client/server programming and open distributed systems. Mark's current research is focused on the development of tools and services for PC-based distributed systems. Mark also tracks international metacomputing efforts and is involved with Java Grande¹ and the definition of a Java interface to MPI².

Mark has recently contributed a number of articles on cluster computing, including a chapter for the Encyclopaedia of Microcomputers³, a paper for Software Practice and Experience⁴ and was the editor and a contributor to a white paper⁵ on cluster computing. Mark is co-chair of the recently established IEEE Computer Society Task Force on Cluster Computing⁶ (TFCC) and is currently a visiting Senior Research Scientist at Oak Ridge National Lab., USA.

Mark is on the international editorial board of the Wiley Journal, Concurrency: Practice and Experience and regularly reviews papers for many journals in his field, including IEEE Computer and Concurrency. Mark gave the Cluster Computing tutorial at HPDC in Los Angeles in 1999. A full list of Mark's activities can be found on his Web site⁷.

Rajkumar Buyya

Monash University, Australia

Rajkumar Buyya is a Research Scholar at the School of Computer Science and Software Engineering, Monash University, Melbourne, Australia. He was awarded Dharma Ratnakara Memorial Trust Gold Medal for his academic excellence during 1992 by Kuvempu/Mysore University. He is co-author of books: Mastering C++ and Microprocessor x86 Programming; and recently, he has edited a two volume book on High Performance Cluster Computing: Architectures and Systems (Vol. 1); Programming and Application (Vol.2) published by Prentice Hall, USA. He served as Guest Editor for the special issues of international journals: Parallel and Distributed Computing Practices, Informatica: An International Journal of Computing and Informatics, and Journal of Supercomputing.

Rajkumar is a speaker in the IEEE Computer Society Chapter Tutorials Program. Along with Mark Baker, he co-chairs the IEEE Computer Society Task Force on Cluster Computing. He has contributed to the development of HPCC system software environment for PARAM supercomputer developed by the Centre for Development of Advanced Computing, India.

-
1. Java Grande – <http://www.javagrande.org/>
 2. mpiJava - <http://www.npac.syr.edu/projects/pcrc/mpiJava/mpiJava.html>
 3. A. Apron and M.A. Baker, Cluster Computing, Encyclopedia of Microcomputers, Edited by Allen Kent and James G. Williams, published by Marcel Dekker, Inc, accepted for publication, Summer 2000 – ISBN 0824727088
 4. M.A. Baker and R. Buyya, Cluster Computing: The Commodity Supercomputer, *International Journal of Software Practice and Experience*, Wiley and Sons Ltd., 29(6), pp. 551-576 1999 – ISSN 0038-0644.
 5. M.A. Baker (editor), A Whitepaper on Cluster Computing, submitted to the *International Journal of High-Performance Applications and Supercomputing*, April 2000
 6. TFCC – <http://www.ieeetfcc.org>
 7. <http://www.dcs.port.ac.uk/~mab/pubs.htm>

Rajkumar conducted tutorials on advanced technologies such as Parallel, Distributed and Multithreaded Computing, Client/Server Computing, Internet and Java, Cluster Computing, and Java and High Performance Computing at international conferences including HPC ASIA'97, Seoul, Korea; HiPC'97, Bangalore, India; NCS'98, Pittsburgh, USA; CATE'98, Cancun, Mexico; ASC'98, Cancun, Mexico; AUUG'98, Sydney, Australia; PDCN'98, Brisbane, Australia; and HICSS-32, Hawaii, USA; HiPer'99, Norway; HPC Asia'2000, Beijing, China; ISCA'2000, Vancouver, Canada; EuroPar'2000, Munich, Germany. He has organised/chaired workshops, symposiums, and conferences at the international level in the areas of Cluster Computing and Grid Computing. He also serves as a reporter for Asian Technology Information Program, Japan/USA. His research papers have appeared in international conferences and journals. His research interests include Programming Paradigms and Operating Environments for Parallel and Distributed Computing.