Author/Editor Information

Full name: RAJKUMAR
Affiliation: Queensland University of Technology, Australia
Address:

School of Computing Science
Queensland University of Technology
825a, ITE Building, Level 8, Gardens Point Campus
Brisbane, Australia. QLD 4001
Phone: +61-7-3864 1290; Fax: +61-7-3864 1801
Email: rajkumar@fit.qut.edu.au
URL: http://www.fit.qut.edu.au/~rajkumar

Author's Professional Background (or resume)

A short biography of a few selected contributing authors is enclosed.

Proposed Book Title: *High Performance Cluster Computing*

Sub-Title (if any): The titles of two volumes of the book can be as follows:

Vol. 1 -- *High Performance Cluster Computing: The System Issues*

Vol. 2 -- *High Performance Cluster Computing: The Programming and Application Issues*

Facts about Your Book

1. Provide a brief and concise overview of your book's subject and purpose.

In the recent years, high speed network and improved microprocessor performance are making clusters or networks of workstations an appealing vehicle for cost effective parallel computing. Clusters built using commodity hardware and software components are playing a major role in redefining the concept of supercomputing.

Parallel computing platforms are shifting from expensive and special purpose MPPs to cheap and general purpose symmetric multiprocessors (SMPs), and clusters of SMPs/workstations/PCs. Using commodity SMPs, workstations and/or PCs, we have many choices to configure a platform for parallel computing for a given budget and for certain types of applications and workloads.

A cluster computer can be viewed from two different perspectives: High Performance and High Availability. Cluster systems can deliver similar or better performance and reliability than traditional mainframes, supercomputers and fault-tolerant systems with a much lower hardware cost. There have been many applications built on cluster systems such as parallel or distributed database applications, telecommunication systems, and recently, Internet or Intranet servers.

In cluster computing environments, parallel and/or distributed computing techniques are applied to the solution of computationally intensive applications across networks of computers.

The proposed book discusses research and technology that unify the fields of parallel processing, distributed computing, and computer networks.

The following sections discuss basics of cluster computing, technology evolution, and trends in cluster computing:

a) Cluster Definition

A cluster is a collection of state-of-the-art, high-performance workstations or PCs that are interconnected with a high-speed and high-capacity data path or network. In other words, a cluster is a group of computers (SMPs/PCs/Workstations) that are loosely connected together to provide fast and reliable services.

b) Technology Transition

Parallel computing has been investigated as a tool for large-scale computing for many years. However, until recently the benefits of this research were confirmed to those few individuals or organizations who had access to a large (and very expensive) parallel computers. Many research organizations have recently moved away from using a few large mainframe computers to a paradigm where individual researchers have access to one or more autonomous personal workstations. Each workstation is dedicated primarily to serving its owner.

We have seen a dramatic increase in workstation performance in the last few years, with workstations doubling in performance every 18-24 months. This is likely to continue for several years, with faster processors and multiprocessor machines. Studies have shown that over half of the workstations CPU cycles are unused even during peak hours. One possible use of these idle cycles is to run parallel applications on a network of computers.

Towards Networked Computing or Cluster Computing: In the 1980s, computer scientists believed computer performance was best improved by creating faster, more efficient processors. This idea was challenged by the idea of clustering which, in essence, means linking together two or more computers to work together on performing functions. The goal is to develop infrastructure (both hardware and software infrastructure) so that end users never know which computer they are actually working-on. This enable multiple computers to work together as one system. Rapid changes in both areas of computing and communication (availability of commodity high-performance microprocessors and high-speed networks) facilitate this transition. These two enabling technologies are making Networks of Workstations (NOW) or Clusters of Workstations (COW) (in short, Network of Computers—PCs/Workstations) an appealing vehicle for parallel processing leading to low-cost *commodity supercomputing*.

The proposed book acts as one point reference for those interested in finding out more about this increasingly common approach to high performance and high availability and of course, with major focus on high performance.

c) Cluster Computer

Cluster Computer is one of the many terms used to describe a distributed system, or network of workstations, when it is used to support loosely-coupled parallel processing and/or high availability. In cluster systems a tension exists between workstation owners, who expect fast response time, and the desires of parallel applications, to soak up as many spare CPU cycles as possible. This can be handled by a new approach called, communal multiprocessing/adaptive parallelism, which emphasizes sharing the processing resources of a cluster system that does away with the concept of node ownership and introduces the need for complex process migration/load balancing systems. This allow clusters to deliver at least the interactive performance of a dedicated workstation while providing the aggregate resources of the network for demanding sequential and parallel programs. Standard programming environments such as Message Passing Interface (MPI) and Parallel Virtual Machine (PVM) are also supported on clusters and thereby allow existing parallel applications to execute on clusters. This means, the existing sequential programs (developed for sequential computer) and parallel programs (developed for specialized parallel computer) run on a cluster computer without any modifications.

d) Clustering for High Performance and High Availability

Currently, cluster technology is increasingly becoming popular and is catching on fast, especially in the corporate world. The reason for this is that clustering offers the following features at low cost:

- High Performance (Commodity Supercomputing)
- Scalability (Scalable Computing)
- High Throughput
- High Availability

Cluster technology permits organizations to boost processing power using commodity hardware and software components—that is, standard technology can be acquired/purchased relatively at low cost. This provides "scalability", an affordable upgrade path that lets organizations increase their computing power, while preserving their existing investment and without incurring a lot of extra expenses. Another benefit of clustering is "failover" capabilities that allow a backup computer to take over the tasks of a failed computer located in its cluster.

e) Levels of Clustering and Future Focus

Individual clusters may be interconnected to form a larger system. In the last few years, several groups have introduced tools to link networked workstations to create a large "virtual" parallel processing computer. High capacity clusters of workstations are becoming an appealing vehicle for parallel computing.

The following are different levels of clustering:

- Department-wide Clustering
- Organization-wide Clustering
- Internet-wide Clustering (Wide-area/Meta Computing)

Researchers are working on a low-cost clustering technology that can run on any kind of network of UNIX/Windows-NT based computer systems.

Other advancements in clustering technology such as "killer-switches" (e.g. Myrinet), ATM, etc., are making computing easier, faster and more reliable. Researchers from many organizations all over the world are approaching clustering from several directions, including improving existing technology to different platforms, and working on Internet-wide clustering that could dramatically improve the reliability of computing over the Internet.

2. Why is your book needed? What is your rationale for writing it? What group(s) constitute the intended audience? In what ways will your book directly benefit this audience (new information, ideas, or philosophy, etc)?

There are many independent efforts in progress in the industry and research labs for developing supercomputers based on the cluster concept. Currently, there are no books available which can act as a reference for researchers (in the industry and academy) and students at the post graduate level. To fill this gap a book on "High Performance Cluster Computing" has been proposed.

Book Audience

The book presents a basic background in cluster computing and illustrate this with a number of examples from current research activities. As there is a strong need for a compact and focused introduction to cluster computing that addresses not only computer scientists but also potential users (especially engineers) in industry.

The targeted audience of the book includes the following:

Graduate and Postgraduate Students, Postdoctoral researchers, Teaching and Research faculties in Industries and Academy, Scientists, Engineers, Mathematicians, and Research/Commercial High Performance Computing and Communication (HPCC) System Developers and Users.

The book will be an important information source for the growing number of researchers, developers and users of High Performance Distributed Computing (HPDC) environments.

3. What will be the most outstanding features of your book?

The book acts as a new standard reference for the entire technical field of Cluster Computing. The most outstanding features of the book are discussed in the following sections:

It is the first of its kind effort to bring together eminent scholars and professional to share their knowledge and experience for the benefit of high performance computing community. It allows the reader to keep up with research trends and innovation in this rapidly growing field.

The Cluster Computing book characterizes the past, present, and future state of theory and practice in the field. In this single volume the reader can find quick reference to various aspects of cluster computing that affect future high performance computing trend. Each chapter describes fundamental principles in generic way (particularly initial chapters and later chapters build on it), best practices, their usage and application with an example systems, research horizons, and references, including key World Wide Web sites. Glossary of key terms, references and sources for further information, including key WWW sites, provide the reader with the most complete information on every topic.

The book's seven sections feature all the most important areas in Cluster Computing:

- Introduction / General Issues
- Networking, Protocols, and I/O
- Programming Environments and Development Tools
- Process Scheduling, Load Sharing and Balancing
- Systems and Architectures
- Java for High Performance Computing
- Algorithms and Applications

More than 100 distinguished scholars and professionals share their knowledge and research and teaching experience of high performance computing as it is practiced around the world. The latest innovations in Architecture, Hardware, Networking, Protocols, Software Design, Programming Languages and Environments, Runtime Environments, Scheduling and Load Balancing, System Management for High Availability and Throughput, and Applications are all in here. The broad scope of the book is to unite theory and practice, filling the gap between introductory textbooks and journals on Parallel, Distributed, and Networked Computing. It is a one-step reference that one working in or studying advanced computer science (parallel, distributed, and networked computing) cannot afford to be without.

4. Are there recent developments in the field that will help us better understand your book and the need for it?

The recent development in the field of cluster computing are adaptive parallelism, high-performance computing with Java, lightweight protocols (active messages), levels of clustering, metacomputing (wide-area computing), and building OS layers on top of traditional operating systems for providing unified access to system resources situated at different locations across the network (Internet/Intranet). The book is covering these recent developments and also its applications (multimedia computing).

The contributing authors of this proposed book are researchers, academics, scientists, research/commercial system investigators having wide experience in designing HPCC hardware and software systems. It is they who have contributed significantly towards the emergence of cluster computing technology through their research work and they have come forward to share their experience in building supercomputers using different architectures and technologies for the benefit of students, co-researchers, commercial system developers, etc.

5. Estimated length in double-spaced manuscript pages or printed book pages (one printed book page is equivalent to approximately 500 words)

Estimated size of printed book pages: 800 (max)

6. Estimated number of drawings/diagrams. Will they be sketches or camera ready (drafted or computer generated)? Estimated number of photographs.

Figures and tables are used for better presentation of concepts dealt in each chapters. They will be provided as camera ready (computer generated) for the publication. The book is expected to have more than 100 computer-generated figures.

7. Expected date of completion of book (month, year)

The camera ready copy of the book is expected to be ready on or before July/August, 1998.

8. Your intended audience (textbook, reference for researchers or practicing professionals, graduate students). Please provide course title and level if it is a textbook.

As the subject area covered by this book is an emerging one, it can serve as a text book for students and reference for researchers or practicing professionals. The proposed book serves as:
1. a tutorial introduction with a review of current research for post-graduate computer scientists, and
2. a concise introduction to cluster computing that addresses potential users, not only computer scientists.

This will be achieved through the closed coordination among contributing authors although they are globally distributed. The book serves as a text book for courses with the following titles (but not limited to):

1. Advances in High Performance Computing (intermediate level)
2. High Performance Computing and Networks (advanced level)
3. An Introduction to Parallel Computing (introductory level)
4. An Introduction to Distributed Computing (intermediate level)
5. Distributed Operating Systems (advanced level)
6. Supercomputing on Clusters (advanced level)
7. Cluster Computing (advanced level)
8. Parallel Programming (intermediate and advanced level)
9. Distributed Systems (advanced level)
10. High Performance Distributed Computing (advanced level)

Note: (i) introductory and intermediate levels means "Undergraduate and Graduate levels. (ii) advanced level means Postgraduate/Advanced-Research level.

9. Competition (any books that you think cover the same material as your book, and why your book is better). Please cite Author, Title, Publisher, and year of publication.

As the subject area covered by this book is an emerging one, we have not noticed many books covering the same material. There is a book entitled, In Search of Clusters, authored by Gregory F Pfister, and published by Prentice Hall, NJ, 1995. This book presents concepts at philosophical/historical level and does not cover in-depth implementation strategies.

The concepts covered in our proposed book is state-of-the-art, timely, and covers a wide range of topics in depth with their design, development, implementation, and customization strategies.

10. Please also attach an outline of the book to this prospectus and enclose as many sample chapters for review as are now available.

The complete table of contents of the book along with a few contributing authors' biographies are enclosed for your kind consideration. The draft of almost all chapters is getting ready and will be provided if required.