Task Force on Cluster Computing (TFCC)

TAB Report Fall 1999

Prepared by
M.A. Baker and R. Buyya
(TFCC Co-Chairs)

Executive Summary

The IEEE Computer Society Task Force on Cluster Computing (TFCC) has been in existence since early in February 1999. In its short life the TFCC has started to have an impact on the cluster computing research and development community in both academia and industry. Some evidence of this can be seen by the numbers influential people willing to get involved in the TFCC activities. Furthermore, the number of TFCC-based events and numbers of willing volunteers prepared to promote our activities all provide ample evidence that the TFCC is not only being successful, but also having an impact on the community

One particularly successful TFCC activity has been our cluster computing educational programme. Here we are attempting to promote the inclusion of cluster computing and its related technologies in the core curriculum of universities around the world. Our efforts in this area also include a book donation programme. In conjunction with influential international authors and publishers we have donated more than 250 books to academic institutions around the world.

The TFCC is also starting to co-organize and sponsor a number of technical events in addition to its own annual event. The first annual (International Workshop on Cluster Computing – IWCC99) event is scheduled to be held in early December 1999, Melbourne, Australia.

Introduction

The TFCC is committed to the advancement of cluster computing research, education, and industry. Towards this end, we have conducted a number of activities including a means for open discussion, bringing together leading experts in the field, book donations, offering advise to academia and industry. We have also hosted workshops and conferences with support from research organisations and industry. This report presents the activities undertaken and planed in the near future by the TFCC.

Organisational Structure

Executive Committee

- Chairs - 2
- Vice Chairs - 2
- Technical Area Coordinators - 17

- Publicity - 6
- Regional Coordinators – 13

Advisory Committee – 59 (includes academia, research labs, and industrial experts)

Technical Areas and Coordinators

- Network Technologies: Salim Hariri (Arizona, USA)
- OS Technologies: Thomas Sterling (Caltech, USA)
- Parallel I/O: Erich Schikuta (Wien, Austria)
- Programming Environments: Tony Skjellum (MPI Softech, USA)
- Java Technologies: Geoffrey Fox (NPAC, Syracuse, USA)
- Algorithms and Applications: Marcin Paprzycki, (USM, USA); and David Bader, Uni. of New Mexico, USA.
- Analysis and Profiling Tools: Dan Reed (UIUC, USA)
- Storage Technologies: Mazin Yousif (IBM, USA)
- High Throughput Computing: Miron Livny (University of Wisconsin at Madison, USA)
- High Availability: Ira Pramanick (Sun Microsystems, USA)
- Single System Image: Toni Cortes (Universitat Politecnica de Catalunya, Spain)
- Performance Evaluation: Jack Dongarra (University of Tennessee at Knoxville and ORNL, USA)
- Software Engineering: Peter Luksch (Technical University, Muenchen, Germany)
- Education: Barry Wilkinson (UNCC, USA)
- Newsletter Editor: Dan Hyde, (Bucknell, USA)
- Industrial Contact: Rawn Shah (Sun World Journal, USA)

Regional Coordinators

TFCC has regional coordinators in all parts of the world and they are actively helping in the establishment of regional activities including the creation of regional web site.

- Africa: Yinong Chen, University of the Witwatersrand, South Africa
- Australia: Hong Shen, Griffith University, Brisbane
- Canada: Harjinder Sandhu, York University, Toronto
- China: Zhiwei Xu, Chinese Academy of Sciences, Beijing
- Europe: Wolfgang Gentzsch, Genias Software, Germany
- Hong Kong: Mounir Hamdi, Hong Kong University of Science And Technology
- India: R. K. Shyamasundar, Tata Institute of Fundamental Research, Bombay
- Japan: Yutaka Ishikawa, Real World Computing Partnership
- Middle-East: Shahriar Shahhoseini, Iran University of Science and Technology
- South America: Ricardo Bianchini, Federal University of Rio de Janeiro, Brazil
- Thailand: Putchong Uthayopas, Kasetsart University, Bangkok
- Taiwan: Chung-Ta King, National Tsing Hua University, Hsinchu
- USA: Amin Vahdat, Duke University

Educational Programme

Educational Resources

Although many educational institutions teach undergraduate and graduate students about the hardware and software technologies that make up a cluster, few courses or programmes concentrate on the wealth of components that constitute the complete cluster environment, from hardware to application development tools. In order to introduce cluster computing into

the curricula of more college programs, the TFCC has set up a web page[1] that has pointers to informative resources, provides links to related journals, books, freely available software, projects from both academia and industry, white papers, and descriptions of hardware components. In addition, with our educational donation programme, we actively support academic faculty members around the world who are interested in introducing new cluster-based courses by providing sample curricular materials.

Book Donation

With the generous cooperation of leading international publishers, we have arranged for the donation of some current books on cluster computing. While the books will be available for faculty members who request them, the TFCC has reserved 50 percent for donation to academic programmes in developing countries. The titles donated include:

Book Title	# Donated
High Performance Cluster Computing: Architectures and Systems, R. Buyya (ed.), *Prentice Hall*, 1999.	50
High Performance Cluster Computing: Programming and Applications, R. Buyya (ed.), *Prentice Hall*, 1999	50
In Search of Clusters, 2nd ed., G. Pfister, *Prentice Hall*, 1998.	25
Metacomputing: Future Generation Computing Systems, W. Gentzsch (ed.), *Elsevier*, 1999	55
Parallel Programming: Techniques and Applications Using Networked Workstations and Parallel Computers, B. Wilkinson and C.M. Allen, *Prentice Hall*, 1998	25
Morgan Kaufmann Publishers	TBA
Cluster Computing: The Journal of Networks, Software Tools and Applications, *Baltzer Science Publishers*, ISSN: 1386-7857	40
SCI - Scalable Coherent Interface: Architecture and Software for High-Performance Compute Clusters, Hermann Hellwagner, Alexander Reinefeld (eds), Springer Verlag, ISBN 3-540-66696-6.	50

Tutorials @ Conferences

- 1999 Conference on High Performance Computing on Hewlett-Packard Systems (HiPer'99), Tromse, Norway, June 27 - 30, 1999.
- Eighth IEEE International Symposium on High Performance Distributed Computing (HPDC-8), Los Angeles, California, USA, August 3-6, 1999.

Tutorials @ IEEE/Computer Society Chapters

- Cluster Computing, IEEE Western Australia Section, Curtin University of Technology, Perth, Australia, March 22, 1999.
- High Performance Cluster Computing, IEEE India Council, Ahmedabad, India, July 11, 1999.
- Parallel Computing on Linux Clusters, IEEE Computer Society Chapter, Baroda, India, July 12, 1999.

[1] `http://www.coe.uncc.edu/~abw/parallel/links.html`

- Java and High Performance Computing, IEEE Section, Bangalore, India, July 13, 1999.
- Performance Cluster Computing, IEEE Hyderabad Section, Hyderabad, India, July 16, 1999.
- Linux Clustering for Parallel Computing, IEEE Computer Society Student Chapter, University of Hyderabad, Hyderabad, India, July 17, 1999.
- Parallel Computing on Linux Clusters, IEEE Computer Society Student Chapter, VNR VJ Institute of Engineering and Technology, Hyderabad, July 17, 1999.
- Supercomputers for All, Annual General Body (AGM) Meeting, IEEE Hyderabad Section, India, July 17, 1999.

TFCC Awareness

Presentations

- *The IEEE Task Force on Cluster Computing (TFCC),* 2nd Workshop on Personal Computer based on Networks of Workstations (PC-NOW), Puerto Rico, USA, April 1999.
- *The IEEE Task Force on Cluster Computing (TFCC),* 4th International Workshop on High-Level Parallel Programming Models and Supportive Environment (HIPS'99), Puerto Rico, USA, April 1999.
- Seminar at University Polytechnica Catalunya, Barcelona, July 1999.
- Seminar at the Institute for High Performance Computing, Singapore, July 1999.
- Seminar at National University of Singapore, July 1999.
- Seminar at Oxford-Brookes University, UK, November 1999.

TFCC Oriented Publications

- M.A. Baker, R. Buyya and D. Hyde, Cluster Computing: A High Performance Contender, *IEEE Computer*, Technical Activities Forum, Vol. 32, No. 7, July 1999 – ISSN 0018-9162.
- M.A. Baker and R. Buyya, Clusters serve up a challenge, *The Age* – newspaper, *Fairfax IT*, August 31st 1999.

Internet Presence

Web Sites

The sample snap shot of the TFCC Web page is shown below. The TFCC also has a logo (see the image between IEEE and CS logos in the TFCC web page). There are three mirrors of the TFCCs Web site:

Australia:
- `http://www.dgs.monash.edu.au/~rajkumar/tfcc/`

UK:
- `http://www.dcs.port.ac.uk/~mab/tfcc/`

USA:
- `http://www-unix.mcs.anl.gov/~buyya/tfcc/`

In additional every technical topic pursued by TFCC has its own web site. The TFCC also has regional web pages managed by local coordinators. These pages highlight cluster computing activities in their regions, they also maintain hot links to projects and research activities in academia and industry.

Mailing lists

We have set up three email lists:

- `tfcc-exe@npac.syr.edu` - a list for inter TFCC executive discussions
- `tfcc-adv@npac.syr.edu` - a list for inter TFCC advisory discussions
- `tfcc-l@bucknell.edu` - a general and open mailing group for discussion and dissemination of TFCC related matter. Currently there are over 361 (as on Oct. 26) people subscribed to this list.

On-Line Magazines

- TFCC newsletter – bi-annual electronic newsletters.
- Genias newsletter

Cluster Computing White Paper

An objective of the TFCC was to act both as a magnet and a focal point for all cluster computing related activities. As such, an early activity that was deemed necessary was to produce a White Paper on cluster computing and its related technologies.

Generally a White Paper is looked upon as a statement of policy on a particular subject. The aim of the White Paper is to provide a relatively unbiased report on the existing, new and emerging technologies as well as the surrounding infrastructure deemed important to the cluster computing community.

This White Paper is meant to provide an authoritative review all the hardware and software technologies that could be used to make up a cluster now or in the near future. These technologies range from the network level, through the operating system and middleware levels up to the application and tools level. The White Paper also tackles the increasingly important area of High Availability as well as Education, which is considered a crucial area for the future success of cluster computing.

The status of the paper is currently at a draft status. We are at the early stage of publicizing its presence and Requesting For Comments (RFC). The paper will be presented to the cluster community at our BOF meeting at SC99 in Portland, Oregon.

White Paper Contents

1. Introduction - Mark Baker (University of Portsmouth) and Rajkumar Buyya (Monash University)
2. Network technologies – Stephen Scott (ORNL), Graham Fagg (UTK) and Salim Hariri (Arizona State)
3. OS - Steve Chapin (Syracuse University) and Joachim Worringen (Aarchen University)
4. SSI – Rajkumar Buyya (Monash University) and Tony Cotes (UPC, Barcelona)
5. Middleware – Mark Baker (University of Portsmouth) and Amy Apon (University of Arkansas)
6. High Availability - Ira Pramanick (Sun Microsystems)
7. Application Tools and libraries – Anne Trefethen (NAG), Shirley Browne (UTK) and Jack Dongarra (ORNL/UTK)
8. Applications – David Bader (UNM) and Robert Pennington (NCSA)
9. Education – Barry Wilkinson (UNC) and Dan Hyde (Bucknell)

International Events

IEEE International Conference on Cluster Computing (IWCC'99)

The IWCC'99 programme consists of number keynote talks, invited talks from industry, regular sessions, a poster session, and also a panel on "Cluster Computing R & D in Australia". The details of the programme organisation and schedule can be found at the workshop web pages[2].

External Grants for TFCC/IWCC'99

- Hewlett Packard (USA)
- MPI Software Technology Inc. (USA)
- Asian Technology Information Program (ATIP, Japan)
- GENIAS Software GmbH (Germany)
- Sun Microsystems Inc. (USA)
- Compaq Computer (Australia)
- Distributed Systems Technology Centre (Australia)
- Monash University (Australia) (Post Graduate Conference Organiser's Grant)

[2] http://www.dgs.monash.edu.au/~rajkumar/tfcc/IWCC99/

TFCC BOFs @ Supercomputing'99 (SC'99) Conference

The TFCC is involved into two BOF meetings at SC99, Portland, Oregon. The first is a fully TFCC-oriented meeting organised by Tim Mattson of Intel Corporation. The purpose of the TFCC BOF is twofold. Firstly we wish to provide a focused activity for the numerous delegates involved or interested in cluster-related activities. Secondly we want to offer a forum to discuss and debate the TFCCs many activities: including the Cluster Computing white paper and future directions and events. The SC99 BOF agenda is shown below:

1. Intro to the BOF: 2 minutes (Tim Mattson, Intel)
2. Intro to the TFCC: 10 minutes (Rajkumar Buyya, Monash University)
3. The white paper: 10 minutes, Mark Baker (University of Portsmouth)
4. Panel - where clusters are going and what's needed. Moderator: (Mark Baker, University of Portsmouth)
    a. Vendor point of view (Tim Mattson, Intel): 10-15 minutes
    b. Integrator point of view (Greg Lindahl, UVa): 10-15 minutes
    c. Research community (Thomas Sterling, JPL): 10-15 minutes
5. Open discussion - 30 minutes

The TFCC is involved in supporting the BOF on "High Speed Interconnects for COTS Cluster Computing[3]".

TFCC Conference Activities

1999

- IEEE International Workshop on Cluster Computing (IWCC'99), TFCC's First Annual Meeting, Melbourne, Australia, Dec. 2-3, 1999.

Cooperative Events:

- IEEE International Symposium on High Performance Distributed Computing (HPDC-8), Los Angeles, California, USA, August 3-6, 1999.

2000

- IEEE International Conference on Cluster Computing (Cluster'2000), TFCC's Second Annual Meeting, Chemnitz, Germany, December 2000 – A merger of German CC and PC-NOW workshop series with IEEE IWCC, Task Force's First Annual Meeting held in Australia.

Cooperative Events:

- International Workshop on Personal Computer based Networks Of Workstations (PC-NOW'2000) – Cancun, Mexico May 2000 along with IPDPS'2000.
- Asia-Pacific International Symposium on Cluster Computing (APSCC'2000)– Bejing, China, May 2000 along with HPC Asia'2000.
- HPCN 2000 Cluster Computing Workshop, May 2000, Amsterdam, The Netherlands.
- Workshop on Cluster Computing for Internet Applications (CCIA2000) – Iwate, Japan, July 2000 along with ICPADS'2000.

[3] http://www.atoll-net.de/news-bofsc99.html

- Session: Cluster Computing- Technologies, Environments, and Applications (CC-TEA'2000) – Las Vegas, USA, July 2000 (along with PDPTA'2000).
- Cluster Computing Workshop – Munich, Germany, September 2000 along with EuroPar'2000.
- International Workshop on Grid Computing, Bangalore, India, December 2000 (along with HiPC'2000).

2001

- IEEE International Conference on Cluster Computing (Cluster'2001), Los Angles, USA

Commercial Collaborations and Activities

The TFCC is actively collaborating with international industry, research, and commercial organizations to promote the field of cluster computing. Their involvement can be seen in the form of personnel support for TFCC activities, strategic collaboration, or financial support to help conduct workshops and conferences. The financial support from universities, research organizations, and commercial vendors has made it possible for us to host TFCC conferences with highly subsidized registration fees and thus enable participation of large number of delegate, including students. The external and commercial collaborations include the following:

- Financial support for TFCC/IWCC99 from a number of commercial entities.
- Conference finance and management support for IWCC'99, from Monash University.
- "Cluster Computing R&D in Australia: A Status Report" – a report being prepared for ATIP, Japan in return for supporting IWCC'99 financially.
- Thomas Sterling is being supported by Microsoft to attend IWCC99.
- International publishers have donated books to the TFCCs educational programme.
- Book donation administrative support is offered by GENIAS Software (Germany).
- Industrial promotion – advice to commercial cluster computing vendors' in particular startup companies.
- Journal collaborations to bring out special issues on cluster computing.
- TFCC White paper.

Conclusions

This report has highlighted major activities of the TFCC since its formation early this year. It has been fruitful year for the TFCC in terms of the many successful initiatives in promoting the field of cluster computing. Interestingly, the TFCC has been able to conduct all the above activities without using the funding allocated to it by the Society (IEEE Computer Society). We strongly believe that without the support of the TFCC volunteers and their employers, it would have been impossible for the TFCC to make such progress within a short period of time. We take this opportunity to acknowledge the support of all the TFCC members, executive and advisory committees, volunteers and those actively participating in open dialogues on the TFCC public mailing list. The success of the TFCC discussion forum is evident from its citation in research publications and on some occasions in Ph.D. theses/dissertations. The support of the IEEE Computer Society, TAB, and volunteer coordinator division is acknowledged.

In 2000, we plan to strengthen the TFCC and initiate new activities by recruiting highly active volunteers for the management committee (executive/advisory). We are confident that the Society will consider TFCCs contributions and promote it to the status of a *Technical Committee* at some stage in the not so distant future. The TFCC is looking forward to a

highly successful future and we believe strongly that its activities will go from strength to strength.