

# The Task Force on Cluster Computing (TFCC)

## TFCC 2003 - Annual Report

December 2003

(Activities over the last year)

**Mark Baker and Rajkumar Buyya**

TFCC Co-Chairs

1. Introduction .....	1
2. Web Presence .....	1
3. Major TFCC Sponsored Conferences and Events .....	2
3.1 International Conference on Cluster Computing .....	2
3.1.1 Steering Committee: .....	2
3.1.2 Events .....	2
3.2 International Conference on Cluster Computing and the Grid .....	3
3.2.1 Steering Committee .....	3
3.2.2 Events: .....	3
3.3 Grid Workshop .....	3
3.3.1 Steering Committee: .....	3
3.3.2 Events .....	3
4. TFCCs International Presence .....	7
5. TFCC Awards .....	7
6. TOP500 and TFCC Initiative on World's TopClusters List .....	7
7. TFCC Membership .....	8
8. Summary .....	8

## 1. Introduction

The TFCC has continued to build on its previous successes. Over the last year, the TFCCs three main events; the International Conference on Cluster Computing (ICCC), the International Symposium on Cluster Computing and the Grid (CCGrid) and the International Grid workshops, held in conjunction with Super Computing, have all continued to be successful, and attract high class submissions and delegates from around the world. All these events are now well established, and significant events in their own rights.

The reformatted TFCC newsletter is big success and attracting a widespread readership and also sponsors from various international vendors. The success of the newsletter is wholly due to the hard work and enthusiasm of its editor. The TFCCs book donation programme continues to run and we have distributed various publications to our less well off colleagues around the world. The following sections of this report detail particular activities that the TFCC has been involved in and with over the last year.

## 2. Web Presence

The TFCCs Web site was redesigned in the fall of 2002. The TFCCs home page is <http://www.ieeetfcc.org>. There are three mirror sites at:

- AU: <http://www.cs.mu.oz.au/~raj/tfcc/>
- UK: <http://dsg.port.ac.uk/ieee-tfcc/>

### 3. Major TFCC Sponsored Conferences and Events

#### 3.1 International Conference on Cluster Computing

##### 3.1.1 *Steering Committee:*

- Mark Baker, University of Portsmouth
- Bill Blake, Compaq Corp.
- Giovanni Chiola, University of Genoa, Italy
- Jack Dongarra, University of Tennessee, USA
- Geoffrey Fox, Indiana University, USA
- Al Geist, Oak Ridge National Lab, USA
- Rusty Lusk, Argonne National Lab, USA
- Greg Pfister, IBM, USA
- Wolfgang Rehm, Technical University of Chemnitz, Germany
- Thomas Sterling, Caltech/JPL, USA,
- Rick Stevens, Argonne National Lab, USA
- Thomas Stricker, ETH, Switzerland
- Satoshi Matsuoka Tokyo Institute of Technology, Japan
- Daniel S. Katz, JPL, USA
- Rajkumar Buyya, University of Melbourne
- William "Bill" Camp, Sandia National Lab, USA
- Pete Beckman, Argonne National Lab, USA

##### 3.1.2 *Events*

IEEE International Conference on Cluster Computing (Cluster 2002)

<http://www-unix.mcs.anl.gov/cluster2002/>

23 – 26, September 2002, Chicago, USA

General Chairs: Daniel Reed, NCSA, USA and Mark Baker, University of Portsmouth, UK

Papers received: 116 (45 accepted as full papers and 12 as posters)

Attendees: approx 300

IEEE International Conference on Cluster Computing (Cluster 2003)

<http://www.csis.hku.hk/cluster2003>

1 - 4, December 2003, Hong Kong

General Chairs: Jack Dongarra, University of Tennessee at Knoxville, USA and Lionel Ni, The University of Hong Kong.

The 6 best papers will be published in a special issue of the International Journal of High Performance Computing and Networking.

IEEE International Conference on Cluster Computing (Cluster 2004)

<http://grail.sdsc.edu/cluster2004/>

September 20-23 2004, San Diego, California

General Co-Chairs: Henri Casanova, University of California, San Diego, USA, Daniel S. Katz, Jet Propulsion Laboratory, USA.

It is likely that Cluster 2005 will be held in Boston and be hosted by either Harvard or MIT.

## 3.2 International Conference on Cluster Computing and the Grid

### 3.2.1 *Steering Committee*

- Mark Baker, University of Portsmouth, United Kingdom
- Henri Bal, Vrije University, The Netherlands
- Rajkuma Buyya, University of Melbourn, Australia (Chair)
- Franck Capello, University of Paris-Sud, France
- Charlie Catlett, Argonne National Laboratory, USA
- Jack Dongarra, University of Tennessee & Oak Ridge National Laboratory, USA
- Ian Foster, Argonne National Laboratory, USA
- Wolfgang Gentsch, Sun Microsystems, USA
- Hai Jin, Huazhong University of Science & Technology, China
- Craig Lee, The Aerospace Corporation, USA
- Laurent Lefevre, INRIA, France
- Satoshi Matsuoaka, Tokyo Institute of Technology, Japan
- Alexander Reinefeld, Konrad-Zuse-Zentrum, Germany
- Satoshi Sekiguchi, AIST, Japan
- David Walker, University of Wales, Cardiff, United Kingdom

### 3.2.2 *Events:*

3<sup>rd</sup> IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2003)  
<http://www.ccgrid.org/ccgrid2003/> | <http://ccgrid2003.apgrid.org/>  
May 12-15, 2003, Toshi Center Hotel, Tokyo, JAPAN  
General Co-chairs, Sangsan Lee (KISTI, Korea) and Satoshi Sekiguchi (AIST, Japan)  
Associated Workshops on Emerging Topics: 7

4<sup>th</sup> IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2004)  
<http://www.ccgrid.org/ccgrid2003/> | <http://www.ccgrid.org/ccgrid2004/>  
April 19-22, 2004, The Drake Hotel, Chicago, Illinois, USA  
General Co-Chairs Charlie Catlett and Pete Beckman, Argonne National Laboratory, USA

## 3.3 Grid Workshop

### 3.3.1 *Steering Committee:*

- o Mark Baker, University of Portsmouth, UK
- o Rajkumar Buyya, Monash University, Australia
- o Craig Lee, The Aerospace Corporation, USA
- o Manish Parashar, Rutgers University, USA
- o Heinz Stockinger, CERN, Switzerland

### 3.3.2 *Events*

3<sup>rd</sup> International Workshop on Grid Computing, <http://www.gridcomputing.org>  
Monday 18th November 2002, Baltimore, MD held in conjunction with Supercomputing 2002  
Program chair: Manish Parashar, Rutgers, The State University of New Jersey

Proceedings published by Springer-Verlag. The best 6 papers in the workshop's focus area of Grid Applications were chosen for journal-length papers in a special issue of the International Journal of High Performance Computing Applications to be published in mid 2003.

Papers: 78 submitted (22 full paper and 6 short papers accepted)

Funding: NSF \$5000 towards publication

Attendance: 150+.

4<sup>th</sup> International Workshop on Grid Computing (Grid2003),  
<http://www.gridcomputing.org/grid2003>, Monday, November 17, 2003, Phoenix, Arizona,  
(held in conjunction with Super Computing 2003)

General Chair: Craig Lee, Aerospace Corp., USA, Program Chair: Heinz Stockinger, CERN, Switzerland

Papers: 67 submitted; 24 accepted (21 as long and 3 as short papers)

Attendance: 396

IEEE Computer Society Press will publish proceedings of the workshop. The best papers are will be extended, peer reviewed again and published in the Journal of Grid Computing published by Kluwer.

5<sup>th</sup> International Workshop on Grid Computing (Grid2004),  
<http://www.gridcomputing.org/grid2004>, Pittsburgh, USA (held in conjunction with Super Computing 2004)

## **1. Book Donation Program**

<http://www.ieeetfcc.org/edupromote.html>

Over 1300 books have been donated to the TFCC who under its Book Donation Programme have passed these onto to University faculties all over the world. This programme is in co-operation with many international publishers and authors. Books distributed over the last year include:

- Cluster Computing Special Issue, R. Buyya, H. Jin, T. Cortes (Guest editors), Future Generation Computer Systems, Vol. 18, No. 3, Elsevier, Jan 2002.
- Alex Vrenios, Linux Cluster Architecture, Sams Publishing, USA, 2002.
- Proceeding of IEEE International Conference on Cluster Computing (Cluster 2003), Hong Kong, CS Press, 1 - 4, December 2003,

## **2. TFCC New (Bi-annual) Letter – April-Oct and Nov-March issues**

The TFCC Newsletter has been reorganised substantially since the beginning of this year. The mission of the Newsletter is to contribute to the TFCC's educational objectives by offering a high-quality information resource on cluster computing to the general technical public. The Newsletter is available in electronic format at <http://www.clustercomputing.org>. Four issues of the TFCC Newsletter were published between September 2002 and November 2003, with another issue slated to go online in late November 2003.

The Newsletter started as an electronic bulletin board, announcing cluster computing-related news items, conferences, and other academic activities twice yearly. The Newsletter's initial editor was Dan Hyde (Bucknell University). In 2002, Frank Sommers (Autospaces, L.L.C.) took over as volunteer Newsletter editor. Currently, two graduate students provide part-time editorial and production support.

In keeping with the Newsletter's mission, throughout 2002 and 2003, several features were added to the Newsletter. The following features were created to bring more timely and relevant information to a growing readership.

**Publication program:** In 2002, we commenced on a publication plan to bring to our readers in-depth technical articles on cluster and grid computing. As of the time of this report, 17 full-length technical articles appeared in the Newsletter, with 2 additional articles to appear in the November/December, 2003, issue. The articles range from academic research reports to industry tutorials. In addition, the Newsletter published 1 book review, and several shorter announcements from various research projects. Several of the articles that appeared in the Newsletter were subsequently discussed on technical and developer Web sites, such as [theserverside.com](http://theserverside.com), [slashdot.org](http://slashdot.org), and [artima.com](http://artima.com).

**Real-time conference listing:** We have automated the listing of academic conferences relevant to cluster computing. Any Newsletter reader can post conference announcements, including all the important dates associated with a conference. That announcement will then immediately appear on the Newsletter's home page, as well as in the Newsletter's conference listings section. As of the time of the present report, 71 academic conferences and workshops have been listed in the Newsletter.

**Industry news and announcements:** Similar to the conference listing, the Newsletter also provides a real-time listing of industry news and announcements. To date, 116 news announcements were entered and published in the Newsletter.

**Resources section:** We have also added a section to the Newsletter that lists cluster and grid resources. Among the most popular items accessed via the Newsletter is the PDF version of the TFCC's 2000 Cluster Computing White Paper, a copy of which is made available in the Resources section.

**Recruitment of new TFCC members:** In addition to being an information resource, the Newsletter's additional aim is to actively spark interest in cluster computing among its readers, and to invite those interested in cluster computing to join the TFCC. To that effect, the Newsletter features an online TFCC membership sign-up form. Between September 2002 and November 2003, 109 new members joined the TFCC via the Newsletter. Those new members represent 28 different countries.

**Readership:** Due in part to the new features introduced, the Newsletter's readership has been on a steady rise. Between December 2002, and October 2003, 48,330 unique visits registered on the Newsletter's Web site, representing 121,565 unique page views. Average monthly readership has been increasing from 620 unique visits in January 2003, comprising 2,074 page views, to 8,558 unique visits, or 28,162 page views, in October 2003. While we do not keep cumulative per country readership data, we do have monthly statistics available based on a reader's country of origin, as determined by the IP address of a reader's Web browser. In October 2003, readers from 82 countries visited the Newsletter's Web site. Leading among those are readers with U.S. commercial domains ("com") (23.83%), U.S. "net" domains (15.06%), German domains (2.68%), U.S. educational domains (1.49%), Italian domains (1.6%), U.K. domains (1.58%), and Brazilian domains (1.21%).

**Commercial Sponsorships:** In order to maintain a high quality publication, and to provide new features to the increasing number and diversity of our readers, in 2003 we sought to

establish a budget for the Newsletter. The Newsletter received initial financial support from the TFCC. At the time of this report, seven sponsors, Cyclades, Myricom, Linux Network, Microsoft, Foundry Networks, and RLX technologies have provided additional support. The objective of seeking commercial sponsorships is to ensure that the Newsletter become a self-sufficient TFCC effort.

**Future Directions:** In 2004, we would like to expand our publication program, and make new issue releases more frequent. Based on the current level interest from authors of articles, we believe that we will be able to publish 1-2 high-quality articles each month. Among the new features planned is a set of online discussion forums dedicated to various topics of interest to cluster computing professionals. Those forums would supplement the TFCC's email list, [tfcc@clustercomputing.org](mailto:tfcc@clustercomputing.org), which is used primarily for conference announcements. The forums will, hopefully, encourage technical discussion about cluster computing. In addition, we are planning significant enhancements to the Newsletter's cluster computing resources section. It is our hope that these initiatives will further contribute to the TFCC's mission.



Figure 1: Cover page for TFCCs November/December, 2003 Newsletter.

### 3. ACM/LANL CoRR and TFCC – Cluster Computing Archive

The Computing Research Repository (CoRR - <http://www.ieeetfcc.org/ClusterArchive.html>) on Cluster Computing is an initiative of the TFCC in collaboration with the ACM, the Los Alamos e-Print archive, and NCSTRL (Networked Computer Science Technical Reference Library). The aim of the CoRR e-Print Archive is to promote science and technology by providing an information repository that is accessible to colleagues all over the world. This Internet-based digital library has been mirrored in many countries to enable faster access to wealth of information.

This archive operates as part of the subject area known as “Distributed, Parallel, and Cluster Computing” and is moderated by Rajkumar Buyya (University of Melbourne) and Tushar Chandra (IBM). Since the starting this Archive program in 2001, a number of articles on Cluster and Grid Computing have been submitted.

#### 4. TFCCs International Presence

The TFCC have extended their regional activities by inviting new regional representative to co-ordinate regional activities. Each of these co-ordinators maintains their own TFCC regional Web sites. Table 1 lists our regional co-ordinators.

<b>Africa</b>	Yinong Chen, University of the Witwatersrand, South Africa
<b>Australia</b>	Hong Shen, Griffith University, Brisbane
<b>Canada</b>	Harjinder Sandhu, York University, Toronto
<b>China</b>	Hai Jin, HUST, China / Uni of Southern California, USA
<b>Europe</b>	Wolfgang Gentzsch, Genias Software, Germany
<b>Hong Kong</b>	Cho-Li Wang Hong Kong University, HK
<b>India</b>	S. Chakravarthi, MPI Software Technology India, Bangalore and R. K. Shyamasundar, Tata Institute of Fundamental Research, Bombay
<b>Israel</b>	Orly Kremien, Bar-Ilan University, Israel
<b>Japan</b>	Yutaka Ishikawa, Real World Computing Partnership
<b>Middle-East</b>	Shahriar Shahhoseini, Iran University of Science and Technology
<b>South America</b>	Ricardo Bianchini, Federal University of Rio de Janeiro, Brazil
<b>Thailand</b>	Putchong Uthayopas, Kasetsart University, Bangkok
<b>Taiwan</b>	Chung-Ta King, National Tsing Hua University, Hsinchu
<b>USA</b>	Amin Vahdat, Duke University

**Table 1: TFCC Regional Co-ordinators**

#### 5. TFCC Awards

In 2003, we have presented two volunteer service awards:

- Frank Sommers, Autospaces and University of Southern California, USA Newsletter Editor, **Citation:** For transforming the IEEE Task Force on Cluster Computing (TFCC) newsletter into a state-of-the-art community technical news portal.
- Ms. Monika Grobecker, Sun Microsystems GmbH, Secretariat, TFCC Book Donation Programme, **Citation:** For providing outstanding secretarial services to the IEEE Task Force on Cluster Computing (TFCC)'s educational book donation program.

The TFCC in conjunction with the TCPP presented six \$1000 travel awards for students to attend Cluster 2003 and HiPC 2003. The \$6000 for these awards was made up of \$2000 given to the TFCC from the profits of the MIT Beowulf books series (with thanks to Drs Thomas Sterling and William Gropp) and \$4000 from the TCPP.

#### 6. TOP500 and TFCC Initiative on World's TopClusters List

The TFCC initiated a collaboration with the organisers of the TOP500 ([www.top500.org](http://www.top500.org)) supercomputers to develop a similar list for clusters (<http://www.TopClusters.org>). In addition to numerical benchmarks, we are looking into other parameters such as System I/O, Web, Database, and Application level performance for clusters. Recently this portal has been re-branded as “Clusters @ TOP 500” and offers a wealth of information on leading trends in cluster computing and hot news.

## 7. TFCC Membership

- Over 300<sup>1</sup> members are registered directly with the CS.
- 700 on the TFCC mailing list (tfcc@clustercomputing.org) as of December 2003.

## 8. Summary

The TFCC sponsors three main events per year, the International Conference on Cluster Computing, the International Conference on Cluster Computing the Grid, and the International Workshop on Grid Computing. These events are now well established, cater for around about 750 delegates per year, attract 350-paper submissions, and are commercially sponsored to the order of \$70K. These TFCC sponsored events are independently managed annual events that have their own steering committees and have selected venues for several years ahead. The popularity of these events demonstrates that the TFCC has a solid presence in both the academic and commercial communities.

The TFCC covers both Cluster and Grid computing for its international community. The profile of the TFCCs membership reflects the importance of the areas that it supports. It can be clearly seen, for example by examining its executive and advisory committees, or by looking at the makeup of the various committees associated with its events, that its members are international leaders and experts in their fields of expertise.

As a CS Task Force is meant to be ephemeral, either dissolving or transitioning to a TC, the TFCC is currently in limbo. The TFCC still has hopes of becoming a full Technical Committee on Clusters and Grids, but our initial request to a change status to a TC, which was tabled in 2002, was turned down. However, over the last year matters have moved on somewhat, and there is a debate underway between members of the TCPP, TCDP and the TCSA, about possible ways to break the current stalemate. Much of the current debate is about the overlaps between the three TCs (mentioned above) and the TFCC. We feel that overlaps are inevitable, but what is more important is that we fulfil the needs of the community that we try to serve. It is clearly evident that the TFCC has made an impact, and should continue to serve its Cluster community. Equally, there is a gap in the coverage of volunteer activities, that of serving the Grid community more fully.

One of the possible ways that is being currently being explored is the possibility of merging the TCSA and the TFCC into a Technical Committee. Our analysis of the recent trends in computing, as well as the objectives of these two volunteer bodies and their focus, indicate that such a move could be very beneficial to the communities they serve. In addition it would solve the important issue at hand – the status of the TFCC. This move, if successful, should be viewed as a precursor and a stepping-stone towards exploring the possibility of creating a council that could assimilate all four (at this time and/or three in the future) bodies.

In summary, the TFCC has continued its very successful activities. The events that the TFCC runs continue to be well attended and their future seem secure. The TFCC continues to help focus both the Cluster and Grid communities, and we feel that this in turn reflects well on the IEEE CS well.

---

<sup>1</sup> This is an estimate, based on old data.