

---

## Special issue: Cluster 2003

---

### Daniel S. Katz\*

Jet Propulsion Laboratory,  
California Institute of Technology,  
Pasadena, CA 91109, USA  
E-mail: d.katz@ieee.org  
\*Corresponding author

### Cho-Li Wang

Department of Computer Science,  
University of Hong Kong, Hong Kong  
E-mail: clwang@cs.hku.hk

---

In December 2003, almost 200 delegates gathered in Hong Kong to participate in Cluster 2003, the 2003 IEEE International Conference on Cluster Computing. The success of this event was partially due to the accumulated hard work of the organisers, in the face of a recent health crisis. In particular, the local organisers did a tremendous job of creating an event that brought the life, food and excitement of Hong Kong and China into the meeting. However, the organisers did not act alone in making this conference unique. The rest of the credit for this successful meeting goes to the contributors, exhibitors and attendees who created a diverse and stimulating environment for those from commercial, academic and government backgrounds.

Cluster 2003 had a very strong technical program, perhaps the best of any of the Cluster conferences. Forty-eight papers were selected from the 164 papers submitted (a 29% acceptance rate). This strength, and the fact that there was a clear grouping of scores by the technical program committee, made it quite easy to choose a short list of what were considered the best papers. The authors of these papers were asked to extend their papers and resubmit them for possible inclusion in this issue. This issue contains the 13 papers that successfully completed this process.

The topics of these 13 papers fall into the areas of applications, scheduling, parallel I/O, communications middleware and cluster management. The applications papers discuss: the performance analysis of a cosmology application on a set of three clusters; the design of a bioinformatics application with a fine-grained, multi-threaded model as an improvement over the application implemented with PVM; and the design and analysis of an efficient out-of-core matrix transpose algorithm. There are three papers on scheduling, with topics: scheduling aperiodic parallel real-time jobs onto a heterogeneous cluster running a set of period real-time jobs; self-adaptive scheduling for heterogeneous clusters; and design and simulation of a strategy for scheduling of mouldable jobs. Within parallel I/O, the papers discuss: an evaluation of structured I/O methods for parallel file systems; the design of a file system to be shared across

administrative domains; and the design of a new redundancy scheme for cluster file systems that includes high-performance user access. The communications middleware papers include topics: design and performance evaluation of mechanisms for latency tolerance in remote memory accesses supported by communications hardware; design of collective operations that bypass the application to avoid reduced performance due to process skew; and a performance comparison of coordinated checkpointing versus message-logging as methods of building fault-tolerant MPI applications. Finally, the cluster management paper discusses the use of reusable mobile agents.

As mentioned previously, this set of papers are those that were considered the best by the technical program committee. The full list of papers in the conference included nine papers on scheduling, six papers each on cluster/job management, communications, and I/O, three papers each on applications, Java, middleware, novel systems, high-availability and performance/analysis, as well as six vendor talks.

Cluster 2003 also included about four tutorials, 20 posters and four keynote presentations on the use of clusters for large scientific simulations, the use of clusters in the financial community, movement toward the federation of clusters and grids and grid security.

As can be seen by this diverse range of topics, there is a wide range of research ongoing in the field of Cluster Computing. This research continues to attempt to increase the performance, usability and acceptance of clusters. There are clearly a wide variety of applications and users that are looking to clusters as the most cost effective solution to their problems today.

Readers are invited to participate in future meetings in the Cluster series, including Cluster 2005 to be held in Boston in September 2005, and Cluster 2006 to be held in Barcelona in September 2006. Further information on the series of conferences can be found at <http://www.clustercomp.org/>.

Finally, the editors would like to thank all the people involved in Cluster 2003, including the organisers, reviewers, contributors, exhibitors and last but not least, the generous sponsors.