

Bernd Mohr
Jesper Larsson Träff
Joachim Worringen
Jack Dongarra (Eds.)

Recent Advances in Parallel Virtual Machine and Message Passing Interface

13th European PVM/MPI User's Group Meeting
Bonn, Germany, September 2006
Proceedings



Springer

Bernd Mohr Jesper Larsson Träff
Joachim Worringen Jack Dongarra (Eds.)

Recent Advances in Parallel Virtual Machine and Message Passing Interface

13th European PVM/MPI User's Group Meeting
Bonn, Germany, September 17-20, 2006
Proceedings



Volume Editors

Bernd Mohr
Forschungszentrum Jülich GmbH
Zentralinstitut für Angewandte Mathematik
52425 Jülich, Germany
E-mail: b.mohr@fz-juelich.de

Jesper Larsson Träff
C&C Research Laboratories NEC Europe Ltd.
Rathausallee 10, 53757 Sankt Augustin, Germany
E-mail: traff@ccrl-nece.de

Joachim Worringen
Dolphin Interconnect Solutions ASA
R&D Germany
Siebengebirgsblick 26, 53343 Wachtberg, Germany
E-mail: joachim@dolphinics.com

Jack Dongarra
University of Tennessee
Computer Science Department
1122 Volunteer Blvd, Knoxville, TN 37996-3450, USA
E-mail: dongarra@cs.utk.edu

Library of Congress Control Number: 2006931769

CR Subject Classification (1998): D.1.3, D.3.2, F.1.2, G.1.0, B.2.1, C.1.2

LNCS Sublibrary: SL 2 – Programming and Software Engineering

ISSN 0302-9743
ISBN-10 3-540-39110-X Springer Berlin Heidelberg New York
ISBN-13 978-3-540-39110-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2006
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11846802 06/3142 5 4 3 2 1 0

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Preface

Since its inception in 1994 as a European PVM user's group meeting, EuroPVM/MPI has evolved into the foremost international conference dedicated to the latest developments concerning MPI (Message Passing Interface) and PVM (Parallel Virtual Machine). These include fundamental aspects of these message passing standards, implementation, new algorithms and techniques, performance and benchmarking, support tools, and applications using message passing. Despite its focus, EuroPVM/MPI is accommodating to new message-passing and other parallel and distributed programming paradigms beyond MPI and PVM. Over the years the meeting has successfully brought together developers, researchers and users from both academia and industry. EuroPVM/MPI has contributed to furthering the understanding of message passing programming in these paradigms, and has positively influenced the quality of many implementations of both MPI and PVM through exchange of ideas and friendly competition.

EuroPVM/MPI takes place each year at a different European location, and the 2006 meeting was the 13th in the series. Previous meetings were held in Sorrento (2005), Budapest (2004), Venice (2003), Linz (2002), Santorini (2001), Balatonfüred (2000), Barcelona (1999), Liverpool (1998), Cracow (1997), Munich (1996), Lyon (1995), and Rome (1994). EuroPVM/MPI 2006 took place in Bonn, Germany, 17 – 20 September, 2006, and was organized jointly by the C&C Research Labs, NEC Europe Ltd., and the Research Center Jülich.

Contributions to EuroPVM/MPI 2006 were submitted in May as either full papers or posters, or (with a later deadline) as full papers to the special session ParSim on "Current Trends in Numerical Simulation for Parallel Engineering Environments" (see page 356). Out of the 75 submitted full papers, 38 were selected for presentation at the conference. Of the 9 submitted poster abstracts, 6 were chosen for the poster session. The ParSim session received 11 submissions, of which 5 were selected for this special session. The task of reviewing was carried out smoothly within very strict time limits by a large program committee and a number of external referees, counting members from most of the American and European groups involved in MPI and PVM development, as well as from significant user communities. Almost all papers received 4 reviews, some even 5, and none fewer than 3, which provided a solid basis for the program chairs to make the final selection for the conference program. The result was a well-balanced and focused program of high quality. All authors are thanked for their contribution to the conference. Out of the accepted 38 papers, 3 were selected as *outstanding contributions* to EuroPVM/MPI 2006, and were presented in a special, plenary session:

- "Issues in Developing a Thread-Safe MPI Implementation" by William Gropp and Rajeev Thakur (page 12)
- "Scalable Parallel Suffix Array Construction" by Fabian Kulla and Peter Sanders (page 22)

- “Formal Verification of Programs That Use MPI One-Sided Communication” by Salman Pervez, Ganesh Gopalakrishnan, Robert M. Kirby, Rajeev Thakur and William Gropp (page 30)

“Late and breaking results”, which were submitted in August as brief abstracts and therefore not included in these proceedings, were presented in the eponymous session. Like the “Outstanding Papers” session, this was a premiere at EuroPVM/MPI 2006.

Complementing the emphasis in the call for papers on new message-passing paradigms and programming models, the invited talks by Richard Graham, William Gropp and Al Geist addressed possible shortcomings of MPI for emerging, large-scale systems, covering issues on fault-tolerance and heterogeneity, productivity and scalability, while the invited talk of Katherine Yelick dealt with advantages of higher-level, partitioned global address space languages. The invited talk of Vaidy Sunderam discussed challenges to message-passing programming in dynamic metacomputing environments. Finally, with the invited talk of Ryutaro Himeno, the audience gained insight into the role and design of the projected Japanese peta-scale supercomputer.

An important part of EuroPVM/MPI is the technically oriented vendor session. At EuroPVM/MPI 2006 eight significant vendors of hard- and software for high-performance computing (Etnus, IBM, Intel, NEC, Dolphin Interconnect Solutions, Hewlett-Packard, Microsoft, and Sun), presented their latest products and developments.

Prior to the conference proper, four tutorials on various aspects of message passing programming (“Using MPI-2: A Problem-Based Approach”, “Performance Tools for Parallel Programming”, “High-Performance Parallel I/O”, and “Hybrid MPI and OpenMP Parallel Programming”) were given by experts in the respective fields.

Information about the conference can be found at the conference Web-site <http://www.pvmmi06.org>, which will be kept available.

The proceedings were edited by Bernd Mohr, Jesper Larsson Träff and Joachim Worringen. The EuroPVM/MPI 2006 logo was designed by Bernd Mohr and Joachim Worringen.

The program and general chairs would like to thank all who contributed to making EuroPVM/MPI 2006 a fruitful and stimulating meeting, be they technical paper or poster authors, program committee members, external referees, participants, or sponsors.

September 2006



Bernd Mohr
Jesper Larsson Träff
Joachim Worringen
Jack Dongarra

Organization

General Chair

Jack Dongarra

University of Tennessee, USA

Program Chairs

Bernd Mohr
Jesper Larsson Träff
Joachim Worringen

Forschungszentrum Jülich, Germany
C&C Research Labs, NEC Europe, Germany
C&C Research Labs, NEC Europe, Germany

Program Committee

George Almasi	IBM, USA
Ranieri Baraglia	CNUCE Institute, Italy
Richard Barrett	ORNL, USA
Gil Bloch	Mellanox, Israel
Arndt Bode	Technical University of Munich, Germany
Marian Bubak	AGH Cracow, Poland
Hakon Bugge	Scali, Norway
Franck Cappello	Université de Paris-Sud, France
Barbara Chapman	University of Houston, USA
Brian Coghlann	Trinity College Dublin, Ireland
Yiannis Cotronis	University of Athens, Greece
Jose Cunha	New University of Lisbon, Portugal
Marco Danelutto	University of Pisa, Italy
Frank Dehne	Carleton University, Canada
Luiz DeRose	Cray, USA
Frederic Despres	INRIA, France
Erik D'Hollander	University of Ghent, Belgium
Beniamino Di Martino	Second University of Naples, Italy
Jack Dongarra	University of Tennessee, USA
Graham Fagg	University of Tennessee, USA
Edgar Gabriel	University of Houston, USA
Al Geist	OakRidge National Laboratory, USA
Patrick Geoffray	Myricom, USA
Michael Gerndt	Tu München, Germany
Andrzej Goscinski	Deakin University, Australia
Richard L. Graham	LANL, USA
William D. Gropp	Argonne National Laboratory, USA
Erez Haba	Microsoft, USA

Program Committee (cont'd)

Rolf Hempel	DLR - German Aerospace Center, Germany
Dieter Kranzlmüller	Johannes Kepler Universität Linz, Austria
Rainer Keller	HLRS, Germany
Stefan Lankes	RWTH Aachen, Germany
Erwin Laure	CERN, Switzerland
Laurent Lefevre	INRIA/LIP, France
Greg Lindahl	QLogic, USA
Thomas Ludwig	University of Heidelberg, Germany
Emilio Luque	Universitat Autònoma de Barcelona, Spain
Ewing Rusty Lusk	Argonne National Laboratory, USA
Tomas Margalef	Universitat Autònoma de Barcelona, Spain
Bart Miller	University of Wisconsin, USA
Bernd Mohr	Forschungszentrum Jülich, Germany
Matthias Müller	Dresden University of Technology, Germany
Salvatore Orlando	University of Venice, Italy
Fabrizio Petrini	PNNL, USA
Neil Pundit	Sandia National Laboratories, USA
Rolf Rabenseifner	HLRS, Germany
Thomas Rauber	Universität Bayreuth, Germany
Wolfgang Rehm	TU Chemnitz, Germany
Casiano Rodriguez-Leon	Universidad de La Laguna, Spain
Michiel Ronsse	University of Ghent, Belgium
Peter Sanders	Universität Karlsruhe, Germany
Martin Schulz	Lawrence Livermore National Laboratory, USA
Jeffrey Squyres	Open System Lab, Indiana
Vaidy Sunderam	Emory University, USA
Bernard Tourancheau	Université de Lyon/INRIA, France
Jesper Larsson Träff	C&C Research Labs, NEC Europe, Germany
Carsten Trinitis	TU München, Germany
Jerzy Wasniewski	Danish Technical University, Denmark
Roland Wismueller	University of Siegen, Germany
Felix Wolf	Forschungszentrum Jülich, Germany
Joachim Worringen	C&C Research Labs, NEC Europe, Germany
Laurence T. Yang	St. Francis Xavier University, Canada

External Referees

(excluding members of the Program Committee)

Dorian Arnold	Ron Brightwell	Rafael Corchuelo
Christian Bell	Michael Brim	Karen Devine
Boris Bierbaum	Carsten Clausen	Frank Dopatka

Gábor Dózsa	Frederic Loulergue	vera
Renato Ferrini	Ricardo Peña Marí	Nathan Rosenblum
Rainer Finocchiaro	Torsten Mehlan	John Ryan
Igor Grobman	Frank Mietke	Carsten Scholtes
Yuri Gurevich	Alexander Mirgorodskiy	Silke Schuch
Torsten Höfler	Francesco Moscato	Stephen F. Siegel
Andreas Hoffmann	Zsolt Nemeth	Nicola Tonellootto
Ralf Hoffmann	Raik Nagel	Gara Miranda Valladares
Sascha Hunold	Raffaele Perego	Salvatore Venticinque
Mauro Iacono	Laura Ricci	John Walsh
Adrian Kacso	Rolf Riesen	Zhaofang Wen
Matthew Legendre	Francisco Fernández Ri-	

For the ParSim session the following external referees provided reviews.

Georg Acher	Michael Ott	Max Walter
Tobias Klug	Daniel Stodden	Josef Weidendorfer

Conference Organization

Bernd Mohr
Jesper Larsson Träff
Joachim Worringen

Sponsors

The conference would have been substantially more expensive and much less pleasant to organize without the generous support of a good many industrial sponsors. Platinum and Gold level sponsors also gave talks at the vendor session on their latest products in parallel systems and message passing software. EuroPVM/MPI 2006 gratefully acknowledges the contributions of the sponsors to a successful conference.

Platinum Level Sponsors

Etnus, IBM, Intel, and NEC.



Gold Level Sponsors

Dolphin Interconnect Solutions, Hewlett-Packard, Microsoft, and Sun.



Microsoft



Standard Level Sponsor

QLogic.



Lecture Notes in Computer Science

For information about Vols. 1–4065

please contact your bookseller or Springer

Vol. 4192: B. Mohr, J.L. Träff, J. Worringen, J. Dongarra (Eds.), Recent Advances in Parallel Virtual Machine and Message Passing Interface. XVI, 414 pages. 2006.

Vol. 4185: R. Mizoguchi, Z. Shi, F. Giunchiglia (Eds.), The Semantic Web – ASWC 2006. XX, 778 pages. 2006.

Vol. 4184: M. Bravetti, M. Nuñes, G. Zavattaro (Eds.), Web Services and Formal Methods. X, 289 pages. 2006.

Vol. 4180: M. Kohlhase, OMDoc – An Open Markup Format for Mathematical Documents [version 1.2]. XIX, 428 pages. 2006. (Sublibrary LNAI).

Vol. 4178: A. Corradini, H. Ehrig, U. Montanari, L. Ribeiro, G. Rozenberg (Eds.), Graph Transformations. XII, 473 pages. 2006.

Vol. 4176: S.K. Katsikas, J. Lopez, M. Backes, S. Gritzalis, B. Preneel (Eds.), Information Security. XIV, 548 pages. 2006.

Vol. 4168: Y. Azar, T. Erlebach (Eds.), Algorithms – ESA 2006. XVIII, 843 pages. 2006.

Vol. 4165: W. Jonker, M. Petkovic (Eds.), Secure, Data Management. X, 185 pages. 2006.

Vol. 4163: H. Bersini, J. Carneiro (Eds.), Artificial Immune Systems. XII, 460 pages. 2006.

Vol. 4162: R. Královič, P. Urzyczyn (Eds.), Mathematical Foundations of Computer Science 2006. XV, 814 pages. 2006.

Vol. 4159: J. Ma, H. Jin, L.T. Yang, J.J.-P. Tsai (Eds.), Ubiquitous Intelligence and Computing. XXII, 1190 pages. 2006.

Vol. 4158: L.T. Yang, H. Jin, J. Ma, T. Ungerer (Eds.), Autonomic and Trusted Computing. XIV, 613 pages. 2006.

Vol. 4156: S. Amer-Yahia, Z. Bellahsène, E. Hunt, R. Ulland, J.X. Yu (Eds.), Database and XML Technologies. IX, 123 pages. 2006.

Vol. 4155: O. Stock, M. Schaerf (Eds.), Reasoning, Action and Interaction in AI Theories and Systems. XVIII, 343 pages. 2006. (Sublibrary LNAI).

Vol. 4153: N. Zheng, X. Jiang, X. Lan (Eds.), Advances in Machine Vision, Image Processing, and Pattern Analysis. XIII, 506 pages. 2006.

Vol. 4152: Y. Manolopoulos, J. Pokorný, T. Sellis (Eds.), Advances in Databases and Information Systems. XV, 448 pages. 2006.

Vol. 4151: A. Iglesias, N. Takayama (Eds.), Mathematical Software - ICMS 2006. XVII, 452 pages. 2006.

Vol. 4150: M. Dorigo, L.M. Gambardella, M. Birattari, A. Martinoli, R. Poli, T. Stützle (Eds.), Ant Colony Optimization and Swarm Intelligence. XVI, 526 pages. 2006.

Vol. 4149: M. Klusch, M. Rovatsos, T.R. Payne (Eds.), Cooperative Information Agents X. XII, 477 pages. 2006. (Sublibrary LNAI).

Vol. 4146: J.C. Rajapakse, L. Wong, R. Acharya (Eds.), Pattern Recognition in Bioinformatics. XIV, 186 pages. 2006. (Sublibrary LNBI).

Vol. 4144: T. Ball, R.B. Jones (Eds.), Computer Aided Verification. XV, 564 pages. 2006.

Vol. 4139: T. Salakoski, F. Ginter, S. Pyysalo, T. Pähikkala, Advances in Natural Language Processing. XVI, 771 pages. 2006. (Sublibrary LNAI).

Vol. 4138: X. Cheng, W. Li, T. Znati (Eds.), Wireless Algorithms, Systems, and Applications. XVI, 709 pages. 2006.

Vol. 4137: C. Baier, H. Hermanns (Eds.), CONCUR 2006 – Concurrency Theory. XIII, 525 pages. 2006.

Vol. 4136: R.A. Schmidt (Ed.), Relations and Kleene Algebra in Computer Science. XI, 433 pages. 2006.

Vol. 4135: C.S. Calude, M.J. Dinneen, G. Păun, G. Rozenberg, S. Stepney (Eds.), Unconventional Computation. X, 267 pages. 2006.

Vol. 4134: K. Yi (Ed.), Static Analysis. XIII, 443 pages. 2006.

Vol. 4133: J. Gratch, M. Young, R. Aylett, D. Ballin, P. Olivier (Eds.), Intelligent Virtual Agents. XIV, 472 pages. 2006. (Sublibrary LNAI).

Vol. 4130: U. Furbach, N. Shankar (Eds.), Automated Reasoning. XV, 680 pages. 2006. (Sublibrary LNAI).

Vol. 4129: D. McGookin, S. Brewster (Eds.), Haptic and Audio Interaction Design. XII, 167 pages. 2006.

Vol. 4128: W.E. Nagel, W.V. Walter, W. Lehner (Eds.), Euro-Par 2006 Parallel Processing. XXXIII, 1221 pages. 2006.

Vol. 4127: E. Damiani, P. Liu (Eds.), Data and Applications Security XX. X, 319 pages. 2006.

Vol. 4126: P. Barahona, F. Bry, E. Franconi, N. Henze, U. Sattler, Reasoning Web. X, 269 pages. 2006.

Vol. 4124: H. de Meer, J.P. G. Sterbenz (Eds.), Self-Organizing Systems. XIV, 261 pages. 2006.

Vol. 4121: A. Biere, C.P. Gomes (Eds.), Theory and Applications of Satisfiability Testing - SAT 2006. XII, 438 pages. 2006.

Vol. 4119: C. Dony, J.L. Knudsen, A. Romanovsky, A. Tripathi (Eds.), Advanced Topics in Exception Handling Components. X, 302 pages. 2006.

Vol. 4117: C. Dwork (Ed.), Advances in Cryptology - CRYPTO 2006. XIII, 621 pages. 2006.

- Vol. 4116: R. De Prisco, M. Yung (Eds.), Security and Cryptography for Networks. XI, 366 pages. 2006.
- Vol. 4115: D.-S. Huang, K. Li, G.W. Irwin (Eds.), Computational Intelligence and Bioinformatics, Part III. XXI, 803 pages. 2006. (Sublibrary LNBI).
- Vol. 4114: D.-S. Huang, K. Li, G.W. Irwin (Eds.), Computational Intelligence, Part II. XXVII, 1337 pages. 2006. (Sublibrary LNAI).
- Vol. 4113: D.-S. Huang, K. Li, G.W. Irwin (Eds.), Intelligent Computing, Part I. XXVII, 1331 pages. 2006.
- Vol. 4112: D.Z. Chen, D. T. Lee (Eds.), Computing and Combinatorics. XIV, 528 pages. 2006.
- Vol. 4111: F.S. de Boer, M.M. Bonsangue, S. Graf, W.-P. de Roever (Eds.), Formal Methods for Components and Objects. VIII, 447 pages. 2006.
- Vol. 4110: J. Díaz, K. Jansen, J.D.P. Rolim, U. Zwick (Eds.), Approximation, Randomization, and Combinatorial Optimization. XII, 522 pages. 2006.
- Vol. 4109: D.-Y. Yeung, J.T. Kwok, A. Fred, F. Roli, D. de Ridder (Eds.), Structural, Syntactic, and Statistical Pattern Recognition. XXI, 939 pages. 2006.
- Vol. 4108: J.M. Borwein, W.M. Farmer (Eds.), Mathematical Knowledge Management. VIII, 295 pages. 2006. (Sublibrary LNAI).
- Vol. 4106: T.R. Roth-Berghofer, M.H. Göker, H. A. Güvenir (Eds.), Advances in Case-Based Reasoning. XIV, 566 pages. 2006. (Sublibrary LNAI).
- Vol. 4104: T. Kunz, S.S. Ravi (Eds.), Ad-Hoc, Mobile, and Wireless Networks. XII, 474 pages. 2006.
- Vol. 4102: S. Dustdar, J.L. Fiadeiro, A. Sheth (Eds.), Business Process Management. XV, 486 pages. 2006.
- Vol. 4099: Q. Yang, G. Webb (Eds.), PRICAI 2006: Trends in Artificial Intelligence. XXVIII, 1263 pages. 2006. (Sublibrary LNAI).
- Vol. 4098: F. Pfenning (Ed.), Term Rewriting and Applications. XIII, 415 pages. 2006.
- Vol. 4097: X. Zhou, O. Sokolsky, L. Yan, E.-S. Jung, Z. Shao, Y. Mu, D.C. Lee, D. Kim, Y.-S. Jeong, C.-Z. Xu (Eds.), Emerging Directions in Embedded and Ubiquitous Computing. XXVII, 1034 pages. 2006.
- Vol. 4096: E. Sha, S.-K. Han, C.-Z. Xu, M.H. Kim, L.T. Yang, B. Xiao (Eds.), Embedded and Ubiquitous Computing. XXIV, 1170 pages. 2006.
- Vol. 4095: S. Nolfi, G. Baldassare, R. Calabretta, D. Marocco, D. Parisi, J.C. T. Hallam, O. Migliano, J.-A. Meyer (Eds.), From Animals to Animats 9. XV, 869 pages. 2006. (Sublibrary LNAI).
- Vol. 4094: O. H. Ibarra, H.-C. Yen (Eds.), Implementation and Application of Automata. XIII, 291 pages. 2006.
- Vol. 4093: X. Li, O.R. Zaïane, Z. Li (Eds.), Advanced Data Mining and Applications. XXI, 1110 pages. 2006. (Sublibrary LNAI).
- Vol. 4092: J. Lang, F. Lin, J. Wang (Eds.), Knowledge Science, Engineering and Management. XV, 664 pages. 2006. (Sublibrary LNAI).
- Vol. 4091: G.-Z. Yang, T. Jiang, D. Shen, L. Gu, J. Yang (Eds.), Medical Imaging and Augmented Reality. XIII, 399 pages. 2006.
- Vol. 4090: S. Spaccapietra, K. Aberer, P. Cudré-Mauroux (Eds.), Journal on Data Semantics VI. XI, 211 pages. 2006.
- Vol. 4089: W. Löwe, M. Südholt (Eds.), Software Composition. X, 339 pages. 2006.
- Vol. 4088: Z.-Z. Shi, R. Sadananda (Eds.), Agent Computing and Multi-Agent Systems. XVII, 827 pages. 2006. (Sublibrary LNAI).
- Vol. 4087: F. Schwenker, S. Marinai (Eds.), Artificial Neural Networks in Pattern Recognition. IX, 299 pages. 2006. (Sublibrary LNAI).
- Vol. 4085: J. Misra, T. Nipkow, E. Sekerinski (Eds.), FM 2006: Formal Methods. XV, 620 pages. 2006.
- Vol. 4084: M.A. Wimmer, H.J. Scholl, Å. Grönlund, K.V. Andersen (Eds.), Electronic Government. XV, 353 pages. 2006.
- Vol. 4083: S. Fischer-Hübner, S. Furnell, C. Lambrounidakis (Eds.), Trust and Privacy in Digital Business. XIII, 243 pages. 2006.
- Vol. 4082: K. Bauknecht, B. Pröll, H. Werthner (Eds.), E-Commerce and Web Technologies. XIII, 243 pages. 2006.
- Vol. 4081: A. M. Tjoa, J. Trujillo (Eds.), Data Warehousing and Knowledge Discovery. XVII, 578 pages. 2006.
- Vol. 4080: S. Bressan, J. Küng, R. Wagner (Eds.), Database and Expert Systems Applications. XXI, 959 pages. 2006.
- Vol. 4079: S. Etalle, M. Truszczyński (Eds.), Logic Programming. XIV, 474 pages. 2006.
- Vol. 4077: M.-S. Kim, K. Shimada (Eds.), Geometric Modeling and Processing - GMP 2006. XVI, 696 pages. 2006.
- Vol. 4076: F. Hess, S. Pauli, M. Pohst (Eds.), Algorithmic Number Theory. X, 599 pages. 2006.
- Vol. 4075: U. Leser, F. Naumann, B. Eckman (Eds.), Data Integration in the Life Sciences. XI, 298 pages. 2006. (Sublibrary LNBI).
- Vol. 4074: M. Burmester, A. Yasinsac (Eds.), Secure Mobile Ad-hoc Networks and Sensors. X, 193 pages. 2006.
- Vol. 4073: A. Butz, B. Fisher, A. Krüger, P. Olivier (Eds.), Smart Graphics. XI, 263 pages. 2006.
- Vol. 4072: M. Harders, G. Székely (Eds.), Biomedical Simulation. XI, 216 pages. 2006.
- Vol. 4071: H. Sundaram, M. Naphade, J.R. Smith, Y. Rui (Eds.), Image and Video Retrieval. XII, 547 pages. 2006.
- Vol. 4070: C. Priami, X. Hu, Y. Pan, T.Y. Lin (Eds.), Transactions on Computational Systems Biology V. IX, 129 pages. 2006. (Sublibrary LNBI).
- Vol. 4069: F.J. Perales, R.B. Fisher (Eds.), Articulated Motion and Deformable Objects. XV, 526 pages. 2006.
- Vol. 4068: H. Schärfe, P. Hitzler, P. Øhrstrøm (Eds.), Conceptual Structures: Inspiration and Application. XI, 455 pages. 2006. (Sublibrary LNAI).
- Vol. 4067: D. Thomas (Ed.), ECOOP 2006 – Object-Oriented Programming. XIV, 527 pages. 2006.
- Vol. 4066: A. Rensink, J. Warmer (Eds.), Model Driven Architecture – Foundations and Applications. XII, 392 pages. 2006.

Table of Contents

Invited Talks

Too Big for MPI?	1
<i>Al Geist</i>	
Approaches for Parallel Applications Fault Tolerance	2
<i>Richard L. Graham</i>	
Where Does MPI Need to Grow?	3
<i>William D. Gropp</i>	
Peta-Scale Supercomputer Project in Japan and Challenges to Life and Human Simulation in Japan	4
<i>Ryutaro Himeno</i>	
Resource and Application Adaptivity in Message Passing Systems	5
<i>Vaidy Sunderam</i>	
Performance Advantages of Partitioned Global Address Space Languages	6
<i>Katherine Yelick</i>	

Tutorials

Using MPI-2: A Problem-Based Approach.....	7
<i>William D. Gropp, Ewing Lusk</i>	
Performance Tools for Parallel Programming	8
<i>Bernd Mohr, Felix Wolf</i>	
High-Performance Parallel I/O	10
<i>Robert Ross, Joachim Worringen</i>	
Hybrid MPI and OpenMP Parallel Programming	11
<i>Rolf Rabenseifner, Georg Hager, Gabriele Jost, Rainer Keller</i>	

Outstanding Papers

Issues in Developing a Thread-Safe MPI Implementation	12
<i>William Gropp, Rajeev Thakur</i>	

Scalable Parallel Suffix Array Construction	22
<i>Fabian Kulla, Peter Sanders</i>	

Formal Verification of Programs That Use MPI One-Sided Communication	30
---	----

*Salman Pervez, Ganesh Gopalakrishnan, Robert M. Kirby,
Rajeev Thakur, William Gropp*

Collective Communication

MPI Collective Algorithm Selection and Quadtree Encoding	40
--	----

*Jelena Pješivac-Grbović, Graham E. Fagg, Thara Angskun,
George Bosilca, Jack J. Dongarra*

Parallel Prefix (Scan) Algorithms for MPI	49
---	----

Peter Sanders, Jesper Larsson Träff

Efficient Allgather for Regular SMP-Clusters	58
--	----

Jesper Larsson Träff

Efficient Shared Memory and RDMA Based Design for MPI_Allgather over InfiniBand	66
--	----

Amith Ranjith Mamidala, Abhinav Vishnu, Dhabaleswar K. Panda

Communication Protocols

High Performance RDMA Protocols in HPC	76
--	----

*Tim S. Woodall, Galen Mark Shipman, George Bosilca,
Richard L. Graham, Arthur B. Maccabe*

Implementation and Shared-Memory Evaluation of MPICH2 over the Nemesis Communication Subsystem	86
---	----

Darius Buntinas, Guillaume Mercier, William Gropp

MPI/CTP: A Reconfigurable MPI for HPC Applications	96
--	----

Manjunath Gorentla Venkata, Patrick G. Bridges

Debugging and Verification

Correctness Checking of MPI One-Sided Communication Using Marmot	105
---	-----

Bettina Krammer, Michael M. Resch

- An Interface to Support the Identification of Dynamic MPI 2 Processes
for Scalable Parallel Debugging 115

*Christopher Gottbrath, Brian Barrett, Bill Gropp,
Ewing "Rusty" Lusk, Jeff Squyres*

- Modeling and Verification of MPI Based Distributed Software 123

Igor Grudenic, Nikola Bogunovic

Fault Tolerance

- FT-MPI, Fault-Tolerant Metacomputing and Generic Name Services:

- A Case Study 133

*David Dewolfs, Jan Broeckhove, Vaidy Sunderam,
Graham E. Fagg*

- Scalable Fault Tolerant Protocol for Parallel Runtime

- Environments 141

*Thara Angskun, Graham E. Fagg, George Bosilca,
Jelena Pešivac-Grbović, Jack J. Dongarra*

- An Intelligent Management of Fault Tolerance in Cluster Using

- RADICMPI 150

Angelo A. Duarte, Dolores Rexachs, Emilio Luque

- Extended mpiJava for Distributed Checkpointing and Recovery 158

Emilio Hernández, Judith Cardinale, Wilmer Pereira

Metacomputing and Grid

- Running PVM Applications on Multidomain Clusters 166

Franco Frattolillo

- Reliable Orchestration of Distributed MPI-Applications in a

- UNICORE-Based Grid with MetaMPICH and MetaScheduling 174

*Boris Bierbaum, Carsten Clauss, Thomas Eickermann,
Lidia Kirtchakova, Arnold Krechel, Stephan Springstubb,
Oliver Wälrich, Wolfgang Ziegler*

- The New Multidevice Architecture of MetaMPICH in the Context of

- Other Approaches to Grid-Enabled MPI 184

*Boris Bierbaum, Carsten Clauss, Martin Pöpke, Stefan Lankes,
Thomas Bemmerl*

Using an Enterprise Grid for Execution of MPI Parallel Applications – A Case Study	194
<i>Adam K.L. Wong, Andrzej M. Goscinski</i>	

Parallel I/O

Self-adaptive Hints for Collective I/O	202
<i>Joachim Worringen</i>	

Exploiting Shared Memory to Improve Parallel I/O Performance	212
<i>Andrew B. Hastings, Alok Choudhary</i>	

High-Bandwidth Remote Parallel I/O with the Distributed Memory Filesystem MEMFS	222
---	-----

*Jan Seidel, Rudolf Berrendorf, Marcel Birkner,
Marc-André Hermanns*

Effective Seamless Remote MPI-I/O Operations with Derived Data Types Using PVFS2	230
<i>Yuichi Tsujita</i>	

Implementation Issues

Automatic Memory Optimizations for Improving MPI Derived Datatype Performance	238
<i>Surendra Byna, Xian-He Sun, Rajeev Thakur, William Gropp</i>	

Improving the Dynamic Creation of Processes in MPI-2	247
<i>Márcia C. Cera, Guilherme P. Pezzi, Elton N. Mathias, Nicolas Maillard, Philippe O.A. Navaux</i>	

Object-Oriented Message Passing

Non-blocking Java Communications Support on Clusters	256
<i>Guillermo L. Taboada, Juan Touriño, Ramón Doallo</i>	

Modernizing the C++ Interface to MPI	266
<i>Prabhanjan Kambadur, Douglas Gregor, Andrew Lumsdaine, Amey Dharurkar</i>	

Limitations and Extensions

Can MPI Be Used for Persistent Parallel Services?	275
<i>Robert Latham, Robert Ross, Rajeev Thakur</i>	

Observations on MPI-2 Support for Hybrid Master/Slave Applications in Dynamic and Heterogeneous Environments	285
<i>Claudia Leopold, Michael Süß</i>	
What MPI Could (and Cannot) Do for Mesh-Partitioning on Non-homogeneous Networks	293
<i>Guntram Berti, Jesper Larsson Träff</i>	

Performance

Scalable Parallel Trace-Based Performance Analysis	303
<i>Markus Geimer, Felix Wolf, Brian J.N. Wylie, Bernd Mohr</i>	
TAUG: Runtime Global Performance Data Access Using MPI	313
<i>Kevin A. Huck, Allen D. Malony, Sameer Shende, Alan Morris</i>	
Tracing the MPI-IO Calls' Disk Accesses	322
<i>Thomas Ludwig, Stephan Krempel, Julian Kunkel, Frank Panse, Dulip Withanage</i>	
Measuring MPI Send and Receive Overhead and Application Availability in High Performance Network Interfaces	331
<i>Douglas Doerfler, Ron Brightwell</i>	
Challenges and Issues in Benchmarking MPI	339
<i>Keith D. Underwood</i>	

Implementation and Usage of the PERUSE-Interface in Open MPI	347
<i>Rainer Keller, George Bosilca, Graham Fagg, Michael Resch, Jack J. Dongarra</i>	

ParSim

Current Trends in Numerical Simulation for Parallel Engineering Environments	356
<i>Martin Schulz, Carsten Trinitis</i>	
MPJ Express Meets Gadget: Towards a Java Code for Cosmological Simulations	358
<i>Mark Baker, Bryan Carpenter, Aamir Shafi</i>	
An Approach for Parallel Fluid-Structure Interaction on Unstructured Meshes	366
<i>Ulrich Küttler, Wolfgang A. Wall</i>	