

Dieter Kranzlmüller
Peter Kacsuk
Jack Dongarra (Eds.)

Recent Advances in Parallel Virtual Machine and Message Passing Interface

11th European PVM/MPI Users' Group Meeting
Budapest, Hungary, September 2004
Proceedings




Springer

Dieter Kranzlmüller Peter Kacsuk
Jack Dongarra (Eds.)

Recent Advances in Parallel Virtual Machine and Message Passing Interface

11th European PVM/MPI Users' Group Meeting
Budapest, Hungary September 14–16, 2004
Proceedings



Volume Editors

Dieter Kranzlmüller
Joh. Kepler Universität Linz
Institut für Graphische und Parallele Datenverarbeitung (GUP)
Altenbergerstr. 69, 4040 Linz, Austria
E-mail: kranzlmueler@gup.jku.at

Peter Kacsuk
Hungarian Academy of Sciences
Computer and Automation Research Institute, MTA SZTAKI
Victor Hugo u. 18-22, 1132 Budapest, Hungary
E-mail: kacsuk@sztaki.hu

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: 2004111520

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

ISSN 0302-9743
ISBN 3-540-23163-3 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

springeronline.com

© Springer-Verlag Berlin Heidelberg 2004
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Olgun Computergrafik
Printed on acid-free paper SPIN: 11324188 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

New York University, NY, 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

The message passing paradigm is the most frequently used approach to develop high-performance computing applications on parallel and distributed computing architectures. Parallel Virtual Machine (PVM) and Message Passing Interface (MPI) are the two main representatives in this domain.

This volume comprises 50 selected contributions presented at the 11th European PVM/MPI Users' Group Meeting, which was held in Budapest, Hungary, September 19–22, 2004. The conference was organized by the Laboratory of Parallel and Distributed Systems (LPDS) at the Computer and Automation Research Institute of the Hungarian Academy of Sciences (MTA SZTAKI).

The conference was previously held in Venice, Italy (2003), Linz, Austria (2002), Santorini, Greece (2001), Balatonfüred, Hungary (2000), Barcelona, Spain (1999), Liverpool, UK (1998), and Krakow, Poland (1997). The first three conferences were devoted to PVM and were held in Munich, Germany (1996), Lyon, France (1995), and Rome, Italy (1994).

In its eleventh year, this conference is well established as the forum for users and developers of PVM, MPI, and other message passing environments. Interactions between these groups have proved to be very useful for developing new ideas in parallel computing, and for applying some of those already existent to new practical fields. The main topics of the meeting were evaluation and performance of PVM and MPI, extensions, implementations and improvements of PVM and MPI, parallel algorithms using the message passing paradigm, and parallel applications in science and engineering. In addition, the topics of the conference were extended to include cluster and grid computing, in order to reflect the importance of this area for the high-performance computing community.

Besides the main track of contributed papers, the conference featured the third edition of the special session “ParSim 04 – Current Trends in Numerical Simulation for Parallel Engineering Environments”. The conference also included three tutorials, one on “Using MPI-2: A Problem-Based Approach”, one on “Interactive Applications on the Grid – the CrossGrid Tutorial”, and another one on “Production Grid Systems and Their Programming”, and invited talks on MPI and high-productivity programming, fault tolerance in message passing and in action, high-performance application execution scenarios in P-GRADE, an open cluster system software stack, from PVM grids to self-assembling virtual machines, the grid middleware of the NorduGrid, next-generation grids, and the Austrian Grid initiative – high-level extensions to grid middleware. These proceedings contain papers on the 50 contributed presentations together with abstracts of the invited and tutorial speakers' presentations.

The 11th Euro PVM/MPI conference was held together with DAPSYS 2004, the 5th Austrian-Hungarian Workshop on Distributed and Parallel Systems. Participants of the two events shared invited talks, tutorials, the vendors' session, and social events, while contributed paper presentations proceeded in separate

tracks in parallel. While Euro PVM/MPI is dedicated to the latest developments of PVM and MPI, DAPSYS was a major event to discuss general aspects of distributed and parallel systems. In this way the two events were complementary to each other and participants of Euro PVM/MPI could benefit from the joint organization of the two events.

The invited speakers of the joint Euro PVM/MPI and DAPSYS conference were Jack Dongarra, Gabor Dozsa, Al Geist, William Gropp, Balazs Konya, Domenico Laforenza, Ewing Lusk, and Jens Volkert. The tutorials were presented by William Gropp and Ewing Lusk, Tomasz Szepieniec, Marcin Radecki and Katarzyna Rycerz, and Peter Kacsuk, Balazs Konya, and Peter Stefan.

We express our gratitude for the kind support of our sponsors (see below) and we thank the members of the Program Committee and the additional reviewers for their work in refereeing the submitted papers and ensuring the high quality of Euro PVM/MPI. Finally, we would like to express our gratitude to our colleagues at MTA SZTAKI and GUP, JKU Linz for their help and support during the conference organization.

September 2004

Dieter Kranzlmüller
Peter Kacsuk
Jack Dongarra

Program Committee

General Chair

Jack Dongarra

University of Tennessee, Knoxville, USA

Program Chairs

Peter Kacsuk
Dieter Kranzlmüller

MTA SZTAKI, Budapest, Hungary
GUP, Joh. Kepler University, Linz, Austria

Program Committee Members

- | | |
|-----------------------------------|--|
| David Abramson | Monash University, Australia |
| Vassil Alexandrov | University of Reading, UK |
| Ranieri Baraglia | CNUCE Institute, Italy |
| Arndt Bode | Technical Univ. of Munich, Germany |
| Marian Bubak | AGH, Cracow, Poland |
| Jacques Chassin
de Kergommeaux | LSR-IMAG, France |
| Yiannis Cotronis | Univ. of Athens, Greece |
| Jose C. Cunha | New University of Lisbon, Portugal |
| Marco Danelutto | Univ. of Pisa, Italy |
| Frederic Desprez | INRIA, France |
| Erik D'Hollander | University of Ghent, Belgium |
| Beniamino Di Martino | UNINA, Italy |
| Jack Dongarra | University of Tennessee, Knoxville, USA |
| Graham Fagg | University of Tennessee, Knoxville, USA |
| Thomas Fahringer | University of Innsbruck, Austria |
| Al Geist | Oak Ridge National Laboratory, USA |
| Michael Gerndt | Technical Univ. of Munich, Germany |
| Andrzej Goscinski | Deakin University, Australia |
| William Gropp | Argonne National Laboratory, USA |
| Rolf Hempel | DLR, Simulation Aerospace Center,
Germany |
| Ladislav Hluchy | Slovak Academy of Sciences, Slovakia |
| Peter Kacsuk | MTA SZTAKI, Hungary |
| Dieter Kranzlmüller | Joh. Kepler University, Linz, Austria |
| Jan Kwiatkowski | Wroclaw University of Technology, Poland |
| Domenico Laforenza | CNUCE, Italy |
| Laurent Lefevre | INRIA, France |
| Thomas Ludwig | University of Heidelberg, Germany |
| Emilio Luque | Universitat Autonoma de Barcelona, Spain |

VIII Program Committee

Ewing Lusk	Argonne National Laboratory, USA
Tomas Margalef	Universitat Autonoma de Barcelona, Spain
Barton Miller	University of Wisconsin, USA
Shirley Moore	University of Tennessee, USA
Wolfgang Nagel	Dresden University of Technology, Germany
Salvatore Orlando	University of Venice, Italy
Benno J. Overeinder	University of Amsterdam, The Netherlands
Raffaele Perego	ISTI, Italy
Neil D. Pundit	Sandia National Labs, USA
Rolf Rabenseifner	University of Stuttgart, Germany
Andrew Rau-Chaplin	Dalhousie University, Canada
Jeff Reeve	University of Southampton, UK
Ralf Reussner	University of Oldenburg, Germany
Yves Robert	ENS Lyon, France
Casiano Rodriguez-Leon	Universidad de La Laguna, Spain
Michiel Ronsse	University of Ghent, Belgium
Wolfgang Schreiner	Joh. Kepler University, Linz, Austria
Miquel Senar	Universitat Autonoma de Barcelona, Spain
Joao Gabriel Silva	University of Coimbra, Portugal
Vaidy Sunderam	Emroy University, USA
Francisco Tirado	Universidad Complutense, Spain
Bernard Tourancheau	SUN Labs Europe, France
Jesper Larsson Träff	NEC Europe, Germany
Pavel Tvrđík	Czech Technical University, Czech Republic
Umberto Villano	University of Sannio, Italy
Jens Volkert	Joh. Kepler University, Linz, Austria
Jerzy Wasniewski	Technical University Denmark
Roland Wismüller	Technical Univ. of Munich, Germany

Sponsoring Institutions

(as of August 3, 2004)

IBM

Intel

NEC

EGEE project

Hungarian Ministry of Education, OM

Hungarian Academy of Sciences, MTA

Foundation for the Technological Progress of the Industry, IMFA (Hungary)

Lecture Notes in Computer Science

For information about Vols. 1–3114

please contact your bookseller or Springer

- Vol. 3241: D. Kranzlmüller, P. Kacsuk, J. Dongarra (Eds.), Recent Advances in Parallel Virtual Machine and Message Passing Interface. XIII, 452 pages. 2004.
- Vol. 3240: I. Jonassen, J. Kim (Eds.), Algorithms in Bioinformatics. IX, 476 pages. 2004. (Subseries LNBI).
- Vol. 3239: G. Nicosia, V. Cutello, P.J. Bentley, J. Timmis (Eds.), Artificial Immune Systems. XII, 444 pages. 2004.
- Vol. 3232: R. Heery, L. Lyon (Eds.), Research and Advanced Technology for Digital Libraries. XV, 528 pages. 2004.
- Vol. 3224: E. Jonsson, A. Valdes, M. Almgren (Eds.), Recent Advances in Intrusion Detection. XII, 315 pages. 2004.
- Vol. 3223: K. Slind, A. Bunker, G. Gopalakrishnan (Eds.), Theorem Proving in Higher Order Logic. VIII, 337 pages. 2004.
- Vol. 3221: S. Albers, T. Radzik (Eds.), Algorithms – ESA 2004. XVIII, 836 pages. 2004.
- Vol. 3220: J.C. Lester, R.M. Vicari, F. Paraguacu (Eds.), Intelligent Tutoring Systems. XXI, 920 pages. 2004.
- Vol. 3210: J. Marcinkowski, A. Tarlecki (Eds.), Computer Science Logic. XI, 520 pages. 2004.
- Vol. 3208: H.J. Ohlbach, S. Schaffert (Eds.), Principles and Practice of Semantic Web Reasoning. VII, 165 pages. 2004.
- Vol. 3207: L.T. Jang, M. Guo, G.R. Gao, N.K. Jha, Embedded and Ubiquitous Computing. XX, 1116 pages. 2004.
- Vol. 3206: P. Sojka, I. Kopecek, K. Pala (Eds.), Text, Speech and Dialogue. XIII, 667 pages. 2004. (Subseries LNAI).
- Vol. 3205: N. Davies, E. Mynatt, I. Siio (Eds.), UbiComp 2004: Ubiquitous Computing. XVI, 452 pages. 2004.
- Vol. 3203: J. Becker, M. Platzner, S. Vernalde (Eds.), Field Programmable Logic and Application. XXX, 1198 pages. 2004.
- Vol. 3202: J.-F. Boulicaut, F. Esposito, F. Giannotti, D. Pedreschi (Eds.), Knowledge Discovery in Databases: PKDD 2004. XIX, 560 pages. 2004. (Subseries LNAI).
- Vol. 3201: J.-F. Boulicaut, F. Esposito, F. Giannotti, D. Pedreschi (Eds.), Machine Learning: ECML 2004. XVIII, 580 pages. 2004. (Subseries LNAI).
- Vol. 3199: H. Scheper (Ed.), Software and Compilers for Embedded Systems. X, 259 pages. 2004.
- Vol. 3198: G.-J. de Vreede, L.A. Guerrero, G. Marín Raventós (Eds.), Groupware: Design, Implementation and Use. XI, 378 pages. 2004.
- Vol. 3194: R. Camacho, R. King, A. Srinivasan (Eds.), Inductive Logic Programming. XI, 361 pages. 2004. (Subseries LNAI).
- Vol. 3193: P. Samarati, P. Ryan, D. Gollmann, R. Molva (Eds.), Computer Security – ESORICS 2004. X, 457 pages. 2004.
- Vol. 3192: C. Bussler, D. Fensel (Eds.), Artificial Intelligence: Methodology, Systems, and Applications. XIII, 522 pages. 2004. (Subseries LNAI).
- Vol. 3190: Y. Luo (Ed.), Cooperative Design, Visualization, and Engineering. IX, 248 pages. 2004.
- Vol. 3189: P.-C. Yew, J. Xue (Eds.), Advances in Computer Systems Architecture. XVII, 598 pages. 2004.
- Vol. 3186: Z. Bellahsène, T. Milo, M. Rys, D. Suciu, R. Unland (Eds.), Database and XML Technologies. X, 235 pages. 2004.
- Vol. 3185: M. Bernardo, F. Corradini (Eds.), Formal Methods for the Design of Real-Time Systems. VII, 295 pages. 2004.
- Vol. 3184: S. Katsikas, J. Lopez, G. Pernul (Eds.), Trust and Privacy in Digital Business. XI, 299 pages. 2004.
- Vol. 3183: R. Traunmüller (Ed.), Electronic Government. XIX, 583 pages. 2004.
- Vol. 3182: K. Bauknecht, M. Bichler, B. Pröll (Eds.), E-Commerce and Web Technologies. XI, 370 pages. 2004.
- Vol. 3181: Y. Kambayashi, M. Mohania, W. Wöß (Eds.), Data Warehousing and Knowledge Discovery. XIV, 412 pages. 2004.
- Vol. 3180: F. Galindo, M. Takizawa, R. Traunmüller (Eds.), Database and Expert Systems Applications. XXI, 972 pages. 2004.
- Vol. 3179: F.J. Perales, B.A. Draper (Eds.), Articulated Motion and Deformable Objects. XI, 270 pages. 2004.
- Vol. 3178: W. Jonker, M. Petkovic (Eds.), Secure Data Management. VIII, 219 pages. 2004.
- Vol. 3177: Z.R. Yang, H. Yin, R. Everson (Eds.), Intelligent Data Engineering and Automated Learning – IDEAL 2004. XVIII, 852 pages. 2004.
- Vol. 3176: O. Bousquet, U. von Luxburg, G. Rätsch (Eds.), Advanced Lectures on Machine Learning. VIII, 241 pages. 2004. (Subseries LNAI).
- Vol. 3175: C.E. Rasmussen, H.H. Bühlhoff, B. Schölkopf, M.A. Giese (Eds.), Pattern Recognition. XVIII, 581 pages. 2004.
- Vol. 3174: F. Yin, J. Wang, C. Guo (Eds.), Advances in Neural Networks - ISNN 2004. XXXV, 1021 pages. 2004.
- Vol. 3173: F. Yin, J. Wang, C. Guo (Eds.), Advances in Neural Networks – ISNN 2004. XXXV, 1041 pages. 2004.
- Vol. 3172: M. Dorigo, M. Birattari, C. Blum, L. M. Gambardella, F. Mondada, T. Stützle (Eds.), Ant Colony Optimization and Swarm Intelligence. XII, 434 pages. 2004.

- Vol. 3170: P. Gardner, N. Yoshida (Eds.), CONCUR 2004 - Concurrency Theory. XIII, 529 pages. 2004.
- Vol. 3166: M. Rauterberg (Ed.), Entertainment Computing - ICEC 2004. XXIII, 617 pages. 2004.
- Vol. 3163: S. Marinai, A. Dengel (Eds.), Document Analysis Systems VI. XII, 564 pages. 2004.
- Vol. 3162: R. Downey, M. Fellows, F. Dehne (Eds.), Parameterized and Exact Computation. X, 293 pages. 2004.
- Vol. 3160: S. Brewster, M. Dunlop (Eds.), Mobile Human-Computer Interaction - MobileHCI 2004. XVIII, 541 pages. 2004.
- Vol. 3159: U. Visser, Intelligent Information Integration for the Semantic Web. XIV, 150 pages. 2004. (Subseries LNAI).
- Vol. 3158: I. Nikolaidis, M. Barbeau, E. Kranakis (Eds.), Ad-Hoc, Mobile, and Wireless Networks. IX, 344 pages. 2004.
- Vol. 3157: C. Zhang, H. W. Guesgen, W.K. Yeap (Eds.), PRICAI 2004: Trends in Artificial Intelligence. XX, 1023 pages. 2004. (Subseries LNAI).
- Vol. 3156: M. Joye, J.-J. Quisquater (Eds.), Cryptographic Hardware and Embedded Systems - CHES 2004. XIII, 455 pages. 2004.
- Vol. 3155: P. Funk, P.A. González Calero (Eds.), Advances in Case-Based Reasoning. XIII, 822 pages. 2004. (Subseries LNAI).
- Vol. 3154: R.L. Nord (Ed.), Software Product Lines. XIV, 334 pages. 2004.
- Vol. 3153: J. Fiala, V. Koubek, J. Kratochvíl (Eds.), Mathematical Foundations of Computer Science 2004. XIV, 902 pages. 2004.
- Vol. 3152: M. Franklin (Ed.), Advances in Cryptology - CRYPTO 2004. XI, 579 pages. 2004.
- Vol. 3150: G.-Z. Yang, T. Jiang (Eds.), Medical Imaging and Augmented Reality. XII, 378 pages. 2004.
- Vol. 3149: M. Danelutto, M. Vanneschi, D. Laforenza (Eds.), Euro-Par 2004 Parallel Processing. XXXIV, 1081 pages. 2004.
- Vol. 3148: R. Giacobazzi (Ed.), Static Analysis. XI, 393 pages. 2004.
- Vol. 3146: P. Érdi, A. Esposito, M. Marinaro, S. Scarpetta (Eds.), Computational Neuroscience: Cortical Dynamics. XI, 161 pages. 2004.
- Vol. 3144: M. Papatriantafilou, P. Hunel (Eds.), Principles of Distributed Systems. XI, 246 pages. 2004.
- Vol. 3143: W. Liu, Y. Shi, Q. Li (Eds.), Advances in Web-Based Learning - ICWL 2004. XIV, 459 pages. 2004.
- Vol. 3142: J. Diaz, J. Karhumäki, A. Lepistö, D. Sannella (Eds.), Automata, Languages and Programming. XIX, 1253 pages. 2004.
- Vol. 3140: N. Koch, P. Fraternali, M. Wirsing (Eds.), Web Engineering. XXI, 623 pages. 2004.
- Vol. 3139: F. Iida, R. Pfeifer, L. Steels, Y. Kuniyoshi (Eds.), Embodied Artificial Intelligence. IX, 331 pages. 2004. (Subseries LNAI).
- Vol. 3138: A. Fred, T. Caelli, R.P.W. Duin, A. Campilho, D.d. Ridder (Eds.), Structural, Syntactic, and Statistical Pattern Recognition. XXII, 1168 pages. 2004.
- Vol. 3137: P. De Bra, W. Nejdl (Eds.), Adaptive Hypermedia and Adaptive Web-Based Systems. XIV, 442 pages. 2004.
- Vol. 3136: F. Meziane, E. Métais (Eds.), Natural Language Processing and Information Systems. XII, 436 pages. 2004.
- Vol. 3134: C. Zannier, H. Erdogmus, L. Lindstrom (Eds.), Extreme Programming and Agile Methods - XP/Agile Universe 2004. XIV, 233 pages. 2004.
- Vol. 3133: A.D. Pimentel, S. Vassiliadis (Eds.), Computer Systems: Architectures, Modeling, and Simulation. XIII, 562 pages. 2004.
- Vol. 3132: B. Demoen, V. Lifschitz (Eds.), Logic Programming. XII, 480 pages. 2004.
- Vol. 3131: V. Torra, Y. Narukawa (Eds.), Modeling Decisions for Artificial Intelligence. XI, 327 pages. 2004. (Subseries LNAI).
- Vol. 3130: A. Syropoulos, K. Berry, Y. Haralambous, B. Hughes, S. Peter, J. Plaice (Eds.), TeX, XML, and Digital Typography. VIII, 265 pages. 2004.
- Vol. 3129: Q. Li, G. Wang, L. Feng (Eds.), Advances in Web-Age Information Management. XVII, 753 pages. 2004.
- Vol. 3128: D. Asanov (Ed.), Querying Databases Privately. IX, 115 pages. 2004.
- Vol. 3127: K.E. Wolff, H.D. Pfeiffer, H.S. Delugach (Eds.), Conceptual Structures at Work. XI, 403 pages. 2004. (Subseries LNAI).
- Vol. 3126: P. Dini, P. Lorenz, J.N.d. Souza (Eds.), Service Assurance with Partial and Intermittent Resources. XI, 312 pages. 2004.
- Vol. 3125: D. Kozen (Ed.), Mathematics of Program Construction. X, 401 pages. 2004.
- Vol. 3124: J.N. de Souza, P. Dini, P. Lorenz (Eds.), Telecommunications and Networking - ICT 2004. XXVI, 1390 pages. 2004.
- Vol. 3123: A. Belz, R. Evans, P. Piwek (Eds.), Natural Language Generation. X, 219 pages. 2004. (Subseries LNAI).
- Vol. 3122: K. Jansen, S. Khanna, J.D.P. Rolim, D. Ron (Eds.), Approximation, Randomization, and Combinatorial Optimization. IX, 428 pages. 2004.
- Vol. 3121: S. Nikoletseas, J.D.P. Rolim (Eds.), Algorithmic Aspects of Wireless Sensor Networks. X, 201 pages. 2004.
- Vol. 3120: J. Shawe-Taylor, Y. Singer (Eds.), Learning Theory. X, 648 pages. 2004. (Subseries LNAI).
- Vol. 3119: A. Aspert, G. Bancerek, A. Trybulec (Eds.), Mathematical Knowledge Management. X, 393 pages. 2004.
- Vol. 3118: K. Miesenberger, J. Klaus, W. Zagler, D. Burger (Eds.), Computer Helping People with Special Needs. XXIII, 1191 pages. 2004.
- Vol. 3116: C. Ratray, S. Maharaj, C. Shankland (Eds.), Algebraic Methodology and Software Technology. XI, 569 pages. 2004.
- Vol. 3115: P. Enser, Y. Kompatsiaris, N.E. O'Connor, A.F. Smeaton, A.W.M. Smeulders (Eds.), Image and Video Retrieval. XVII, 679 pages. 2004.

Table of Contents

Invited Talks

PVM Grids to Self-assembling Virtual Machines	1
<i>A. Geist</i>	
The Austrian Grid Initiative –	
High Level Extensions to Grid Middleware	5
<i>J. Volkert</i>	
Fault Tolerance in Message Passing and in Action	6
<i>J.J. Dongarra</i>	
MPI and High Productivity Programming	7
<i>W.D. Gropp</i>	
High Performance Application Execution Scenarios in P-GRADE	8
<i>G. Dózsa</i>	
An Open Cluster System Software Stack	9
<i>E. Lusk</i>	
Advanced Resource Connector (ARC) –	
The Grid Middleware of the NorduGrid	10
<i>B. Kónya</i>	
Next Generation Grid: Learn from the Past, Look to the Future	11
<i>D. Laforenza</i>	

Tutorials

Production Grid Systems and Their Programming	13
<i>P. Kacsuk, B. Kónya, and P. Stefán</i>	
Tools and Services for Interactive Applications on the Grid –	
The CrossGrid Tutorial	14
<i>T. Szepieniec, M. Radecki, K. Rycerz, M. Bubak, and M. Malawski</i>	

Extensions and Improvements

Verifying Collective MPI Calls	18
<i>J.L. Träff and J. Worringen</i>	
Fast Tuning of Intra-cluster Collective Communications	28
<i>L.A. Barchet-Estefanel and G. Mounié</i>	

More Efficient Reduction Algorithms for Non-Power-of-Two Number of Processors in Message-Passing Parallel Systems	36
<i>R. Rabenseifner and J.L. Träff</i>	
Zero-Copy MPI Derived Datatype Communication over InfiniBand	47
<i>G. Santhanaraman, J. Wu, and D.K. Panda</i>	
Minimizing Synchronization Overhead in the Implementation of MPI One-Sided Communication	57
<i>R. Thakur, W.D. Gropp, and B. Toonen</i>	
Efficient Implementation of MPI-2 Passive One-Sided Communication on InfiniBand Clusters	68
<i>W. Jiang, J. Liu, H.-W. Jin, D.K. Panda, D. Buntinas, R. Thakur, and W.D. Gropp</i>	
Providing Efficient I/O Redundancy in MPI Environments	77
<i>W.D. Gropp, R. Ross, and N. Miller</i>	
The Impact of File Systems on MPI-IO Scalability	87
<i>R. Latham, R. Ross, and R. Thakur</i>	
Open MPI: Goals, Concept, and Design of a Next Generation MPI Implementation	97
<i>E. Gabriel, G.E. Fagg, G. Bosilca, T. Angskun, J.J. Dongarra, J.M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R.H. Castain, D.J. Daniel, R.L. Graham, and T.S. Woodall</i>	
Open MPI's TEG Point-to-Point Communications Methodology: Comparison to Existing Implementations	105
<i>T.S. Woodall, R.L. Graham, R.H. Castain, D.J. Daniel, M.W. Sukalski, G.E. Fagg, E. Gabriel, G. Bosilca, T. Angskun, J.J. Dongarra, J.M. Squyres, V. Sahay, P. Kambadur, B. Barrett, and A. Lumsdaine</i>	
The Architecture and Performance of WMPI II	112
<i>A.L. Christensen, J. Brito, and J.G. Silva</i>	
A New MPI Implementation for Cray SHMEM	122
<i>R. Brightwell</i>	
Algorithms	
A Message Ordering Problem in Parallel Programs	131
<i>B. Uçar and C. Aykanat</i>	
BSP/CGM Algorithms for Maximum Subsequence and Maximum Subarray	139
<i>C.E.R. Alves, E.N. Cáceres, and S.W. Song</i>	

A Parallel Approach for a Non-rigid Image Registration Algorithm	147
<i>G. Román-Alonso, N.P. Castellanos-Abrego, and L. Zamora-Venegas</i>	
Neighborhood Composition: A Parallelization of Local Search Algorithms	155
<i>Y. Handa, H. Ono, K. Sadakane, and M. Yamashita</i>	
Asynchronous Distributed Broadcasting in Cluster Environment	164
<i>S. Juhász and F. Kovács</i>	
A Simple Work-Optimal Broadcast Algorithm for Message-Passing Parallel Systems	173
<i>J.L. Träff</i>	

Nesting OpenMP and MPI in the Conjugate Gradient Method for Band Systems	181
<i>L.F. Romero, E.M. Ortigosa, S. Romero, and E.L. Zapata</i>	

An Asynchronous Branch and Bound Skeleton for Heterogeneous Clusters	191
<i>J.R. González, C. León, and C. Rodríguez</i>	

Applications

Parallelization of GSL: Architecture, Interfaces, and Programming Models	199
<i>J. Aliaga, F. Almeida, J.M. Badía, S. Barrachina, V. Blanco, M. Castillo, U. Dorta, R. Mayo, E.S. Quintana, G. Quintana, C. Rodríguez, and F. de Sande</i>	

Using Web Services to Run Distributed Numerical Applications	207
<i>D. Puppin, N. Tonellootto, and D. Laforenza</i>	

A Grid-Based Parallel Maple	215
<i>D. Petcu, D. Dubu, and M. Paprzycki</i>	

A Pipeline-Based Approach for Mapping Message-Passing Applications with an Input Data Stream	224
<i>F. Guirado, A. Ripoll, C. Roig, and E. Luque</i>	

Parallel Simulations of Electrophysiological Phenomena in Myocardium on Large 32 and 64-bit Linux Clusters	234
<i>P. Czarnul and K. Grzęda</i>	

Tools and Environments

MPI I/O Analysis and Error Detection with MARMOT	242
<i>B. Krammer, M.S. Müller, and M.M. Resch</i>	

Parallel I/O in an Object-Oriented Message-Passing Library	251
<i>S. Pinkenburg and W. Rosenstiel</i>	
Detection of Collective MPI Operation Patterns	259
<i>A. Knüpfer, D. Kranzlmüller, and W.E. Nagel</i>	
Detecting Unaffected Race Conditions in Message-Passing Programs	268
<i>M.-Y. Park and Y.-K. Jun</i>	
MPI Cluster System Software	277
<i>N. Desai, R. Bradshaw, A. Lusk, and E. Lusk</i>	
A Lightweight Framework for Executing Task Parallelism on Top of MPI	287
<i>P.E. Hadjidoukas</i>	
Easing Message-Passing Parallel Programming Through a Data Balancing Service	295
<i>G. Román-Alonso, M.A. Castro-García, and J. Buenabad-Chávez</i>	
TEG: A High-Performance, Scalable, Multi-network Point-to-Point Communications Methodology	303
<i>T.S. Woodall, R.L. Graham, R.H. Castain, D.J. Daniel, M.W. Sukalski, G.E. Fagg, E. Gabriel, G. Bosilca, T. Angskun, J.J. Dongarra, J.M. Squyres, V. Sahay, P. Kambadur, B. Barrett, and A. Lumsdaine</i>	
Cluster and Grid	
Efficient Execution on Long-Distance Geographically Distributed Dedicated Clusters	311
<i>E. Argollo, J.R. de Souza, D. Rexachs, and E. Luque</i>	
Identifying Logical Homogeneous Clusters for Efficient Wide-Area Communications	319
<i>L.A. Barchet-Estefanel and G. Mounié</i>	
Coscheduling and Multiprogramming Level in a Non-dedicated Cluster	327
<i>M. Hanzich, F. Giné, P. Hernández, F. Solsona, and E. Luque</i>	
Heterogeneous Parallel Computing Across Multidomain Clusters	337
<i>P. Hwang, D. Kurzyniec, and V. Sunderam</i>	
Performance Evaluation and Monitoring of Interactive Grid Applications	345
<i>B. Baliś, M. Bubak, W. Funika, R. Wismüller, M. Radecki, T. Szepieniec, T. Arodź, and M. Kurdziel</i>	
A Domain Decomposition Strategy for GRID Environments	353
<i>B. Otero, J.M. Cela, R.M. Badia, and J. Labarta</i>	

A PVM Extension to Exploit Cluster Grids	362
<i>F. Frattolillo</i>	
Performance	
An Initial Analysis of the Impact of Overlap and Independent Progress for MPI	370
<i>R. Brightwell, K.D. Underwood, and R. Riesen</i>	
A Performance-Oriented Technique for Hybrid Application Development ..	378
<i>E. Mancini, M. Rak, R. Torella, and U. Villano</i>	
A Refinement Strategy for a User-Oriented Performance Analysis	388
<i>J. Lemeire, A. Crijns, J. Crijns, and E. Dirkx</i>	
What Size Cluster Equals a Dedicated Chip	397
<i>S. Höfinger</i>	
Architecture and Performance of the BlueGene/L Message Layer	405
<i>G. Almási, C. Archer, J. Gunnels, P. Heidelberger, X. Martorell, and J.E. Moreira</i>	
Special Session: ParSim 2004	
Current Trends in Numerical Simulation for Parallel Engineering Environments. ParSim 2004	415
<i>C. Trinitis and M. Schulz</i>	
Parallelization of a Monte Carlo Simulation for a Space Cosmic Particles Detector	417
<i>F. Almeida, C. Delgado, R.J. García López, and F. de Sande</i>	
On the Parallelization of a Cache-Optimal Iterative Solver for PDEs Based on Hierarchical Data Structures and Space-Filling Curves	425
<i>F. Günther, A. Krahne, M. Langlotz, M. Mehl, M. Pögl, and C. Zenger</i>	
Parallelization of an Adaptive Vlasov Solver	430
<i>O. Hoenen, M. Mehrenberger, and É. Violard</i>	
A Framework for Optimising Parameter Studies on a Cluster Computer by the Example of Micro-system Design	436
<i>D. Fey, M. Komann, and C. Kauhaus</i>	
Numerical Simulations on PC Graphics Hardware	442
<i>J. Krüger, T. Schiwiets, P. Kipfer, and R. Westermann</i>	
Author Index	451

PVM Grids to Self-assembling Virtual Machines

Al Geist

Oak Ridge National Laboratory
PO Box 2008, Oak Ridge, TN 37831-6016
gst@ornl.gov
<http://www.csm.ornl.gov/~geist>

Abstract. Oak Ridge National Laboratory (ORNL) leads two of the five big Genomes-to-Life projects funded in the USA. As a part of these projects researchers at ORNL have been using PVM to build a computational biology grid that spans the USA. This talk will describe this effort, how it is built, and the unique features in PVM that led the researchers to choose PVM as their framework. The computations such as parallel BLAST are run on individual supercomputers or clusters within this P2P grid and are themselves written in PVM to exploit PVM's fault tolerant capabilities.

We will then describe our recent progress in building an even more adaptable distributed virtual machine package called Harness. The Harness project includes research on a scalable, self-adapting core called H₂O, and research on fault tolerant MPI. Harness software framework provides parallel software "plug-ins" that adapt the run-time system to changing application needs in real time. This past year we have demonstrated Harness' ability to self-assemble into a virtual machine specifically tailored for particular applications.

Finally we will describe DOE's plan to create a National Leadership Computing Facility, which will house a 100 TF Cray X2 system, and a Cray Red Storm at ORNL, and an IBM Blue Gene system at Argonne National Lab. We will describe the scientific missions of this facility and the new concept of "computational end stations" being pioneered by the Facility.

1 Genomics Grid Built on PVM

The United States Department of Energy (DOE) has embarked on an ambitious computational biology program called Genomes to Life[1]. The program is using DNA sequences from microbes and higher organisms, for systematically tackling questions about the molecular machines and regulatory pathways of living systems. Advanced technological and computational resources are being employed to identify and understand the underlying mechanisms that enable living organisms to develop and survive under a wide variety of environmental conditions.

ORNL is a leader in two of the five Genomes to Life centers. As part of this effort ORNL is building a Genomics Computational Grid across the U.S. connecting ORNL, Argonne National Lab, Pacific Northwest National Lab, and

Lawrence Berkley National Lab. The software being deployed is called The Genome Channel[2]. The Genome Channel is a computational biology workbench that allows biologists to run a wide range of genome analysis and comparison studies transparently on resources at the grid sites. Genome Channel is built on top of the PVM software. When a request comes in to the Genome Channel, PVM is used to track the request, create a parallel virtual machine combining database servers, Linux clusters, and supercomputer nodes tailored to the nature of the request, spawning the appropriate analysis code on the virtual machine, and then returning the results.

The creators of this Genomics Grid require that their system be available 24/7 and that analyses that are running when a failure occurs are reconfigured around the problem and automatically restarted. PVM's dynamic programming model and fail tolerance features are ideal for this use. The Genome Channel has been cited in "Science" and used by thousands of researchers from around the world.

2 Harness: Self-assembling Virtual Machine

Harness[3] is the code name for the next generation heterogeneous distributed computing package being developed by the PVM team at Oak Ridge National Laboratory, the University of Tennessee, and Emory University. The basic idea behind Harness is to allow users to dynamically customize, adapt, and extend a virtual machine's features to more closely match the needs of their application and to optimize the virtual machine for the underlying computer resources, for example, taking advantage of a high-speed I/O. As part of the Harness project, the University of Tennessee is developing a fault tolerant MPI called FT-MPI. Emory has taken the lead in the architectural design of Harness and development of the H2O core[4].

Harness was envisioned as a research platform for investigating the concepts of parallel plug-ins, distributed peer-to-peer control, and merging and splitting of multiple virtual machines. The parallel plug-in concept provides a way for a heterogeneous distributed machine to take on a new capability, or replace an existing capability with a new method across the entire virtual machine. Parallel plug-ins are also the means for a Harness virtual machine to self-assemble. The peer-to-peer control eliminates all single points of failure and even multiple points of failure. The merging and splitting of Harness virtual machines provides a means of self healing.

The project has made good progress this past year and we have demonstrated the capability for a self-assembling virtual machine with capabilities tuned to the needs of a particular chemistry application. The second part of this talk will describe the latest Harness progress and results.

3 DOE National Leadership Computing Facility

In May of 2004 it was announced that Oak Ridge National Laboratory (ORNL) had been chosen to provide the USA's most powerful open resource for capability