

**Lecture Notes in Computer Science**      2131  
Edited by G. Goos, J. Hartmanis and J. van Leeuwen

**Springer**

*Berlin*

*Heidelberg*

*New York*

*Barcelona*

*Hong Kong*

*London*

*Milan*

*Paris*

*Tokyo*

Yiannis Cotronis Jack Dongarra (Eds.)

# Recent Advances in Parallel Virtual Machine and Message Passing Interface

8th European PVM/MPI Users' Group Meeting  
Santorini/Thera, Greece, September 23-26, 2001  
Proceedings



Springer

## **Series Editors**

Gerhard Goos, Karlsruhe University, Germany  
Juris Hartmanis, Cornell University, NY, USA  
Jan van Leeuwen, Utrecht University, The Netherlands

## **Volume Editors**

Yiannis Cotronis  
University of Athens  
Department of Informatics and Telecommunications  
Panepistimiopolis, 157 84, Greece  
E-mail: cotronis@di.uoa.gr

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

## Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Recent advances in parallel virtual machine and message passing interface :  
proceedings / 8th European PVM MPI Users' Group Meeting, Santorini/Thera,  
Greece, September 23 - 26, 2001. Yiannis Cotronis ; Jack Dongarra (ed.). -  
Berlin ; Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ;  
Paris ; Tokyo : Springer, 2001  
(Lecture notes in computer science ; Vol. 2131)  
ISBN 3-540-42609-4

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 3-540-42609-4 Springer-Verlag 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-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York  
a member of BertelsmannSpringer Science+Business Media GmbH

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2001  
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Olgun Computergrafik  
Printed on acid-free paper      SPIN 10840070      06/3142      5 4 3 2 1 0

# Preface

Parallel Virtual Machine (PVM) and Message Passing Interface (MPI) are the most frequently used tools for programming according to the message passing paradigm, which is considered one of the best ways to develop parallel applications.

This volume comprises 50 revised contributions presented at the Eighth European PVM/MPI Users' Group Meeting, which was held on Santorini (Thera), Greece, 23–26 September 2001. The conference was organized by the Department of Informatics and Telecommunications, University of Athens, Greece.

This conference has been previously held in Balatofüred, Hungary (2000), Barcelona, Spain (1999), Liverpool, UK (1998), and Krakow, Poland (1997). The first three conferences were devoted to PVM and were held at the TU Munich, Germany (1996), the ENS Lyon, France (1995), and the University of Rome (1994).

This conference has become a forum for users and developers of PVM, MPI, and other message passing environments. Interaction between these groups has 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 and improvements to PVM and MPI, algorithms using the message passing paradigm, and applications in science and engineering based on message passing. The conference included one tutorial on MPI and 9 invited talks on advances in MPI, cluster computing, network computing, Grid computing, and parallel programming and programming systems. These proceedings contain papers on the 46 oral presentations together with 4 poster presentations.

Invited speakers of Euro PVM/MPI were Frederica Darema, Al Geist, Bill Gropp, Domenico Laforenza, Phil Papadopoulos, Alexander Reinefeld, Thomas Sterling, and Vaidy Sunderam.

We would like to express our gratitude for the kind support of Compaq, HP-Germany, IBM, Microsoft, MPI Software Technology, M-Data, Silicon Computers Ltd, SUN Microsystems, TurboLinux, the Ministry of Education and Religious Affairs of Greece, and the University of Athens. Also, we would like to thank the members of the Program Committee and the other reviewers for their work in refereeing the submitted papers and ensuring the high quality of Euro PVM/MPI.

September 2001

Yiannis Cotronis  
Jack Dongarra

# Program Committee

Vassil Alexandrov	University of Reading, UK
Ranieri Baraglia	CNUCE, Pisa, Italy
Arndt Bode	LRR – Technische Universität München, Germany
Marian Bubak	Institute of Computer Science and ACC CYFRONET, AGH, Krakow, Poland
Jacques Chassin-de-Kergommeaux	ID IMAG – Grenoble, France
Yiannis Cotronis	University of Athens, Greece
José Cunha	Universidade Nova de Lisboa, Portugal
Erik D'Hollander	University of Gent, Belgium
Frederic Desprez	LIP – ENS Lyon and INRIA France
Jack Dongarra	University of Tennessee and ORNL, USA
Graham Fagg	University of Tennessee, USA
Al Geist	Oak Ridge National Labs, USA
Michael Gerndt	LRR – Technische Universität München, Germany
Andrzej Goscinski	Deakin University, Australia
Rolf Hempel	C&C Research Labs, NEC Europe Ltd., Germany
Ladislav Hluchý	Slovak Academy of Science – Bratislava, Slovakia
Peter Kacsuk	SZTAKI, Hungary
Jan Kwiatkowski	Wroclaw University of Technology, Poland
Domenico Laforeza	CNUCE-Inst. of the Italian National Res. Council, Italy
Miron Livny	University of Wisconsin – Madison, USA
Thomas Ludwig	LRR – Technische Universität München, Germany
Emilio Luque	Universitat Autònoma de Barcelona, Spain
Tomàs Margalef	Universitat Autònoma de Barcelona, Spain
Hermann Mierendorff	GMD, Germany
Shirley Moore	University of Tennessee, USA
Benno Overeinder	University of Amsterdam, The Netherlands
Andrew Rau-Chaplin	Dalhousie University – Halifax, Canada
Jeff Reeve	University of Southampton, UK
Yves Robert	LIP – Ecole Normale Supérieure de Lyon, France
Casiano Rodríguez	Universidad de La Laguna, SPAIN
Wolfgang Schreiner	RISC-Linz, – Johannes Kepler University, Linz, Austria
Miquel A. Senar	Universitat Autònoma de Barcelona, Spain
João Gabriel Silva	Universidade de Coimbra, Portugal
Vaidy Sunderam	Emory University – Atlanta, USA
Francisco Tirado	Universidad Computense de Madrid, Spain
Bernard Tourancheau	RESAM, UCB-Lyon and INRIA, France
Pavel Tvrđík	Czech Technical University, Czech Republic
Jerzy Wąsniewski	The Danish Computing Centre for Research and Education, Lyngby, Denmark
Roland Wismüller	LRR – Technische Universität München, Germany

## Additional Reviewers

Astaloš Jan	Slovak Academy of Science – Bratislava, Slovakia
Balis Bartosz	Institute of Computer Science AGH, Cracow, Poland
Balogh Zoltan	Slovak Academy of Science – Bratislava, Slovakia
Birra Fernando	Universidade Nova de Lisboa, Portugal
Buhmann Martin	Universität Giessen, Germany
Caron Eddy	LIP ENS Lyon, France
Dobrucky Miroslav	Slovak Academy of Science – Bratislava, Slovakia
Eavis Todd	Dalhousie University, Canada
Ferrini Renato	CNUCE, Pisa, Italy
Funika Włodzimierz	Institute of Computer Science AGH, Cracow, Poland
Hernandez Mario	Universidad de las Palmas de Gran Canaria, Spain
Leon Coromoto	Universidad de La Laguna, Spain
Lindermeier Markus	LRR – Technische Universität München, Germany
Luksch Peter	LRR – Technische Universität München, Germany
Martakos Drakoulis	University of Athens, Greece
Miranda Javier	Universidad de las Palmas de Gran Canaria, Spain
Namyst Raymond	LIP ENS Lyon, France
Palmerini Paolo	CNUCE, Pisa, Italy
Perego Raffaele	CNUCE, Pisa, Italy
Perez Juan C.	Universidad de La Laguna, Spain
Quinson Martin	LIP ENS Lyon, France
Rackl Gunther	LRR – Technische Universität München, Germany
Sayas Francisco-Javier	Universidad de Zaragoza, Spain
Suter Frederic	LIP ENS Lyon, France
Tran D.V.	Slovak Academy of Science – Bratislava, Slovakia
Zajac Katarzyna	Institute of Computer Science AGH, Cracow, Poland

# Table of Contents

## Invited Speakers

<i>The SPMD Model: Past, Present and Future.....</i>	1
<i>F. Darema</i>	
<i>Building a Foundation for the Next PVM: Petascale Virtual Machines .....</i>	2
<i>G.A. Geist</i>	
<i>Challenges and Successes in Achieving the Potential of MPI .....</i>	7
<i>W.D. Gropp</i>	
<i>Programming High Performance Applications in Grid Environments .....</i>	8
<i>D. Laforenza</i>	
<i>NPACI Rocks Clusters: Tools for Easily Deploying and Maintaining Manageable High-Performance Linux Clusters .....</i>	10
<i>P.M. Papadopoulos, M.J. Katz, and G. Bruno</i>	
<i>Clusters for Data-Intensive Applications in the Grid .....</i>	12
<i>A. Reinefeld</i>	
<i>A Comparative Analysis of PVM/MPI and Computational Grids .....</i>	14
<i>V. Sunderam and Z. Németh</i>	
<b>Implementation, Evaluation and Performance of PVM/MPI</b>	
<i>MPI-2 One-Sided Communications on a Gigabit Cluster .....</i>	16
<i>M. Gołębiewski and J.L. Träff</i>	
<i>Effective Communication and File-I/O Bandwidth Benchmarks .....</i>	24
<i>R. Rabenseifner and A.E. Koniges</i>	
<i>Performance of PENTRAN™ 3-D Parallel Particle Transport Code on the IBM SP2 and PCTTRAN Cluster .....</i>	36
<i>V. Kucukboyaci, A. Haghagh, and G.E. Sjoden</i>	
<i>Layering SHMEM on Top of MPI .....</i>	44
<i>L.P. Huse</i>	
<i>Support for MPI at the Network Interface Level .....</i>	52
<i>B. Tourancheau and R. Westrelin</i>	
<i>The Implementation of One-Sided Communications for WMPI II.....</i>	61
<i>T. Baptista, H. Pedroso, and J.G. Silva</i>	

Assessment of PVM Suitability to Testbed Client-Agent-Server Applications . . . . .	69
<i>M.R. Matuszek</i>	
<b>Extensions and Improvements on PVM/MPI</b>	
TH-MPI: OS Kernel Integrated Fault Tolerant MPI . . . . .	75
<i>Y. Chen, Q. Fang, Z. Du, and S. Li</i>	
CPPvm – C++ and PVM . . . . .	83
<i>S. Görzig</i>	
Persistent and Non-persistent Data Objects on Top of PVM and MPI . . . . .	91
<i>G. Manis</i>	
System Area Network Extensions to the Parallel Virtual Machine . . . . .	98
<i>M. Fischer</i>	
Adding Dynamic Coscheduling Support to PVM . . . . .	106
<i>A. Gaito, M. Rak, and U. Villano</i>	
A Model to Integrate Message Passing and Shared Memory Programming .	114
<i>J.A. González, C. León, C. Rodríguez, and F. Sande</i>	
An Architecture for a Multi-threaded Harness Kernel . . . . .	126
<i>W.R. Elwasif, D.E. Bernholdt, J.A. Kohl, and G.A. Geist</i>	
Parallel IO Support for Meta-computing Applications: MPI_Connect IO Applied to PACX-MPI . . . . .	135
<i>G.E. Fagg, E. Gabriel, M. Resch, and J.J. Dongarra</i>	
<b>Tools for PVM and MPI</b>	
TOPPER: A Tool for Optimizing the Performance of Parallel Applications . . . . .	148
<i>D. Konstantinou, N. Koziris, and G. Papakonstantinou</i>	
Programming Parallel Applications with LAMGAC in a LAN-WLAN Environment . . . . .	158
<i>E.M. Macías, A. Suárez, C.N. Ojeda-Guerra, and E. Robayna</i>	
A Dynamic Load Balancing Architecture for PDES Using PVM on Clusters . . . . .	166
<i>A.N. Pears and N. Thong</i>	
Dynamic Partitioning of the Divide-and-Conquer Scheme with Migration in PVM Environment . . . . .	174
<i>P. Czarnul, K. Tomko, and H. Krawczyk</i>	

Using Monitoring Techniques to Support the Cooperation of Software Components .....	183
<i>R. Wismüller</i>	
An Integrated Record&Replay Mechanism for Nondeterministic Message Passing Programs .....	192
<i>D. Kranzlmüller, C. Schaubenschläger, and J. Volkert</i>	
Fast and Scalable Real-Time Monitoring System for Beowulf Clusters .....	201
<i>P. Uthayopas and S. Phatanapherom</i>	
Dynamic Process Management in KSIX Cluster Middleware .....	209
<i>T. Angskun, P. Uthayopas, and A. Rungsawang</i>	
Adaptive Execution of Pipelines .....	217
<i>L.M. Moreno, F. Almeida, D. González, and C. Rodríguez</i>	
MemTo: A Memory Monitoring Tool for a Linux Cluster .....	225
<i>F. Giné, F. Solsona, X. Navarro, P. Hernández, and E. Luque</i>	
A Community Databank for Performance Tracefiles .....	233
<i>K. Ferschweiler, M. Calzarossa, C. Pancake, D. Tessera, and D. Keon</i>	
Review of Performance Analysis Tools for MPI Parallel Programs .....	241
<i>S. Moore, D. Cronk, K. London, and J. Dongarra</i>	
<b>Algorithms Using Message Passing</b>	
PVM Computation of the Transitive Closure: The Dependency Graph Approach .....	249
<i>A. Pagourtzis, I. Potapov, and W. Rytter</i>	
Parallizing 1-Dimensional Estuarine Model .....	257
<i>J. Luo, S. Rajasekaran, and C. Qiu</i>	
A Parallel ADI and Steepest Descent Methods .....	265
<i>I.V. Schevtschenko</i>	
Distributed Numerical Markov Chain Analysis .....	272
<i>M. Fischer and P. Kemper</i>	
A Parallel Algorithm for Connected Components on Distributed Memory Machines .....	280
<i>L. Buš and P. Tvrďák</i>	
Biharmonic Many Body Calculations for Fast Evaluation of Radial Basis Function Interpolants in Cluster Environments .....	288
<i>G. Roussos and B.J.C. Baxter</i>	

Heterogeneous Networks of Workstations and the Parallel Matrix Multiplication . . . . .	296
<i>F. Tinetti, A. Quijano, A. De Giusti, and E. Luque</i>	
Collecting Remote Data in Irregular Problems with Hierarchical Representation of the Domain . . . . .	304
<i>F. Baiardi, P. Mori, and L. Ricci</i>	
Parallel Image Matching on PC Cluster . . . . .	312
<i>H. Krawczyk and J. Saif</i>	
Computing Partial Data Cubes for Parallel Data Warehousing Applications . . . . .	319
<i>F. Dehne, T. Eavis, and A. Rau-Chaplin</i>	
PDES: A Case Study Using the Switch Time Warp . . . . .	327
<i>R. Suppi, F. Cores, and E. Luque</i>	
Application of MPI in Displacement Based Multilevel Structural Optimization . . . . .	335
<i>C.L. Plunkett, A.G. Striz, and J. Sobieszczanski-Sobieski</i>	
<b>Algorithms in Science and Engineering</b>	
Parallelization of Characteristics Solvers for 3D Neutron Transport . . . . .	344
<i>G.J. Wu and R. Roy</i>	
Using a Network of Workstations to Enhance Database Query Processing Performance . . . . .	352
<i>M. Al Haddad and J. Robinson</i>	
Towards a Portable, Fast Parallel AP <sup>3</sup> M-SPH Code: HYDRA_MPI . . . . .	360
<i>G.J. Pringle, S.P. Booth, H.M.P. Couchman, F.R. Pearce, and A.D. Simpson</i>	
Efficient Mapping for Message-Passing Applications Using the TTIG Model: A Case Study in Image Processing . . . . .	370
<i>C. Roig, A. Ripoll, J. Borrás, and E. Luque</i>	
Text Searching on a Heterogeneous Cluster of Workstations . . . . .	378
<i>P.D. Michailidis and K.G. Margaritis</i>	
Simulation of Forest Fire Propagation on Parallel & Distributed PVM Platforms . . . . .	386
<i>J. Jorba, T. Margalef, and E. Luque</i>	
A Data and Task Parallel Image Processing Environment . . . . .	393
<i>C. Nicolescu and P. Jonker</i>	

Evaluating the DIPORSI Framework: Distributed Processing of Remotely Sensed Imagery .....	401
<i>J.A. Gallud, J. García-Consuegra, J.M. García, and L. Orozco</i>	
Scalable Unix Commands for Parallel Processors: A High-Performance Implementation .....	410
<i>E. Ong, E. Lusk, and W. Gropp</i>	
Low-Cost Parallel Text Retrieval Using PC-Cluster .....	419
<i>A. Rungsawang, A. Laohakanniyom, and M. Lertprasertkune</i>	
Parallelization of Finite Element Package by MPI Library .....	427
<i>F. Okulicka-Dlużewska</i>	
<b>Author Index .....</b>	<b>437</b>