

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Singapore

Tokyo

Jack Dongarra Emilio Luque
Tomàs Margalef (Eds.)

Recent Advances in Parallel Virtual Machine and Message Passing Interface

6th European PVM/MPI Users' Group Meeting
Barcelona, Spain, September 26-29, 1999
Proceedings



Springer

Series Editors

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

Volume Editors

Jack Dongarra
University of Tennessee and Oak Ridge National Laboratory
107, Ayres Hall, Knoxville, TN 37996-1301, USA
E-mail: dongarra@cs.utk.edu

Emilio Luque
Tomàs Margalef
Universitat Autònoma of Barcelona, Computer Science Department
Computer Architecture and Operating Systems Group
E-08193 Bellaterra, Barcelona, Spain
E-mail: {e.luque, t.margalef}@cc.uab.es

Cataloging-in-Publication data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Recent advances in parallel virtual machine and message passing interface :
proceedings / 6th European PVM-MPI Users' Group Meeting, Barcelona, Spain,
September 26 - 29, 1999. Jack Dongarra . . . (ed.). - Berlin ; Heidelberg ; New
York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ; Singapore ; Tokyo :
Springer, 1999
(Lecture notes in computer science ; Vol. 1697)
ISBN 3-540-66549-8

CR Subject Classification (1998): D.1.3, D.3.2, F.1.2, G.1.0, B.2.1, C.1.2, C.2.4

ISSN 0302-9743

ISBN 3-540-66549-8 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 1999
Printed in Germany

Typesetting: Camera-ready by author
SPIN: 10704892 06/3142 - 5 4 3 2 1 0 Printed on acid-free paper

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 67 revised contributions presented at the Sixth European PVM/MPI Users' Group Meeting, which was held in Barcelona, Spain, 26-29 September 1999. The conference was organized by the Computer Science Department of the Universitat Autònoma de Barcelona.

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

This conference has become a forum for users and developers of PVM, MPI, and other message passing environments. Interaction between those 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 3 tutorials on advances in PVM and MPI, 6 invited talks on various message passing issues such as advanced use of MPI and PVM and integration of tools and environments, and high performance clusters; and 33 oral presentations together with 34 poster presentations. These proceedings reflect the final results of the meeting.

This final version of the invited talks and presentations has been made possible by the kind contributions of the members of the PVM/MPI '99 program committee. Each paper submitted to the conference has been reviewed by at least two reviewers. We would like to congratulate the reviewers for their efforts.

This conference has received the kind support of the following companies and administrations: Ministerio de Educación y Ciencia, Generalitat de Catalunya, Universitat Autònoma de Barcelona, EuroTools consortium, SGI and Microsoft.

September 1999

Jack Dongarra
Emilio Luque
Tomàs Margalef

Program Committee

Vassil Alexandrov	University of Liverpool, UK
Ranieri Baraglia	CNUCE, Pisa, ITALY
Arndt Bode	LRR - Technische Universität München, GERMANY
Peter Brezany	University of Vienna, AUSTRIA
Marian Bubak	Institute of Computer Science - University of Mining and Metallurgy - Krakow, POLAND
Shirley Browne	University of Tennessee, USA
Jacques Chassin	IMAG – LMC - Grenoble, FRANCE
Jens Clausen	Technical University of Denmark, DENMARK
Jeremy Cook	Parallab - University of Bergen, NORWAY
Yiannis Cotronis	University of Athens, GREECE
José Cunha	Universidade Nova de Lisboa, PORTUGAL
Ivan Dimov	Bulgarian Academy of Science - Sofia, BULGARIA
Jack Dongarra	University of Tennessee and ORNL, USA
Graham Fagg	University of Tennessee, USA
Afonso Ferreira	INRIA - Sophia-Antipolis, FRANCE
Al Geist	Oak Ridge National Labs, USA
Michael Gerndt	Forschungszentrum Juelich, GERMANY
Rolf Hempel	C&C, Research Labs. - NEC Europe Ltd., GERMANY
Erik D'Hollander	University of Gent, BELGIUM
Ladislav Hluchy	Slovak Academy of Science - Bratislava, SLOVAKIA
Robert Hood	NASA Ames Research Center, USA
Peter Kacsuk	SZTAKI, HUNGARY
Henryk Krawczyk	Technical University of Gdansk, POLAND
Jan Kwiatkowski	Wroclaw University of Technology, POLAND
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
Barton Miller	University of Wisconsin – Madison, USA
Benno Overeinder	University of Amsterdam, THE NETHERLANDS
Andrew Rau-Chaplin	Dalhousie University - Halifax, CANADA
Yves Robert	Ecole Normale Supérieure de Lyon, FRANCE
Casiano Rodríguez	Universidad de La Laguna, SPAIN
Subhash Saini	NASA Ames Research Center, USA
Wolfgang Schreiner	RISC - University of Linz, AUSTRIA
Miquel A.Senar	Universitat Autònoma de Barcelona, SPAIN
Joao Gabriel Silva	Universidade de Coimbra, PORTUGAL
Vaidy Sunderam	Emory University - Atlanta, USA
Francisco Tirado	Universidad Complutense de Madrid, SPAIN
Bernard Tourancheau	Université Claude Bernard de Lyon, FRANCE

Pavel Tvrdik	Czech Technical Univeristy, CZECH REPUBLIC
Marian Vajtersic	Slovak Academy of Science - Bratislava, SLOVAKIA
Stephen Winter	University of Westminster, UK
Jerzy Wásniewski	The Danish Computing Centre for Research and Education, Lyngby, DENMARK
Roland Wismueller	LRR - Technische Universität München, GERMANY
Zahari Zlatev	National Environmental Research Institute - Copenhagen, DENMARK

Additional Reviewers

Jan Astalos	Institute of Informatics, SAS, Slovakia
Oliver Briant	IMAG- LMC, Grenoble, France
Nuno Correia	Universidade Nova de Lisboa, Portugal
K. De Bosschere	University of Ghent, Belgium
Francisco de Sande	Universidad de La Laguna, Spain
Miroslav Dobrucky	Institute of Informatics, SAS, Slovakia
Philipp Drum	Lehrstuhl für Rechnertechnik und Rechnerorganisation, Germany
Vitor Duarte	Universidade Nova de Lisboa, Portugal
Milagros Fernández	Universidad Complutense de Madrid, Spain
R. Ferrini	CNUCE, Italy
Wlodzimierz Funika	Institute of Computer Science, AGH, Krakow, Poland
Jerome Galtier	INRIA, Sophia-Antipolis, France
Jesús A. González	Universidad de La Laguna, Spain
Owen Kaser	University of New Brunswick, Canada
D. Laforenza	CNUCE, Italy
Erwin Laure	Institute of Software Technology & Parallel Systems, University of Vienna, Austria
Coromoto León	Universidad de La Laguna, Spain
Markus Lindermeier	Technische Universität München, Germany
Ignacio M. Llorente	Universidad Complutense de Madrid, Spain
Róbert Lovas	SZTAKI, Hungary
Zsolt Nemeth	SZTAKI, Hungary
Karl L. Paap	GMD-SET, Germany
R. Perego	CNUCE, Italy
C. D. Pham	University of Lyon, France
Norbert Podhorszki	SZTAKI, Hungary
Paula Prata	Universidade da Beira Interior, Portugal
Günter Raki	Technische Universität München, Germany
José Luis Roda	Universidad de La Laguna, Spain
David Sagnol	INRIA- Sophia-Antipolis, France
Krzysztof Sowa	Institute for Software Technology & Parallel Systems, University of Vienna, Austria
T. Theoharis	University of Athens, Greece
Carsten Trinitis	Technische Universität München, Germany
Sipkova Viera	University of Vienna, Austria
Tran D. Viet	Institute of Informatics, SAS, Slovakia
J. M. Vincent	University of Grenoble, France
Roland Westrelin	University Claude Bernard, Lyon, France
Yijun Yu	University of Ghent, Belgium

Local Committee

Miquel A. Senar	Universitat Autònoma de Barcelona, Spain
José Antonio Marco	Universitat Autònoma de Barcelona, Spain
Bahjat Moh'd Qazzaz	Universitat Autònoma de Barcelona, Spain
Antonio Espinosa	Universitat Autònoma de Barcelona, Spain

Table of Contents

1 Evaluation and Performance

Performance Issues of Distributed MPI Applications in a German Gigabit Testbed	3
T. Eickermann, H. Grund, and J. Henrichs	
Reproducible Measurements of MPI Performance Characteristics	11
W. Gropp and E. Lusk	
Performance Evaluation of the MPI/MBCF with the NAS Parallel Benchmarks	19
K. Morimoto, T. Matsumoto, and K. Hiraki	
Performance and Predictability of MPI and BSP Programs on the CRAY T3E	27
J.A. González, C. Rodríguez, J.L. Roda, D.G. Morales, F. Sande, F. Almeida, and C. León	
Automatic Profiling of MPI Applications with Hardware Performance Counters	35
R. Rabenseifner	
Monitor Overhead Measurement with SKaMPI	43
D. Kranzlmüller, R. Reussner, and Ch. Schaubschläger	
A Standard Interface for Debugger Access to Message Queue Information in MPI	51
J. Cownie and W. Gropp	
Towards Portable Runtime Support for Irregular and Out-of-Core Computations	59
M. Bubak and P. Łuszczek	
Enhancing the Functionality of Performance Measurement Tools for Message Passing Environments	67
M. Bubak, W. Funika, K. Iskra, R. Maruszewski, and R. Wismüller	

Performance Modeling Based on PVM H. Mierendorff and H. Schwamborn	75
Efficient Replay of PVM Programs M. Neyman, M. Bukowski, and P. Kuzora	83
Relating the Execution Behaviour with the Structure of the Application A. Espinosa, F. Parcerisa, T. Margalef, and E. Luque	91
 2. Extensions and Improvements	
Extending PVM with Consistent Cut Capabilities: Application Aspects and Implementation Strategies A. Clematis and V. Gianuzzi	101
Flattening on the Fly: Efficient Handling of MPI Derived Datatypes J. L. Träff, R. Hempel, H. Ritzdorf, and F. Zimmermann	109
PVM Emulation in the Harness Metacomputing System: A Plug-In Based Approach M. Migliardi and V. Sunderam	117
Implementing MPI-2 Extended Collective Operations P. Silva and J. G. Silva	125
Modeling MPI Collective Communications on the AP3000 Multicomputer J. Touriño and R. Doallo	133
MPL*: Efficient Record/Replay of Nondeterministic Features of Message Passing Libraries J. Chassin de Kergommeaux, M. Ronsse, and K. De Bosschere	141
Comparison of PVM and MPI on SGI Multiprocessors in a High Bandwidth Multimedia Application R. Kutil and A. Uhl	149

On Line Visualization or Combining the Standard ORNL PVM with a Vendor PVM Implementation	157
J. Borkowski	
Native Versus Java Message Passing	165
N. Stankovic and K. Zhang	
JPT: A Java Parallelization Tool	173
K. Beyls, E. D'Hollander, and Y. Yu	
Facilitating Parallel Programming in PVM Using Condensed Graphs	181
J. P. Morrison and R. W. Connolly	
Nested Bulk Synchronous Parallel Computing	189
F. de Sande, C. León, C. Rodríguez, J. Roda, and J. A. González	
3. Implementation Issues	
An MPI Implementation on the Top of the Virtual Interface Architecture	199
M. Bertozzi, F. Boselli, G. Conte, and M. Reggiani	
MiMPI: A Multithred-Safe Implementation of MPI	207
F. García, A. Calderón, and J. Carretero	
Building MPI for Multi-Programming Systems Using Implicit Information	215
F. C. Wong, A.C. Arpaci-Dusseau, and D.E. Culler	
The Design for a High Performance MPI Implementation on the Myrinet Network	223
L. Prylli, B. Tourancheau, and R. Westrelin	
Implementing MPI's One-Sided Communications for WMPI	231
F. E. Mourão and J. G. Silva	

4. Tools

A Parallel Genetic Programming Tool Based on PVM	241
F. Fernández, J. M. Sánchez, M. Tomassini, and J.A. Gómez	
Net-Console: A Web-Based Development Environment for MPI Programs	249
A. Papagapiou, P. Evripidou, and G. Samaras	
Visual MPI, A knowledge-Based System for Writing Efficient MPI Applications	257
D. Ferenc, J. Nabrzyski, M. Stroiński, and P. Wierzejewski	

5. Algorithms

Solving Generalized Boundary Value Problems with Distributed Computing and Recursive Programming	267
I. Szeberényi and G. Domokos	
Hyper-Rectangle Distribution Algorithm for Parallel Multi-Dimensional Numerical Integration	275
R. Čiegis, R. Šablinskas, and J. Waśniewski	
Parallel Monte Carlo Algorithms for Sparse SLAE Using MPI	283
V. Alexandrov and A. Karaivanova	
A Method for Model Parameter Identification Using Parallel Genetic Algorithms	291
J. I. Hidalgo, M. Prieto, J. Lanchares, F. Tirado, B. de Andrés, S. Esteban, and D. Rivera	
Large-Scale FE Modelling in Geomechanics: A Case Study in Parallelization	299
R. Blaheta, O. Jakl, and J. Starý	
A Parallel Robust Multigrid Algorithm Based on Semi-Coarsening	307
M. Prieto, R. Santiago, I. M. Llorente, and F. Tirado	

6. Applications in Science and Engineering

PLIERS: A Parallel Information Retrieval System Using MPI	317
A. MacFarlane, J. A. McCann , and S.E. Robertson	
Parallel DSIR Text Retrieval System	325
A. Rungsawang, A. Tangpong , and P. Laohawee	
PVM Implementation of Heterogeneous ScaLAPACK Dense Linear Solvers	333
V. Boudet, F. Rastello, and Y. Robert	
Using PMD to Parallel Solve Large-Scale Navier-Stokes Equations. Performance Analysis on SGI/CRAY-T3E Machine	341
J. Chergui	
Implementation Issues of Computational Fluid Dynamics Algorithms on Parallel Computers	349
J. Płazek, K. Banaś, and J. Kitowski	
A Scalable Parallel Gauss-Seidel and Jacobi Solver for Animal Genetics	356
M. Larsen and P. Madsen	
Parallel Approaches to a Numerically Intensive Application Using PVM	364
R. Baraglia, R. Ferrini, D. Laforenza, and A. Laganà	
Solving the Inverse Toeplitz Eigenproblem Using ScaLAPACK and MPI	372
J. M. Badía and A. M. Vidal	
A Parallel Implementation of the Eigenproblem for Large, Symmetric and Sparse Matrices	380
E.M. Garzón and I. García	
Parallel Computation of the SVD of a Matrix Product	388
J. M. Claver, M. Mollar, and V. Hernández	
Porting Generalized Eigenvalue Software on Distributed Memory Machines Using Systolic Model Principles	396
P. Bassomo, I. Sakho, and A. Corbel	

Heading for an Asynchronous Parallel Ocean Model J. Schuele	404
Distributed Collision Handling for Particle-Based Simulation G. Frugoli, A. Fava, E. Fava, and G. Conte	410
Parallel Watershed Algorithm on Images from Cranial CT-Scans Using PVM and MPI on a Distributed Memory System C. Nicolescu, B. Albers, and P. Jonker	418
MPIPOV: A Parallel Implementation of POV-Ray Based on MPI A. Fava, M. Fava, and M. Bertozzi	426
Minimum Communication Cost Fractal Image Compression on PVM P. -Y. Wu	434
Cluster Computing Using MPI and Windows NT to Solve the Processing of Remotely Sensed Imagery J. A. Gallud, J. M. García, and J. García-Consuegra	442
Ground Water Flow Modelling in PVM L. Hluchý, V. D. Tran, L. Halada, and M. Dobrucký	450
 7. Networking 	
Virtual BUS: A Simple Implementation of an Effortless Networking System Based on PVM S. Ishihara, S. Tani, and A. Takahara	461
Collective Communication on Dedicated Clusters of Workstations L. P. Huse	469
Experiences Deploying a Distributed Parallel Processing Environment over a Broadband Multiservice Network J. Corbacho-Lozano., O.-I. Lepe-Aldama., J. Solé-Pareta, and J. Domingo-Pascual	477

Asynchronous Communications in MPI – the BIP/Myrinet Approach	485
F. Chaussumier, F. Desprez, and L. Prylli	
Parallel Computing on PC Clusters – An Alternative to Supercomputers for Industrial Applications	493
M. Eberl, W. Karl, C. Trinitis, and A. Blaszczyk	
Benchmarking the PVM Group Communication Efficiency	499
M.R.Matuszek, A. Mazurkiewicz, and P. W. Umiński	
8. Heterogeneous Distributed Systems	
Dynamic Assignment with Process Migration in Distributed Environments	509
P. Czarnul and H. Krawczyk	
Parallelizing of Sequential Annotated Programs in PVM Environment	517
A. Godlevsky, M. Gažák, and L. Hluchý	
Di_pSystem: A Parallel Programming System for Distributed Memory Architectures	525
F. Silva, H. Paulino, and L. Lopes	
Parallel NLP Strategies Using PVM on Heterogeneous Distributed Environments	533
G. E. Vazquez and N. B. Brignole	
Using PVM for Distributed Logic Minimization in a Network of Computers	541
L. Parrilla, J. Ortega, and A. Lloris	
Author Index	549