

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. Dózsza</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. Worringer</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. Puppín, N. Tonello, 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. Heidelberg, 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. Krahnke, M. Langlotz, M. Mehl, M. Pögl, and C. Zenger</i> | |
| Parallelization of an Adaptive Vlasov Solver | 430 |
| <i>O. Hoenen, M. Mehrenberger, and É. Violdard</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. Schiwietz, P. Kipfer, and R. Westermann</i> | |
| Author Index | 451 |

Recent Advances in Parallel Virtual Machine and
Message Passing Interface

11th European PVM/MPI Users' Group Meeting,
Budapest, Hungary, September 19-22, 2004,
Proceedings

Kranzlmüller, D.; Kacsuk, P.; Dongarra, J. (Eds.)

2004, XIV, 458 p., Softcover

ISBN: 978-3-540-23163-9