

Table of Contents

Invited Talks

The Anatomy of the Grid: Enabling Scalable Virtual Organizations	1
<i>Ian Foster</i>	
Software Component Technology for High Performance Parallel and Grid Computing	5
<i>Dennis Gannon</i>	
Macro- and Micro-parallelism in a DBMS	6
<i>Martin Kersten, Stefan Manegold, Peter Boncz and Niels Nes</i>	
An Introduction to the Gilgamesh PIM Architecture	16
<i>Thomas Sterling</i>	
High Performance Computing and Trends: Connecting Computational Requirements with Computing Resources	33
<i>Jack Dongarra</i>	

Topic 01

Support Tools and Environments	34
<i>Michael Gerndt</i>	
Dynamic Performance Tuning Environment	36
<i>Anna Morajko, Eduardo César, Tomàs Margalef, Joan Sorribes and Emilio Luque</i>	
Self-Organizing Hierarchical Cluster Timestamps	46
<i>Paul A.S. Ward and David J. Taylor</i>	
A Tool for Binding to Threads Processors	57
<i>Magnus Broberg, Lars Lundberg and Håkan Grahn</i>	
VizzScheduler – A Framework for the Visualization of Scheduling Algorithms	62
<i>Welf Löwe and Alex Liebrich</i>	
A Distributed Object Infrastructure for Interaction and Steering	67
<i>Rajeev Muralidhar and Manish Parashar</i>	
Checkpointing Facility on a Metasystem	75
<i>Yudith Cardinale and Emilio Hernández</i>	
Optimising the MPI Library for the T3E	80
<i>Stephen Booth</i>	

Topic 02

Performance Evaluation and Prediction 84
*Allen D. Malony, Graham D. Riley, Bernd Mohr, Mark Bull
and Tomàs Margalef*

Optimal Polling for Latency-Throughput Tradeoffs
in Queue-Based Network Interfaces for Clusters 86
Dmitry Ponomarev, Kanad Ghose and Eugeny Saksonov

Performance Prediction of Oblivious BSP Programs 96
*Jesús A. González, Coromoto León, Fabiana Piccoli, Marcela Printista,
José L. Roda, Casiano Rodríguez and Francisco de Sande*

Performance Prediction of Data-Dependent Task Parallel Programs 106
Hasyim Gautama and Arjan J. C. van Gemund

The Tuning Problem on Pipelines 117
*Luz Marina Moreno, Francisco Almeida,
Daniel González and Casiano Rodríguez*

The Hardware Performance Monitor Toolkit 122
Luiz A. DeRose

VIA Communication Performance on a Gigabit Ethernet Cluster 132
Mark Baker, Paul A. Farrell, Hong Ong and Stephen L. Scott

Performance Analysis of Intel’s MMX and SSE: A Case Study 142
Alfred Strey and Martin Bange

Group-Based Performance Analysis
for Multithreaded SMP Cluster Applications 148
Holger Brunst, Wolfgang E. Nagel and Hans-Christian Hoppe

Topic 03

Scheduling and Load Balancing 154
Ishfaq Ahmad, Henri Casanova, Rupert Ford and Yves Robert

On Minimising the Processor Requirements of LogP Schedules 156
Cristina Boeres, Gerson N. da Cunha and Vinod E. F. Rebello

Exploiting Unused Time Slots
in List Scheduling Considering Communication Contention 166
Oliver Sinnen and Leonel Sousa

An Evaluation of Partitioners for Parallel SAMR Applications 171
Sumir Chandra and Manish Parashar

Load Balancing on Networks with Dynamically Changing Topology 175
Jacques M. Bahi and Jaafar Gaber

A Fuzzy Load Balancing Service for Network Computing Based on Jini ...	183
<i>Lap-Sun Cheung and Yu-Kwong Kwok</i>	
Approximation Algorithms for Scheduling Independent Malleable Tasks ...	191
<i>J. Błażewicz, M. Machowiak, G. Mounié and D. Trystram</i>	
The Way to Produce the Quasi-workload in a Cluster	198
<i>Fumie Costen and John Brooke</i>	

Topic 04

Compilers for High Performance	204
<i>Jens Knoop, Manish Gupta, Keshav K. Pingali and Michael F. P. O'Boyle</i>	
Handling Irreducible Loops: Optimized Node Splitting vs. DJ-Graphs	207
<i>Sebastian Unger and Frank Mueller</i>	
Load Redundancy Elimination on Executable Code	221
<i>Manel Fernández, Roger Espasa and Saumya Debray</i>	
Loop-Carried Code Placement	230
<i>Peter Faber, Martin Griebl and Christian Lengauer</i>	
Using a Swap Instruction to Coalesce Loads and Stores	235
<i>Apan Qasem, David Whalley, Xin Yuan and Robert van Engelen</i>	
Data-Parallel Compiler Support for Multipartitioning	241
<i>Daniel Chavarría-Miranda, John Mellor-Crummey and Trushar Sarang</i>	
Cache Models for Iterative Compilation	254
<i>Peter M. W. Knijnenburg, Toru Kisuki and Kyle Gallivan</i>	
Data Sequence Locality: A Generalization of Temporal Locality	262
<i>Vincent Loechner, Benoît Meister and Philippe Clauss</i>	
Efficient Dependence Analysis for Java Arrays	273
<i>Vivek Sarkar and Stephen Fink</i>	

Topic 05

Parallel and Distributed Databases, Data Mining and Knowledge Discovery	278
<i>Harald Kosch, Pedro R. Falcone Sampaio, Abdelkader Hameurlain and Lionel Brunie</i>	
An Experimental Performance Evaluation of Join Algorithms for Parallel Object Databases	280
<i>Sandra de F. Mendes Sampaio, Jim Smith, Norman W. Paton and Paul Watson</i>	

A Classification of Skew Effects in Parallel Database Systems 291
Holger Märtens

Improving Concurrency Control in Distributed Databases
with Predeclared Tables 301
Azzedine Boukerche and Terry Tuck

Parallel Tree Projection Algorithm for Sequence Mining 310
Valerie Guralnik, Nivea Garg and George Karypis

Parallel Pruning for K-Means Clustering
on Shared Memory Architectures 321
Attila Gürsoy and İlker Cengiz

Experiments in Parallel Clustering with DBSCAN 326
Domenica Arlia and Massimo Coppola

Topic 06

Complexity Theory and Algorithms 332
Gianfranco Bilardi, Rainer Feldmann, Kieran Herley and Bruce Maggs

Beyond External Computing:
Analysis of the Cycle Structure of Permutations 333
Jörg Keller and Jop F. Sibeyn

Heaps Are Better than Buckets:
Parallel Shortest Paths on Unbalanced Graphs 343
Ulrich Meyer

Efficient Synchronization of Asynchronous Processes 352
Sandeep Lodha, Punit Chandra, Ajay Kshemkalyani and Mayank Rawat

Topic 07

Applications on High-Performance Computers 358
Yoichi Muraoka, Randall Bramley, David F. Snelling and Harry Wijshoff

Scanning Biosequence Databases on a Hybrid Parallel Architecture 360
Bertil Schmidt, Heiko Schröder and Manfred Schimmler

A Parallel Computation of Power System Equations 371
Y. F. Fung, M.F. Ercan, T.K. Ho and W. L. Cheung

Level-3 Trigger for a Heavy Ion Experiment at LHC 375
*U. Frankenfeld, H. Helstrup, J. Lien, V. Lindenstruth, D. Röhrich,
M. Schulz, B. Skaali, T. Steinbeck, K. Ullaland, A. Vestbø and A. Wiebalck*

Experiences in Using MPI-IO on Top of GPFS
for the IFS Weather Forecast Code 380
Nicholas K. Allsopp, John F. Hague and Jean-Pierre Prost

Topic 08+13

Instruction-Level Parallelism and Computer Architecture	385
<i>Eduard Ayguadé, Fredrik Dahlgren, Christine Eisenbeis, Roger Espasa, Guang R. Gao, Henk Muller, Rizos Sakellariou and André Seznec</i>	
Branch Prediction Using Profile Data	386
<i>Alex Ramirez, Josep L. Larriba-Pey and Mateo Valero</i>	
An Efficient Indirect Branch Predictor	394
<i>Yul Chu and M. R. Ito</i>	
The Behavior of <i>Efficient</i> Virtual Machine Interpreters on Modern Architectures	403
<i>M. Anton Ertl and David Gregg</i>	
Improving Conditional Branch Prediction on Speculative Multithreading Architectures	413
<i>Chitaka Iwama, Niko Demus Barli, Shuichi Sakai and Hidehiko Tanaka</i>	
Instruction Wake-Up in Wide Issue Superscalars	418
<i>Soner Önder and Rajiv Gupta</i>	
Execution Latency Reduction via Variable Latency Pipeline and Instruction Reuse	428
<i>Toshinori Sato and Itsujiro Arita</i>	
Memory Bandwidth: The True Bottleneck of SIMD Multimedia Performance on a Superscalar Processor	439
<i>Julien Sebot and Nathalie Drach-Temam</i>	
Macro Extension for SIMD Processing	448
<i>Patricio Bulić and Veselko Guštin</i>	
Performances of a Dynamic Threads Scheduler	452
<i>Smail Niar and Mahamed Adda</i>	

Topic 09

Distributed Systems and Algorithms	457
<i>Bertil Folliot, Giovanni Chiola, Peter Druschel and Anne-Marie Kermarrec</i>	
Self-stabilizing Neighborhood Unique Naming under Unfair Scheduler	458
<i>Maria Gradinariu and Colette Johnen</i>	
Event List Management in Distributed Simulation	466
<i>Jörgen Dahl, Malolan Chethur and Philip A. Wilsey</i>	
Performance Evaluation of Plausible Clocks	476
<i>Francisco J. Torres-Rojas</i>	

Building TMR-Based Reliable Servers Despite Bounded Input Lifetimes ... 482
Paul Ezhilchelvan, Jean-Michel H elary and Michel Raynal

Fractional Weighted Reference Counting 486
Erik Klinskog, Anna Neiderud, Per Brand and Seif Haridi

Topic 10

Parallel Programming: Models, Methods and Programming Languages 491
Scott B. Baden, Paul H. J. Kelly, Sergei Gorlatch and Calvin Lin

Accordion Clocks: Logical Clocks for Data Race Detection 494
Mark Christiaens and Koen De Bosschere

Partial Evaluation of Concurrent Programs 504
Matthieu Martel and Marc Gengler

A Transparent Operating System Infrastructure
 for Embedding Adaptability to Thread-Based Programming Models 514
*Ioannis E. Venetis, Dimitrios S. Nikolopoulos
 and Theodore S. Papatheodorou*

Nepal – Nested Data Parallelism in Haskell 524
*Manuel M. T. Chakravarty, Gabriele Keller, Roman Lechtchinsky
 and Wolf Pfannenstiel*

Introduction of Static Load Balancing
 in Incremental Parallel Programming 535
Joy Goodman and John O’Donnell

A Component Framework for HPC Applications 540
*Nathalie Furmento, Anthony Mayer, Stephen McGough,
 Steven Newhouse and John Darlington*

Towards Formally Refining BSP *Barriers*
 into Explicit *Two – Sided* Communications 549
Alan Stewart, Maurice Clint, Joquim Gabarr o and Maria J. Serna

Solving Bi-knapsack Problem Using Tiling Approach
 for Dynamic Programming 560
Benamar Sidi Boulenouar

Topic 11

Numerical Algorithms 566
*Henk A. van der Vorst, Rob Bisseling, Iain S. Duff
 and Bernard J. Philippe*

Parallel Implementation of a Block Algorithm for Matrix 1-Norm Estimation	568
<i>Sheung Hun Cheng and Nicholas J. Higham</i>	
Eigenvalue Spectrum Estimation and Photonic Crystals	578
<i>Ken S. Thomas, Simon J. Cox, Duan H. Beckett, Ben P. Hiatt, Jasek Generowicz and Geoffrey J. Daniell</i>	
Polynomial Preconditioning for Specially Structured Linear Systems of Equations	587
<i>Y. Liang, J. Weston and M. Szularz</i>	
Parallel Application of a Novel Domain Decomposition Preconditioner for the Stable Finite Element Solution of Three-Dimensional Convection-Dominated PDEs	592
<i>Peter K. Jimack and Sarfraz A. Nadeem</i>	
Performance of High-Accuracy PDE Solvers on a Self-Optimizing NUMA Architecture	602
<i>Sverker Holmgren and Dan Wallin</i>	

Topic 12

Routing and Communication in Interconnection Networks	611
<i>Ramón Beivide, Chris Jesshope, Antonio Robles and Cruz Izu</i>	
An Analytical Model of Deterministic Routing in the Presence of Hot-Spot Traffic	613
<i>Samia Loucif and Mohamed Ould-Khaoua</i>	
Improving the Accuracy of Reliability Models for Direct Interconnection Networks	621
<i>Rosa Alcover, Vicente Chirivella and José Duato</i>	
On Deadlock Frequency during Dynamic Reconfiguration in NOWs	630
<i>Lorenzo Fernández, José M. García and Rafael Casado</i>	
Analysis of Broadcast Communication in 2D Tori	639
<i>A. Shahrabadi, M. Ould-Khaoua and L. M. Mackenzie</i>	
Optimal Many-to-One Routing on the Mesh with Constant Queues	645
<i>Andrea Pietracaprina and Geppino Pucci</i>	

Topic 15+20

Multimedia and Embedded Systems	651
<i>Stamatis Vassiliadis, Francky Catthoor, Mateo Valero and Sorin Cotofana</i>	

XXVIII Table of Contents

A Software Architecture
for User Transparent Parallel Image Processing on MIMD Computers 653
Frank Seinstra, Dennis Koelma and Jan-Mark Geusebroek

A Case Study of Load Distribution
in Parallel View Frustum Culling and Collision Detection 663
Ulf Assarsson and Per Stenström

Parallelisable Zero-Tree Image Coding with Significance Maps 674
Rade Kutil

Performance of the Complex Streamed Instruction Set
on Image Processing Kernels 678
*Dmitri Tcheressiz, Ben Juurlink, Stamatis Vassiliadis
and Harry Wijshoff*

A Two Dimensional Vector Architecture for Multimedia 687
Ahmed El-Mahdy and Ian Watson

Multiprocessor Clustering for Embedded Systems 697
Vida Kianzad and Shuvra S. Bhattacharyya

Topic 16

Cluster Computing 702
Mark Baker, John Brooke, Ken Hawick and Rajkumar Buyya

Prioritizing Network Event Handling in Clusters of Workstations 704
Jørgen S. Hansen and Eric Jul

Fault Tolerance for Cluster Computing Based on Functional Tasks 712
Wolfgang Schreiner, Gabor Kuster and Karoly Bosa

PAPI Message Passing Library: Comparison of Performance
in User and Kernel Level Messaging 717
Eric Renault and Pierre David

Implementing Java on Clusters 722
Yariv Aridor, Michael Factor and Avi Teperman

Predictive Coscheduling Implementation
in a Non-dedicated Linux Cluster 732
Francesc Solsona, Francesc Giné, Porfidio Hernández and Emilio Luque

Self-Adjusting Scheduling of Master-Worker Applications
on Distributed Clusters 742
Elisa Heymann, Miquel A. Senar, Emilio Luque and Miron Livny

Smooth and Efficient Integration of High-Availability
in a Parallel Single Level Store System 752
Anne-Marie Kermarrec and Christine Morin

Optimal Scheduling of Aperiodic Jobs on Cluster	764
<i>Ligang He, Hai Jin, Ying Chen and Zongfen Han</i>	
HMM: A Cluster Membership Service	773
<i>Francesc D. Muñoz-Escó, Óscar Gomis, Pablo Galdámez and José M. Bernabéu-Aubán</i>	
Dynamic Processor Allocation in Large Mesh-Connected Multicomputers	783
<i>César A. F. De Rose and Hans-Ulrich Heiss</i>	
A New Communication Mechanism for Cluster Computing	793
<i>Andres Ibañez, Valentin Puente, Jose Angel Gregorio and Ramón Bevide</i>	
Isolated Dynamic Clusters for Web Hosting	801
<i>Michael Kalantar and Jun Fong</i>	

Topic 17

Metacomputing and Grid Computing	805
<i>Alexander Reinefeld, Omer F. Rana, Jarek Nabrzyski and David W. Walker</i>	
Cactus Application: Performance Predictions in Grid Environments	807
<i>Matei Ripeanu, Adriana Iamnitchi and Ian Foster</i>	
Cactus Grid Computing: Review of Current Development	817
<i>Gabrielle Allen, Werner Benger, Thomas Dramlitsch, Tom Goodale, Hans-Christian Hege, Gerd Lanfermann, André Merzky, Thomas Radke and Edward Seidel</i>	
UNICORE: A Grid Computing Environment	825
<i>Dietmar W. Erwin and David F. Snelling</i>	
Portable Parallel CORBA Objects: An Approach to Combine Parallel and Distributed Programming for Grid Computing	835
<i>Alexandre Denis, Christian Pérez and Thierry Priol</i>	
CORBA <i>Lightweight Components</i> :	
A Model for Distributed Component-Based Heterogeneous Computation ..	845
<i>Diego Sevilla, José M. García and Antonio Gómez</i>	
Building Computational Communities from Federated Resources	855
<i>Nathalie Furmento, Steven Newhouse and John Darlington</i>	
Scalable Causal Message Logging for Wide-Area Environments	864
<i>Karan Bhatia, Keith Marzullo and Lorenzo Alvisi</i>	
From Cluster Monitoring to Grid Monitoring Based on GRM	874
<i>Zoltán Balaton, Péter Kacsuk, Norbert Podhorszki and Ferenc Vajda</i>	

Use of Agent-Based Service Discovery
for Resource Management in Metacomputing Environment882
Junwei Cao, Darren J. Kerbyson and Graham R. Nudd

Topic 18

Parallel I/O and Storage Technology887
Peter Brezany, Marianne Winslett, Denis A. Nicole and Toni Cortes

Optimal Partitioning for Efficient I/O in Spatial Databases 889
Hakan Ferhatosmanoglu, Divyakant Agrawal and Amr El Abbadi

Improving Network Performance by Efficiently Dealing
with Short Control Messages in Fibre Channel SANs901
Xavier Molero, Federico Silla, Vicente Santonja and José Duato

Improving MPI-I/O Performance on PVFS 911
Jonathan Iroy, Cyrille Randriamaro and Gil Utard

Topic 19

Problem Solving Environments 916
*David W. Walker, Ken Hawick, Domenico Laforenza
and Efstratios Gallopoulos*

Remote Visualization of Distributed Electro-Magnetic Simulations 918
Erik Engquist

Solving Initial Value Problems with Parallel Maple Processes 926
Dana Petcu

Design of Problem-Solving Environment for Contingent Claim Valuation ..935
F. Oliver Bunnin, Yike Guo and John Darlington

Author Index939