

Table of Contents – Part II

HiBB 2011: 2nd Workshop on High-Performance Bioinformatics and Biomedicine

Introduction	1
<i>Mario Cannataro</i>	
On Parallelizing On-Line Statistics for Stochastic Biological Simulations	3
<i>Marco Aldinucci, Mario Coppo, Ferruccio Damiani, Maurizio Drocco, Eva Sciacca, Salvatore Spinella, Massimo Torquati, and Angelo Troina</i>	
Scalable Sequence Similarity Search and Join in Main Memory on Multi-cores	13
<i>Astrid Rheinländer and Ulf Leser</i>	
Enabling Data and Compute Intensive Workflows in Bioinformatics	23
<i>Gaurang Mehta, Ewa Deelman, James A. Knowles, Ting Chen, Ying Wang, Jens Vöckler, Steven Buyske, and Tara Matise</i>	
Homogenizing Access to Highly Time-Consuming Biomedical Applications through a Web-Based Interface	33
<i>Luigi Grasso, Nuria Medina-Medina, Rosana Montes-Soldado, and María M. Abad-Grau</i>	
Distributed Management and Analysis of Omics Data	43
<i>Mario Cannataro and Pietro Hiram Guzzi</i>	
Managing and Delivering Grid Services (MDGS)	
Introduction	53
<i>Thomas Schaaf, Adam S.Z. Belloum, Owen Appleton, Joan Serrat-Fernández, and Tomasz Szepieniec</i>	
Resource Allocation for the French National Grid Initiative	55
<i>Gilles Mathieu and Hélène Cordier</i>	
On Importance of Service Level Management in Grids	64
<i>Tomasz Szepieniec, Joanna Kocot, Thomas Schaaf, Owen Appleton, Matti Heikkurinen, Adam S.Z. Belloum, Joan Serrat-Fernández, and Martin Metzker</i>	
On-Line Monitoring of Service-Level Agreements in the Grid	76
<i>Bartosz Balis, Renata Slota, Jacek Kitowski, and Marian Bubak</i>	

Challenges of Future e-Infrastructure Governance	86
<i>Dana Petcu</i>	
Influences between Performance Based Scheduling and Service Level Agreements	96
<i>Antonella Galizia, Alfonso Quarati, Michael Schaffers, and Mark Yampolskiy</i>	
User Centric Service Level Management in mOSAIC Applications	106
<i>Massimiliano Rak, Rocco Aversa, Salvatore Venticinquè, and Beniamino Di Martino</i>	
Service Level Management for Executable Papers	116
<i>Reginald Cushing, Spiros Koulouzis, Rudolf Strijkers, Adam S.Z. Belloum, and Marian Bubak</i>	
Change Management in e-Infrastructures to Support Service Level Agreements	124
<i>Silvia Knittl, Thomas Schaaf, and Ilya Saverchenko</i>	
PROPER 2011: Fourth Workshop on Productivity and Performance: Tools for HPC Application Development	
Introduction	135
<i>Michael Gerndt</i>	
Scout: A Source-to-Source Transformator for SIMD-Optimizations	137
<i>Olaf Krzikalla, Kim Feldhoff, Ralph Müller-Pfefferkorn, and Wolfgang E. Nagel</i>	
Scalable Automatic Performance Analysis on IBM BlueGene/P Systems	146
<i>Yury Oleynik and Michael Gerndt</i>	
An Approach to Creating Performance Visualizations in a Parallel Profile Analysis Tool	156
<i>Wyatt Spear, Allen D. Malony, Chee Wai Lee, Scott Biersdorff, and Sameer Shende</i>	
INAM - A Scalable InfiniBand Network Analysis and Monitoring Tool	166
<i>N. Dandapanthula, H. Subramoni, J. Vienne, K. Kandalla, S. Sur, Dhableswar K. Panda, and Ron Brightwell</i>	
Auto-tuning for Energy Usage in Scientific Applications	178
<i>Ananta Tiwari, Michael A. Laurenzano, Laura Carrington, and Allan Snaveley</i>	

Automatic Source Code Transformation for GPUs Based on Program Comprehension	188
<i>Pasquale Cantiello and Beniamino Di Martino</i>	
Enhancing Brainware Productivity through a Performance Tuning Workflow	198
<i>Christian Iwainsky, Ralph Altenfeld, Dieter an Mey, and Christian Bischof</i>	
Workshop on Resiliency in High-Performance Computing (Resilience) in Clusters, Clouds, and Grids	
Introduction.....	209
<i>Stephen L. Scott and Chokchai (Box) Leangsuksun</i>	
The Malthusian Catastrophe Is Upon Us! Are the Largest HPC Machines Ever Up?	211
<i>Patricia Kovatch, Matthew Ezell, and Ryan Braby</i>	
Simulating Application Resilience at Exascale	221
<i>Rolf Riesen, Kurt B. Ferreira, Maria Ruiz Varela, Michela Taufer, and Arun Rodrigues</i>	
Framework for Enabling System Understanding	231
<i>J. Brandt, F. Chen, A. Gentile, Chokchai (Box) Leangsuksun, J. Mayo, P. Pebay, D. Roe, N. Taerat, D. Thompson, and M. Wong</i>	
Cooperative Application/OS DRAM Fault Recovery	241
<i>Patrick G. Bridges, Mark Hoemmen, Kurt B. Ferreira, Michael A. Heroux, Philip Soltero, and Ron Brightwell</i>	
A Tunable, Software-Based DRAM Error Detection and Correction Library for HPC	251
<i>David Fiala, Kurt B. Ferreira, Frank Mueller, and Christian Engelmann</i>	
Reducing the Impact of Soft Errors on Fabric-Based Collective Communications	262
<i>José Carlos Sancho, Ana Jokanovic, and Jesus Labarta</i>	
Evaluating Application Vulnerability to Soft Errors in Multi-level Cache Hierarchy	272
<i>Zhe Ma, Trevor Carlson, Wim Heirman, and Lieven Eeckhout</i>	
Experimental Framework for Injecting Logic Errors in a Virtual Machine to Profile Applications for Soft Error Resilience	282
<i>Nathan DeBardleben, Sean Blanchard, Qiang Guan, Ziming Zhang, and Song Fu</i>	

Automatic Source Code Transformation for GPUs Based on Program Comprehension	188
<i>Pasquale Cantiello and Beniamino Di Martino</i>	
Enhancing Brainware Productivity through a Performance Tuning Workflow	198
<i>Christian Iwainsky, Ralph Altenfeld, Dieter an Mey, and Christian Bischof</i>	
Workshop on Resiliency in High-Performance Computing (Resilience) in Clusters, Clouds, and Grids	
Introduction	209
<i>Stephen L. Scott and Chokchai (Box) Leangsuksun</i>	
The Malthusian Catastrophe Is Upon Us! Are the Largest HPC Machines Ever Up?	211
<i>Patricia Kovatch, Matthew Ezell, and Ryan Braby</i>	
Simulating Application Resilience at Exascale	221
<i>Rolf Riesen, Kurt B. Ferreira, Maria Ruiz Varela, Michela Tafer, and Arun Rodrigues</i>	
Framework for Enabling System Understanding	231
<i>J. Brandt, F. Chen, A. Gentile, Chokchai (Box) Leangsuksun, J. Mayo, P. Pebay, D. Roe, N. Tuerat, D. Thompson, and M. Wong</i>	
Cooperative Application/OS DRAM Fault Recovery	241
<i>Patrick G. Bridges, Mark Hoemmen, Kurt B. Ferreira, Michael A. Heroux, Philip Soltero, and Ron Brightwell</i>	
A Tunable, Software-Based DRAM Error Detection and Correction Library for HPC	251
<i>David Fiala, Kurt B. Ferreira, Frank Mueller, and Christian Engelmann</i>	
Reducing the Impact of Soft Errors on Fabric-Based Collective Communications	262
<i>José Carlos Sancho, Ana Jokanovic, and Jesus Labarta</i>	
Evaluating Application Vulnerability to Soft Errors in Multi-level Cache Hierarchy	272
<i>Zhe Ma, Trevor Carlson, Wim Heirman, and Lieven Eeckhout</i>	
Experimental Framework for Injecting Logic Errors in a Virtual Machine to Profile Applications for Soft Error Resilience	282
<i>Nathan DeBardleben, Sean Blanchard, Qiang Guan, Ziming Zhang, and Song Fu</i>	

High Availability on Cloud with HA-OSCAR	292
<i>Thanadech Thanakornworakij, Rajan Sharma, Blaine Scroggs, Chokchai (Box) Leangsuksun, Zeno Dixon Greenwood, Pierre Riteau, and Christine Morin</i>	
On the Viability of Checkpoint Compression for Extreme Scale Fault Tolerance	302
<i>Dewan Ibtesham, Dorian Arnold, Kurt B. Ferreira, and Patrick G. Bridges</i>	
Can Checkpoint/Restart Mechanisms Benefit from Hierarchical Data Staging?	312
<i>Raghunath Rajachandrasekar, Xiangyong Ouyang, Xavier Besseron, Vilobh Meshram, and Dhabaleswar K. Panda</i>	
Impact of Over-Decomposition on Coordinated Checkpoint/Rollback Protocol	322
<i>Xavier Besseron and Thierry Gautier</i>	
 UCHPC 2011: Fourth Workshop on UnConventional High-Performance Computing	
Introduction	333
<i>Anders Hast, Josef Weidendorfer, and Jan-Philipp Weiss</i>	
PACUE: Processor Allocator Considering User Experience	335
<i>Tetsuro Horikawa, Michio Honda, Jin Nakazawa, Kazunori Takashio, and Hideyuki Tokuda</i>	
Workload Balancing on Heterogeneous Systems: A Case Study of Sparse Grid Interpolation	345
<i>Alin Murarazu, Josef Weidendorfer, and Arndt Bode</i>	
Performance Evaluation of a Multi-GPU Enabled Finite Element Method for Computational Electromagnetics	355
<i>Tristan Cabel, Joseph Charles, and Stéphane Lanteri</i>	
Study of Hierarchical N-Body Methods for Network-on-Chip Architectures	365
<i>Thomas Canhao Xu, Pasi Liljeberg, and Hannu Tenhunen</i>	
Extending a Highly Parallel Data Mining Algorithm to the Intel® Many Integrated Core Architecture	375
<i>Alexander Heinecke, Michael Klemm, Dirk Pflüger, Arndt Bode, and Hans-Joachim Bungartz</i>	

VHPC 2011: 6th Workshop on Virtualization in High-Performance Cloud Computing

Introduction	385
<i>Michael Alexander and Gianluigi Zanetti</i>	
Group-Based Memory Deduplication for Virtualized Clouds	387
<i>Sangwook Kim, Hwanju Kim, and Joonwon Lee</i>	
A Smart HPC Interconnect for Clusters of Virtual Machines	398
<i>Anastassios Nanos, Nikos Nikoleris, Stratos Psomadakis, Elisavet Kozyri, and Nectarios Koziris</i>	
Coexisting Scheduling Policies Boosting I/O Virtual Machines	407
<i>Dimitris Aragiorgis, Anastassios Nanos, and Nectarios Koziris</i>	
PIGA-Virt: An Advanced Distributed MAC Protection of Virtual Systems	416
<i>J. Briffaut, E. Lefebvre, J. Rouzaud-Cornabas, and C. Toinard</i>	
An Economic Approach for Application QoS Management in Clouds	426
<i>Stefania Costache, Nikos Parlavantzas, Christine Morin, and Samuel Kortas</i>	
Evaluation of the HPC Challenge Benchmarks in Virtualized Environments	436
<i>Piotr Luszczek, Eric Meek, Shirley Moore, Dan Terpstra, Vincent M. Weaver, and Jack Dongarra</i>	
DISCOVERY, Beyond the Clouds: DIStributed and COoperative Framework to Manage Virtual EnviRonments autonomically: A Prospective Study	446
<i>Adrien Lèbre, Paolo Anedda, Massimo Gaggero, and Flavien Quesnel</i>	
Cooperative Dynamic Scheduling of Virtual Machines in Distributed Systems	457
<i>Flavien Quesnel and Adrien Lèbre</i>	
Large-Scale DNA Sequence Analysis in the Cloud: A Stream-Based Approach	467
<i>Romeo Kienzler, Rémy Bruggmann, Anand Ranganathan, and Nesime Tatbul</i>	
Author Index	477

Table of Contents – Part I

CCPI 2011: Workshop on Cloud Computing Projects and Initiatives

Introduction	1
<i>Beniamino Di Martino and Dana Petcu</i>	
Towards Cross-Platform Cloud Computing	5
<i>Magdalena Slawinska, Jaroslaw Slawinski, and Vaidy Sunderam</i>	
QoS Monitoring in a Cloud Services Environment: The SRT-15 Approach	15
<i>Giuseppe Cicotti, Luigi Coppolino, Rosario Cristaldi, Salvatore D’Antonio, and Luigi Romano</i>	
Enabling e-Science Applications on the Cloud with COMPSs	25
<i>Daniele Lezzi, Roger Rafanell, Abel Carrión, Ignacio Blanquer Espert, Vicente Hernández, and Rosa M. Badia</i>	
OPTIMIS and VISION Cloud: How to Manage Data in Clouds	35
<i>Spyridon V. Gogouvitis, George Kousiouris, George Vafiadis, Elliot K. Kolodner, and Dimosthenis Kyriazis</i>	
Integrated Monitoring of Infrastructures and Applications in Cloud Environments	45
<i>Roberto Palmieri, Pierangelo di Sanzo, Francesco Quaglia, Paolo Romano, Sebastiano Peluso, and Diego Didona</i>	
Towards Collaborative Data Management in the VPH-Share Project	54
<i>Siegfried Benkner, Jesus Bisbal, Gerhard Engelbrecht, Rod D. Hose, Yuriy Kaniouviskyi, Martin Koehler, Carlos Pedrinaci, and Steven Wood</i>	
SLM and SDM Challenges in Federated Infrastructures	64
<i>Matti Heikkurinen and Owen Appleton</i>	
Rapid Prototyping of Architectures on the Cloud Using Semantic Resource Description	73
<i>Houssam Haitof</i>	

Cloud Patterns for mOSAIC-Enabled Scientific Applications	83
<i>Teodor-Florin Fortiș, Gorka Esnal Lopez, Imanol Padillo Cruz, Gábor Ferschl, and Tamás Máhr</i>	
Enhancing an Autonomic Cloud Architecture with Mobile Agents	94
<i>A. Cuomo, M. Rak, S. Venticinque, and U. Villano</i>	
Mapping Application Requirements to Cloud Resources	104
<i>Yih Leong Sun, Terence Harmer, Alan Stewart, and Peter Wright</i>	
CoreGRID/ERCIM Workshop on Grids, Clouds and P2P Computing – CGWS2011	
Introduction	113
<i>Marco Danelutto, Frédéric Desprez, Vladimir Getov, and Wolfgang Ziegler</i>	
A Perspective on the <i>CoreGRID</i> Grid Component Model	115
<i>Françoise Baude</i>	
Towards Scheduling Evolving Applications	117
<i>Cristian Klein and Christian Pérez</i>	
Model Checking Support for Conflict Resolution in Multiple Non-functional Concern Management	128
<i>Marco Danelutto, P. Kilpatrick, C. Montangero, and L. Semini</i>	
Consistent Rollback Protocols for Autonomic ASSISTANT Applications	139
<i>Carlo Bertolli, Gabriele Mencagli, and Marco Vanneschi</i>	
A Dynamic Resource Management System for Real-Time Online Applications on Clouds	149
<i>Dominik Meiländer, Alexander Ploss, Frank Glinka, and Sergei Gorlatch</i>	
Cloud Federations in Contrail	159
<i>Emanuele Carlini, Massimo Coppola, Patrizio Dazzi, Laura Ricci, and Giacomo Righetti</i>	
Semi-automatic Composition of Ontologies for ASKALON Grid Workflows	169
<i>Muhammad Junaid Malik, Thomas Fahringer, and Radu Prodan</i>	

The Chemical Machine: An Interpreter for the Higher Order Chemical Language	181
<i>Vilmos Rajcsányi and Zsolt Németh</i>	
Design and Performance of the OP2 Library for Unstructured Mesh Applications	191
<i>Carlo Bertolli, Adam Betts, Gihan Mudalige, Mike Giles, and Paul Kelly</i>	
Mining Association Rules on Grid Platforms	201
<i>Raja Thili and Yahya Slimani</i>	
 5th Workshop on System-Level Virtualization for High-Performance Computing (HPCVirt 2011)	
Introduction	211
<i>Stephen L. Scott, Geoffroy Vallée, and Thomas Naughton</i>	
Performance Evaluation of HPC Benchmarks on VMware's ESXi Server	213
<i>Qasim Ali, Vladimir Kiriansky, Josh Simons, and Puneet Zaroo</i>	
Virtualizing Performance Counters	223
<i>Benjamin Serebrin and Daniel Hecht</i>	
A Case for Virtual Machine Based Fault Injection in a High-Performance Computing Environment	234
<i>Thomas Naughton, Geoffroy Vallée, Christian Engelmann, and Stephen L. Scott</i>	
 HPPC 2010: 5th Workshop on Highly Parallel Processing on a Chip	
Introduction	245
<i>Martti Forsell and Jesper Larsson Träff</i>	
Thermal Management of a Many-Core Processor under Fine-Grained Parallelism	249
<i>Fuat Keceli, Tali Moreshet, and Uzi Vishkin</i>	
Mainstream Parallel Array Programming on Cell	260
<i>Paul Keir, Paul W. Cockshott, and Andrew Richards</i>	

Generating GPU Code from a High-Level Representation for Image
Processing Kernels 270
Richard Membarth, Anton Lokhtotov, and Jürgen Teich

A Greedy Heuristic Approximation Scheduling Algorithm for 3D
Multicore Processors 281
Thomas Canhao Xu, Pasi Liljeberg, and Hannu Tenhunen

**Algorithms and Programming Tools for
Next-Generation High-Performance Scientific
Software HPSS 2011**

Introduction 293
Stefania Corsaro, Pasqua D'Ambra, and Francesca Perla

European Exascale Software Initiative: Numerical Libraries, Solvers
and Algorithms 295
Iain S. Duff

On Reducing I/O Overheads in Large-Scale Invariant Subspace
Projections 305
*Hasan Metin Aktulga, Chao Yang, Ümit V. Çatalyürek,
Pieter Maris, James P. Vary, and Esmond G. Ng*

Enabling Next-Generation Parallel Circuit Simulation with Trilinos 315
*Chris Baker, Erik Boman, Mike Heroux, Eric Keiter,
Siva Rajamanickam, Rich Schiek, and Heidi Thornquist*

DAG-Based Software Frameworks for PDEs 324
*Martin Berzins, Qingyu Meng, John Schmidt, and
James C. Sutherland*

On Partitioning Problems with Complex Objectives 334
Kamer Kaya, François-Henry Rouet, and Bora Uçar

A Communication-Avoiding Thick-Restart Lanczos Method on a
Distributed-Memory System 345
Ichitaro Yamazaki and Kesheng Wu

Spherical Harmonic Transform with GPUs 355
Ioan Ovidiu Hupca, Joel Falcou, Laura Grigori, and Radek Stompor

Design Patterns for Scientific Computations on Sparse Matrices 367
*Davide Barbieri, Valeria Cardellini, Salvatore Filippone, and
Damian Rouson*

High-Performance Matrix-Vector Multiplication on the GPU	377
<i>Hans Henrik Brandenburg Sørensen</i>	
Relaxed Synchronization with Ordered Read-Write Locks	387
<i>Jens Gustedt and Emmanuel Jeanvoine</i>	
The Parallel C++ Statistical Library ‘QUESO’: Quantification of Uncertainty for Estimation, Simulation and Optimization	398
<i>Ernesto E. Prudencio and Karl W. Schulz</i>	
Use of HPC-Techniques for Large-Scale Data Migration	408
<i>Jan Dünnweber, Valentin Mihaylov, René Glettler, Volker Maiborn, and Holger Wolff</i>	
 Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar 2011)	
Introduction	417
<i>George Bosilca</i>	
A Genetic Algorithm with Communication Costs to Schedule Workflows on a SOA-Grid	419
<i>Jean-Marc Nicod, Laurent Philippe, and Lamiel Toch</i>	
An Extension of XcalableMP PGAS Lanaguage for Multi-node GPU Clusters	429
<i>Jinpil Lee, Minh Tuan Tran, Tetsuya Odajima, Taisuke Boku, and Mitsuhisa Sato</i>	
Performance Evaluation of List Based Scheduling on Heterogeneous Systems	440
<i>Hamid Arabnejad and Jorge G. Barbosa</i>	
Column-Based Matrix Partitioning for Parallel Matrix Multiplication on Heterogeneous Processors Based on Functional Performance Models	450
<i>David Clarke, Alexey Lastovetsky, and Vladimir Rychkov</i>	
A Framework for Distributing Agent-Based Simulations	460
<i>Gennaro Cordasco, Rosario De Chiara, Ada Mancuso, Dario Mazzeo, Vittorio Scarano, and Carmine Spagnuolo</i>	
Parallel Sparse Linear Solver GMRES for GPU Clusters with Compression of Exchanged Data	471
<i>Jacques M. Bahi, Raphaël Couturier, and Lilia Ziane Khodja</i>	
Two-Dimensional Discrete Wavelet Transform on Large Images for Hybrid Computing Architectures: GPU and CELL	481
<i>Marek Błażewicz, Miłosz Ciżnicki, Piotr Kopta, Krzysztof Kurowski, and Paweł Lichoński</i>	

Scheduling Divisible Loads on Heterogeneous Desktop Systems with
Limited Memory 491
Aleksandar Ilic and Leonel Sousa

Peer Group and Fuzzy Metric to Remove Noise in Images Using
Heterogeneous Computing 502
Ma. Guadalupe Sánchez, Vicente Vidal, and Jordi Bataller

Estimation of MPI Application Performance on Volunteer
Environments 511
Girish Nandagudi, Jaspal Subhlok, Edgar Gabriel, and Judit Gimenez

Author Index 521