

Częstochowa

PPAM

2013



01010000 01001101

10th

01000000

01010000 01001101

01000001 01001101

01000001


01000001 01010000 01001101

01000000

International Conference on Parallel Processing & Applied Mathematics

Poland
WARSAW

September 8-11, 2013



PPAM 2013

10th International
Conference on
**Parallel
Processing &
Applied
Mathematics**

Book of abstracts

Poland, Warsaw
September 8-11, 2013

Contents

Main Track: Numerical Algorithms and Parallel Scientific Computing

Performance of Dense Eigensolvers on BlueGene/Q	1
<i>Inge Gutheil, Jan Felix Münchhalfen, Johannes Grotendorst</i>	
Experiences with a Lanczos Eigensolver in High-Precision Arithmetic	2
<i>Alexander Alperovich, Alex Druinsky, Sivan Toledo</i>	
Adaptive load balancing for massively parallel multi-level Monte Carlo solvers	3
<i>Jonas Šukys</i>	
A Simple Implementation of Parareal-in-Time on a Parallel Bucket-Brigade Interface	4
<i>Toshiya Takami, Daiki Fukudome</i>	
Methods for High-Throughput Computation of Elementary Functions	5
<i>Marat Dukhan, Richard Vuduc</i>	
Engineering Nonlinear Pseudorandom Number Generators	6
<i>Samuel Neves, Filipe Araujo</i>	
Extending the Generalized Fermat Prime Number Search Beyond One Million Digits Using GPUs	7
<i>Iain Bethune, Michael Goetz</i>	
Iterative solution of singular systems with applications	8
<i>Radim Blaheta, Ondrej Jakl, Jiri Starý, Erhan Turan</i>	
Statistical estimates for the conditioning of linear least squares problems	9
<i>Marc Baboulin, Serge Gratton, Rémi Lacroix, Alan Laub</i>	
Numerical treatment of a cross-diffusion model of biofilm exposure to antimicrobials	10
<i>Kazi Rahman, Hermann Eberl</i>	
Performance analysis for stencil-based 3D MPDATA algorithm on hybrid CPU-GPU platform	11
<i>Krzysztof Rojek, Lukasz Szustak, Roman Wyrzykowski</i>	
Elliptic solver performance evaluation on modern hardware architectures	12
<i>Milosz Ciznicki, Piotr Kopta, Michal Kulczewski, Krzysztof Kurowski, Pawel Gepner</i>	
Parallel Geometric Multigrid Preconditioner for Three-Dimensional FEM in NUSCA Software Package	13
<i>Tomasz Olas</i>	

Scalable Parallel Generation of Very Large Sparse Benchmark Matrices	14
<i>Daniel Langr, Ivan Šimeček, Pavel Tvrdík, Tomáš Dytrych</i>	

Main Track: Parallel Non-Numerical Algorithms

Co-operation Schemes for the Parallel Memetic Algorithm	15
<i>Jakub Nalepa, Mirosław Blocho, Zbigniew J. Czech</i>	

Optimal diffusion for load balancing in heterogeneous networks	16
<i>Katerina Dimitrakopoulou, Nikolaos Missirlis</i>	

Parallel bounded model checking of security protocols	17
<i>Olga Siedlecka-Lamch, Mirosław Kurkowski, Sabina Szymoniak, Henryk Piech</i>	

Efficient Parallel Selection	18
<i>Christian Siebert</i>	

Main Track: Environments and Tools for Distributed/Cloud/Grid Computing

Development of Domain-Specific Solutions within the Polish Infrastructure for Advanced Scientific Research	19
<i>J. Kitowski, P. Bała, M. Borcz, A. Czyżewski, Ł. Dutka, R. Kluszczyński, J. Kotus, P. Kuśtra, N. Meyer, A. Milenin, Z. Mosurska, R. Pajak, Ł. Rauch, M. Sterzel, D. Stokłosa, T. Szepieniec</i>	

Cost Optimization of Execution of Multi-level Deadline-constrained Scientific Workflows on Clouds	20
<i>Maciej Malawski, Kamil Figiela, Marian Bubak, Ewa Deelman, Jarek Nabrzyski</i>	

Parallel Computations in the Volunteer based Comcute System	21
<i>Paweł Czarnul, Jarosław Kuchta, Mariusz Matuszek</i>	

Secure storage and processing of confidential data on public clouds	22
<i>Jan Meizner, Marian Bubak, Maciej Malawski, Piotr Nowakowski</i>	

Efficient service delivery in complex heterogeneous and distributed environment	23
<i>Jan Kwiatkowski, Mariusz Fras</i>	

Domain-driven Visual Query Formulation over RDF Data Sets	24
<i>Bartosz Balis, Tomasz Grabiec, Marian Bubak</i>	

Towards High Reliability of a Multi-Agent System Designed for Intrusion Detection in MANET	25
<i>Leila Mechtri, Fatiha Djemili Tolba, Salim Ghanemi</i>	

Distributed program execution control based on application global states monitoring in PEGASUS DA framework	26
<i>Marek Tudruj, Damian Kopanski, Lukasz Masko</i>	

Main Track: Applications of Parallel Computing

- New Scalable SIMD-based Ray Caster Implementation for Virtual Machining 27
Torsten Welsch, Alexander Leutgeb, Michael Hava
- Parallelization of Permuting Schema-less XML Compressors 28
Tyler Corbin, Tomasz Muldner, Jan Krzysztof Miziołek
- Parallel Processing Model for Syntactic Pattern Recognition-Based Electrical Load Forecast 29
Mariusz Flasiński, Janusz Jurek, Tomasz Peszek
- Parallel event-driven simulation based on application global state monitoring 30
Lukasz Masko, Marek Tudruj

Main Track: Applied Mathematics, Evolutionary Computing and Metaheuristics

- It's Not a Bug, It's a Feature. Wait-free Asynchronous Cellular Genetic Algorithm 31
Frédéric Pinel, Bernabé Dorronsoro, Pascal Bouvry, Samee U. Khan
- Evolutionary algorithms for abstract planning 32
Jarosław Skaruz, Artur Niewiadomski, Wojciech Penczek
- Solution of the Inverse Continuous Casting Problem with the Aid of Modified Harmony Search Algorithm 33
Edyta Hetmaniok, Damian Słota, Adam Zielonka
- Influence of a Topology of a Spring Network on its Ability to Learn Mechanical Behaviour 34
Maja Czoków, Jacek Miękiś

Minisymposium on GPU Computing

- Evaluation of autoparallelization toolkits for commodity graphics hardware 35
David Williams, Valeriu Codreanu, Po Yang, Baoquan Liu, Feng Dong, Burhan Yasar, Babak Mahdian, Alessandro Chiarini, Xia Zhao, Jos Roerdink
- Real-Time Multiview Human Body Tracking using GPU-accelerated PSO 36
Bogusław Rymut, Bogdan Kwolek
- Implementation of a heterogeneous image reconstruction system for clinical Magnetic Resonance 37
Grzegorz Kowalik, Jennifer Steeden, David Atkinson, Andrew Taylor, Vivek Muthurangu
- X-ray laser imaging of biomolecules using multiple GPUs 38
Stefan Engblom, Jing Liu

Out-of-Core Solution of Eigenproblems for Macromolecular Simulations on GPUs	39
<i>José I. Aliaga, Davor Davidovic, Enrique S. Quintana-Ortí</i>	
GPU implementation of the Monte-Carlo simulations of the extended Ginzburg–Landau mode	40
<i>Piotr Bialas, Jakub Kowal, Adam Strzelecki</i>	
Using GPUs for parallel stencil computations in relativistic hydrodynamic simulation	41
<i>Sebastian Cygert, Daniel Kikoła, Joanna Porter-Sobieraj, Jan Sikorski, Marcin Słodkowski</i>	
Special Session on Multicore Systems	
PDNOC: An Efficient Partially DiagonalNetwork-on-Chip Design	43
<i>Thomas Canhao Xu, Ville Leppänen, Pasi Liljeberg, Juha Plosila, Hannu Tenhunen</i>	
Adaptive Fork-Heuristics for SoftwareThread-Level Speculation	44
<i>Zhen Cao, Clark Verbrugge</i>	
Inexact sparse matrix vector multiplication in Krylov subspace methods: An application-oriented reduction method	45
<i>Ahmad Mansour, Jürgen Götz</i>	
The regular expression matching algorithm for the energy efficient reconfigurable SoC	46
<i>Paweł Russek, Kazimierz Wiatr</i>	
Workshop on Models, Algorithms and Methodologies for Hierarchical Parallelism in New HPC Systems	
Transparent application acceleration by intelligent scheduling of shared library calls on heterogeneous systems	47
<i>João Colaço, Adrian Matoga, Aleksandar Ilić, Nuno Roma, Pedro Tomás, Ricardo Chaves</i>	
Improving Parallel I/O Performance Using Multithreaded Two-Phase I/O with Processor Affinity Management	48
<i>Yuichi Tsujita, Kazumi Yoshinaga, Atsushi Hori, Mikiko Sato, Mitaro Namiki, Yutaka Ishikawa</i>	
Storage Systems for Organizationally Distributed Environments - PLGrid PLUS Case Study	49
<i>Renata Słota, Łukasz Dutka, Bartosz Kryza, Darin Nikolow, Dariusz Król, Michał Wrzeszcz, Jacek Kitowski</i>	
The High Performance Internet of Things: using GVirtuS for gluing cloud computing and ubiquitous connected devices	50
<i>Raffaele Montella, Giuliano Laccetti</i>	

Workshop on Numerical Algorithms on Hybrid Architectures

Performance Evaluation of Sparse Matrix Multiplication Kernels on Intel Xeon Phi	51
<i>Erik Saule, Kamer Kaya, Ümit Çatalyürek</i>	

Portable HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi	52
<i>Jack Dongarra, Mark Gates, Azzam Haidar, Yulu Jia, Khairul Kabir, Piotr Luszczek, Stan Tomov</i>	

Accelerating a Massively Parallel Numerical Simulation in Electromagnetism using a Cluster of GPUs	53
<i>Cédric Augonnet, David Goudin, Agnès Pujols, Muriel Sesques</i>	

Multidimensional Monte Carlo Integration on Clusters with Hybrid GPU-accelerated Nodes	54
<i>Dominik Szalkowski, Przemysław Stpiczynski</i>	

Efficient Execution of Erasure Codes on AMD APU Architecture	55
<i>Roman Wyrzykowski, Marcin Woźniak, Lukasz Kuczynski</i>	

AVX acceleration of DD arithmetic between a sparse matrix and vector	56
<i>Toshiaki Hishinuma, Akihiro Fujii, Teruo Tanaka, Hidehiko Hasegawa</i>	

Using Quadruple Precision Arithmetic to Accelerate Krylov Subspace Methods on GPUs	57
<i>Daichi Mukunoki, Daisuke Takahashi</i>	

Effectiveness of sparse data structure for double-double and quad-double arithmetics	58
<i>Tsubasa Saito, Satoko Kikkawa, Emiko Ishiwata, Hidehiko Hasegawa</i>	

Efficient heuristic adaptive quadrature on GPUs: Design and Evaluation	59
<i>Daniel Thuerck, Sven Widmer, Arjan Kuijper, Michael Goesele</i>	

Square Block Code for Positive Definite Symmetric Cholesky Band Routines	60
<i>Fred G. Gustavson, José R. Herrero, Enric Moráncho</i>	

Minisymposium on Communication Avoiding Algorithms for Linear Algebra

Exploiting Data Sparsity in Parallel Matrix Powers Computations	63
<i>Nicholas Knight, Erin Carson, James Demmel</i>	

Communication Avoiding ILU0 Preconditioner	64
<i>Laura Grigori, Sophie Moufawad</i>	

Parallel Design and Performance of Nested Filtering Factorization Preconditioner	65
<i>Laura Grigori, Frédéric Nataf, Long Qu</i>	
Hiding global communication latency in the GMRES algorithm on massively parallel machines	66
<i>Wim Vanroose, Pieter Ghysels, Karl Meerbergen, Tom Ashby</i>	
Workshop on Applied High Performance Numerical Algorithms in PDEs	
A Domain Decomposition Method for Discretization of Multiscale Elliptic Problems by Discontinuous Galerkin Method	67
<i>Max Dryja</i>	
Parallel preconditioner for Finite Volume Element discretization of elliptic problem	68
<i>Leszek Marcinkowski, Talal Rahman</i>	
Abstract Schwarz Method for Nonsymmetric Local Discontinuous Galerkin Discretization of Elliptic Problem	69
<i>Filip Klawe</i>	
Fast Numerical Method for 2D Initial-Boundary Value Problems for the Boltzmann Equation	70
<i>Alexei Heintz, Piotr Kowalczyk</i>	
Simulating phase transition dynamics on nontrivial domains	71
<i>Maria Gokieli, Łukasz Bolikowski</i>	
Variable block multilevel iterative solution of general sparse linear systems	72
<i>Bruno Carpentieri, Jia Liao, Masha Sosonkina</i>	
An automatic way of finding optimal elimination trees for sequential and parallel multi-frontal direct solver for adaptive finite element method	73
<i>Hassan AbouEisha, Piotr Gurgul, Anna Paszyńska, Maciej Paszyński, Mikhail Moshkov, Krzysztof Kuźnik</i>	
Parallel efficiency of an adaptive, dynamically balanced flow solver	74
<i>Stanisław Gepner, Jerzy Majewski, Jacek Rokicki</i>	
Modification of the Newton's method for the simulations of gallium nitride semiconductor devices	75
<i>Konrad Sakowski, Leszek Marcinkowski, Stanisław Krukowski</i>	
A project of numerical realization of the one-dimensional model of burning methanol	76
<i>Krzysztof Moszyński</i>	

Workshop on Scheduling for Parallel Computing

Scheduling Bag-of-Tasks Applications to Optimize Computation Time and Cost	77
<i>Anastasia Grekioni, Natalia V. Shakhlevich</i>	

Scheduling Moldable Tasks with Precedence Constraints and Arbitrary Speedup Functions on Multiprocessors	78
<i>Sascha Hunold</i>	

OStrich: Fair Scheduling for Multiple Submissions	79
<i>Joseph Emeras, Vinicius Pinheiro, Krzysztof Rządca, Denis Trystram</i>	

Fair share is not enough: measuring fairness in scheduling with cooperative game theory	80
<i>Piotr Skowron, Krzysztof Rządca</i>	

Setting up clusters of computing units to process several data streams efficiently	81
<i>Daniel Millot, Christian Parrot</i>	

A New Multi-Criteria based Divisible Load Scheduling Algorithm	82
<i>Shamsollah Ghanbari, Mohamed Othman, Wah June Leong, Mohd Rizam Abu Bakar</i>	

Workshop on Complex Collective Systems

Bridging the Gap: From Cellular Automata to Differential Equation Models for Pedestrian Dynamics	83
<i>Felix Dietrich, Gerta Koester, Michael Seitz, Isabella von Sivers</i>	

Cellular model of pedestrian dynamics with adaptive time span	84
<i>Marek Bukáček, Pavel Hrabák, Milan Krbálek</i>	

The use of GPGPU in continuous and discrete models of crowd dynamics	85
<i>Hubert Mróz, Jarosław Wąs, Paweł Topa</i>	

Modeling Behavioral Traits of Employees in a Workplace with Cellular Automata	86
<i>Petros Saravakos, Georgios Ch. Sirakoulis</i>	

Probabilistic Pharmaceutical Modelling: A Comparison Between Synchronous and Asynchronous Cellular Automata	87
<i>Marija Bezbradica, Heather J. Ruskin, Martin Crane</i>	

Coupling Lattice Boltzmann Gas and Level Set Method for Simulating Free Surface Flow in GPU/CUDA Environment	88
<i>Tomir Kryza, Witold Dzwiniel</i>	

Creation of Agent's Vision of Social Network through Episodic Memory	89
<i>Michał Wrzeszcz, Jacek Kitowski</i>	
The influence of multi-agent cooperation on the efficiency of taxi dispatching	90
<i>Michał Maciejewski, Kai Nagel</i>	
Basic Endogenous-Money Economy: an Agent-Based Approach	91
<i>Ivan Blečić, Arnaldo Cecchini, Giuseppe A. Trunfio</i>	
Minisymposium on High Performance Computing Interval Methods	
A shaving method for interval linear systems of equations	93
<i>Milan Hladík, Jaroslav Horáček</i>	
Inner Estimation of Linear Parametric AE-Solution Sets	94
<i>Evgenija Popova</i>	
Finding Enclosures for Linear Systems using Interval Matrix Multiplication in CUDA	95
<i>Alexander Dallmann, Philip-Daniel Beck</i>	
GPU accelerated metaheuristics for solving large scale parametric interval algebraic systems	96
<i>Iwona Skalna and Jerzy Duda</i>	
Parallel approach to Monte Carlo simulation for Option Price Sensitivities using the Adjoint and Interval Analysis	97
<i>Grzegorz Kozikowski, Bartłomiej Kubica</i>	
Subsquares Approach - Simple Scheme for Solving Overdetermined Interval Linear Systems	98
<i>Jaroslav Horáček, Milan Hladík</i>	
Using quadratic approximations in an interval method of solving underdetermined and well-determined nonlinear systems	99
<i>Bartłomiej Jacek Kubica</i>	
Interval Finite Difference Method for Solving the Problem of Bioheat Transfer between Blood Vessel and Tissue	100
<i>Malgorzata A. Jankowska</i>	
Chosen Interval Methods For Solving An Interval Linear System of Equations	101
<i>Barbara Szyszka</i>	
Numerical reproducibility in HPC: the interval point of view	102
<i>Nathalie Revol, Philippe Theveny</i>	

Minisymposium on Applications of Parallel Computation in Industry and Engineering

A parallel solver for the time-periodic Navier-Stokes equations 103
Peter Arbenz, Daniel Hupp, Dominik Obrist

Parallel Numerical Algorithms for Simulation of Rectangular Waveguides
 by Using GPU 104
Raimondas Čiegis, Andrej Bugajev, Žilvinas Kancleris, Gediminas Šlekas

OpenACC Parallelisation For Diffusion Problems, Applied To Temperature
 Distribution On A Honeycomb Around the Bee Brood: A Worked Example
 Using BiCGSTAB 105
Hermann Eberl, Rangarajan Sudarsan

Application of CUDA for Acceleration of Calculations in Boundary Value
 Problems Solving Using PIES 106
Andrzej Kuźelewski, Eugeniusz Zieniuk, Agnieszka Bołtuć

Modeling and simulations of beam stabilization in edge-emitting broad area
 semiconductor devices 107
Mindaugas Radziunas, Raimondas Čiegis

Concurrent nomadic and bundle search: A class of parallel algorithms
 for local optimization 108
Costas Voglis, Dimitrios Papageorgiou, Isaac Lagaris

Parallel Multi-Objective Memetic Algorithm for Competitive Facility Location 109
Algirdas Lančinskas, Julius Žilinskas

Parallelization of Encryption Algorithm Based on Chaos System
 and Neural Networks 110
Dariusz Burak

Workshop on Language-Based Parallel Programming Models

Towards Standardization of Measuring the Usability of Parallel Languages 111
Ami Marowka

Experiences with Implementing Task Pools in Chapel and X10 112
Claudia Fohry, Jens Breitbart

Parampl: A simple approach for parallel and distributed execution
 of AMPL programs 113
Artur Olszak, Andrzej Karbowski

Optimization of an OpenCL-Based Multi-Swarm PSO Algorithm on an APU 114
Wayne Franz, Parimala Thulasiraman, Ruppa Thulasiram

Algorithms for In-Place Matrix Transposition	115
<i>Fred G. Gustavson, David W. Walker</i>	
FooPar: A Functional Object Oriented Parallel Framework in Scala	116
<i>Felix P. Hargreaves, Daniel Merkle</i>	
Prototyping framework for parallel numerical computations	117
<i>Ondřej Meca, Stanislav Böhm, Marek Běhálék, Martin Šurkovský</i>	
Core Allocation Strategies on Multicore Platforms to Accelerate Forest Fire Spread Predictions	118
<i>Tomàs Artés, Andrés Cencerrado, Ana Cortés, Tomàs Margalef</i>	
Effects of Segmented Finite Difference Time Domain on GPU	119
<i>Jose Mijares, Parimala Thulasiraman, Ruppa Thulasiram, Gagan Battoo</i>	
Workshop on Parallel Computational Biology	
Resolving Load Balancing Issues in BWA on NUMA Multicore Architectures	121
<i>Charlotte Herzeel, Thomas J. Ashby, Pascal Costanza, Wolfgang De Meuter</i>	
K-mulus: Strategies for BLAST in the Cloud	122
<i>Christopher Hill, Carl Albach, Sebastian Angel, Mihai Pop</i>	
Faster GPU-accelerated Smith-Waterman Algorithm with Alignment Backtracking for Short DNA Sequences	123
<i>Yongchao Liu, Bertil Schmidt</i>	
Accelerating string matching on MIC architecture for motif extraction	124
<i>Solon Pissis, Christian Goll, Pavlos Pavlidis, Alexandros Stamatakis</i>	
A Parallel, Distributed-memory Framework for Comparative Motif Discovery	125
<i>Dieter De Witte, Michiel Van Bel, Pieter Audenaert, Piet Demeester, Bart Dhoedt, Klaas Vandepoele, Jan Fostier</i>	
Parallel seed-based approach to protein structure similarity detection	126
<i>Guillaume Chapuis, Mathilde Le Boudic - Jamin, Rumen Andonov, Hristo Djidjev, Dominique Lavenier</i>	
Workshop on Power and Energy Aspects of Computation	
Monitoring Performance and Power for Application Characterization with Cache-aware Roofline Model	127
<i>Leonel Sousa</i>	
Energy and Deadline Constrained Robust Stochastic Static Resource Allocation	128
<i>Mark Oxley, Sudeep Pasricha, Howard Siegel, Anthony Maciejewski</i>	

Performance and Energy Analysis of the Iterative Solution of Sparse Linear Systems on Multicore and Manycore Architectures	129
<i>José I. Aliaga, Maribel Castillo, Juan C. Fernández, Germán León, Joaquín Pérez, Enrique S. Quintana-Ortí</i>	
Energy and Power Consumption Trends in HPC	130
<i>Piotr Luszczek</i>	
Measuring the Sensitivity of Graph Metrics to Missing Data	131
<i>Anita Zakrzewska, David A. Bader</i>	
The Energy/Frequency Convexity Rule: Modeling and Experimental Validation on Mobile Devices	132
<i>Karel De Voeleer, Gerard Memmi, Pierre Jouvelot, Fabien Coelho</i>	
Minisymposium on HPC Applications in Physical Sciences	
Simulations of the adsorption behavior of dendrimers	133
<i>Jaroslav Kłos, Jens -Uwe Sommer</i>	
An optimized Lattice Boltzmann code for BlueGene/Q	134
<i>Marcello Pivanti, Filippo Mantovani, Sebastiano Fabio Schifano, Raffaele Tripiccone, Luca Zenesini</i>	
A Parallel and Scalable Iterative Solver for Sequences of Dense Eigenproblems Arising in FLAPW	135
<i>Mario Berljafa, Edoardo Angelo Di Napoli</i>	
Sequential Monte Carlo in Bayesian assessment of contaminant source localization based on the sensors concentration measurements	136
<i>Anna Wawrzynczak-Szaban, Piotr Kopka, Mieczysław Borysiewicz</i>	
Effective parallelization of quantum simulations: nanomagnetic molecular rings	137
<i>Piotr Kozłowski, Grzegorz Musiał, Michał Antkowiak, Dante Gatteschi</i>	
DFT study of the Cr ₈ molecular magnet within chain-model approximations	138
<i>V. Bellini, D. M. Tomecka, Bartosz Brzostowski, M. Wojciechowski, F. Troiani, F. Manghi, M. Affronte</i>	
Workshop on Performance Evaluation of Parallel Applications on Large-Scale Systems	
The effect of parallelization on a tetrahedral mesh optimization method	139
<i>Domingo Benítez, Eduardo Rodríguez, Jose M. Escobar, Rafael Montenegro</i>	
Analysis of Partitioning Models and Metrics in Parallel Sparse Matrix-Vector Multiplication	140
<i>Kamer Kaya, Bora Uçar, Ümit Çatalyürek</i>	

Achieving Memory Scalability in the Gysela Code to Fit Exascale Constraints	141
<i>Guillaume Latu, Jean Roman, Fabien Rozar</i>	
Probabilistic analysis of barrier eliminating method applied to load-imbalanced parallel application	142
<i>Naoki Yonezawa, Ken'ichi Katou, Issei Kino, Koichi Wada</i>	
Multi-GPU parallel memetic algorithm for capacitated vehicle routing problem	143
<i>Mieczysław Wodecki, Wojciech Bożejko, Michał Karpiński, Maciej Pacut</i>	
Parallel Applications Performance Evaluation using the Concept of Granularity	144
<i>Jan Kwiatkowski</i>	
Poster session	
Improving Perfect Parallelism	145
<i>Lars Karlsson, Carl Christian Kjelgaard Mikkelsen, Bo Kågström</i>	
Parallel One-Sided Jacobi SVD Algorithm with Variable Blocking Factor	146
<i>Martin Bečka, Gabriel Okša</i>	
Using Intel Xeon Phi coprocessor to accelerate computations in MPDATA algorithm	147
<i>Lukasz Szustak, Krzysztof Rojek, Pawel Gepner</i>	
Genetic Programming in Automatic Discovery of Relationships in Computer System Monitoring Data	148
<i>Włodzimierz Funika, Pawel Koperek</i>	
Genetic Algorithms Execution Control Under a Global Application State Monitoring Infrastructure	149
<i>Adam Smyk, Marek Tudruj</i>	
Non-perturbative methods in phenomenological simulations of ring-shape molecular nanomagnets	150
<i>Piotr Kozłowski, Grzegorz Musiał, Monika Haglauer, Wojciech Florek, Michał Antkowiak, Filippo Esposito, Dante Gatteschi</i>	
Non-uniform quantum spin chains: static and dynamic properties	151
<i>Artur Barasiński, Bartosz Brzostowski, Ryszard Matysiak, Pawel Sobczak, Dariusz Woźniak</i>	
Neighborhood selection and rules identification for cellular automata: a rough sets approach	152
<i>Bartłomiej Płaczek</i>	
The Graph of Cellular Automata Applied for Modelling Tumour Induced Angiogenesis	153
<i>Paweł Topa</i>	

Preconditioning iterative substructuring methods using inexact local solvers	154
<i>Piotr Krzyżanowski</i>	
An Efficient Representation on GPU for Transition Rate Matrices for Markov Chains	155
<i>Jaroslav Bylina, Beata Bylina, Marek Karwacki</i>	
Eigen-G: GPU-based eigenvalue solver for real-symmetric dense matrices	156
<i>Toshiyuki Imamura, Susumu Yamada, Masahiko Machida</i>	
A study on adaptive algorithms for numerical quadrature on hybrid GPU and multicore based systems	157
<i>Giuliano Laccetti, Marco Lapegna, Valeria Mele, Diego Romano</i>	
The definition of interval-valued intuitionistic fuzzy sets in the framework of Dempster-Shafer theory	158
<i>Ludmila Dymova, Pavel Sevastjanov</i>	

Performance of Dense Eigensolvers on BlueGene/Q

Inge Gutheil, Jan Felix Münchhalfen, Johannes Grotendorst
Institute for Advanced Simulation
Jülich Supercomputing Centre
Forschungszentrum Jülich
GmbH, 52425 Jülich, Germany
i.gutheil@fz-juelich.de

Many scientific applications require the computation of about 10-30% of the eigenvalues and eigenvectors of large full dense symmetric or complex hermitian matrices. In this paper we will present performance evaluation results of the eigensolvers of the three libraries Elemental, ELPA, and ScaLAPACK on the BlueGene/Q architecture. All libraries include solvers for the computation of only a part of the spectrum. The most time-consuming part of the eigensolver is the reduction of the full eigenproblem to a tridiagonal one. Whereas Elemental and ScaLAPACK only offer routines to directly reduce the full matrix to a banded one, which only allows the use of BLAS 2 matrix-vector operations and needs a lot of communication, ELPA also offers a two-step reduction routine, first transforming the full matrix to banded form and thereafter to tridiagonal form. This two-step reduction shortens the reduction time significantly but at the cost of a higher complexity of the back transformation step. We will show up to which part of the eigenspectrum the use of the two-step reduction pays off.

Keywords: eigenvalue and eigenvector computation, BlueGene/Q, ScaLAPACK, ELPA, Elemental

Experiences with a Lanczos Eigensolver in High-Precision Arithmetic

Alexander Alperovich¹, Alex Druinsky², Sivan Toledo²

¹Microsoft, Haifa, Israel

²Tel Aviv University, Tel Aviv, Israel

alexdrui@post.tau.ac.il, stoledo@tau.ac.il

We investigate the behavior of the Lanczos process when it is used to find all the eigenvalues of large sparse symmetric matrices. We study the convergence of classical Lanczos (i.e., without re-orthogonalization) to the point where there is a cluster of Ritz values around each eigenvalue of the input matrix A . At that point, convergence to all the eigenvalues can be ascertained if A has no multiple eigenvalues. To eliminate multiple eigenvalues, we disperse them by adding to A a random matrix with a small norm; using high-precision arithmetic, we can perturb the eigenvalues and still compute accurate double-precision eigenvalues. Our experiments indicate that the speed with which Ritz clusters form depends on the local density of eigenvalues *and on unit roundoff*, which implies that we can accelerate convergence by using high-precision arithmetic in computations involving the Lanczos iterates.

Keywords: Lanczos, mixed precision arithmetic, Ritz clusters

Adaptive load balancing for massively parallel multi-level Monte Carlo solvers

Jonas Šukys
ETH Zürich, Switzerland
`jonas.sukys@sam.math.ethz.ch`

The Multi-Level Monte Carlo (MLMC) algorithm was shown to be a robust and fast solver for uncertainty quantification in the solutions of multi-dimensional systems of stochastic conservation laws. A novel static load balancing procedure is already developed to ensure scalability of the MLMC algorithm on massively parallel hardware up to 40 000 cores. However, for random fluxes or random initial data with large variances, the time step of the explicit time stepping scheme becomes also random due to the random CFL stability restriction. Such sample path dependent complexity of the underlying deterministic solver renders the aforementioned static load balancing very inefficient. We introduce an improved, adaptive load balancing procedure which is based on two key ingredients: 1) pre-computation of the time step size for each realization, 2) distribution of the obtained loads using the greedy algorithm for non-uniformly sized workers (core groups). Numerical experiments in multi-dimensions showing strong and weak scaling of our implementation are presented.

Keywords: uncertainty quantification, conservation laws, multi-level Monte Carlo, random time step, load balancing, linear scaling

A Simple Implementation of Parareal-in-Time on a Parallel Bucket-Brigade Interface

Toshiya Takami¹, Daiki Fukudome²

¹Research Institute for Information Technology, Kyushu University

6-10-1 Hakozaki, Higashi-ku, Fukuoka 812-8581, Japan

²Graduate School of Information Science and Electrical Engineering

Kyushu University, 744 Motooka, Nishi-ku, Fukuoka 819-0395, Japan

takami@cc.kyushu-u.ac.jp

A new simplified definition of time-domain parallelism is introduced for explicit time evolution calculations, and is implemented on parallel machines with bucket-brigade type communications. By the use of an identity operator instead of introducing an approximate solver, a recurrence formula for the parareal-in-time algorithm is much simplified. In spite of such a simple definition, it is applicable to many of explicit time-evolution calculations. In addition, this approach overcomes several drawbacks known in the original parareal-in-time method. In order to implement this algorithm on parallel machines, a parallel bucket-brigade interface is introduced, which reduces programming and tuning costs for complicated space-time parallel programs.

Keywords: parareal-in-time, bucket-brigade communication, strong scaling, massively parallel machine, scientific computing

Methods for High-Throughput Computation of Elementary Functions

Marat Dukhan, Richard Vuduc
Georgia Institute of Technology, College of Computing
266 Ferst Dr NW, Atlanta, GA 30332, USA
mdukhan3@gatech.edu

Computing elementary functions on large arrays is an essential part of many machine learning and signal processing algorithms. Since the introduction of floating-point computations in mainstream processors, table lookups, division, square root, and piecewise approximations were essential components of elementary functions implementations. However, we suggest that these operations can not deliver high throughput on modern processors, and argue that algorithms which rely only on multiplication, addition, and integer operations would achieve higher performance. We propose 4 design principles for high-throughput elementary functions, and describe how they can be applied to implementation of log, exp, sin, and tan functions. We evaluate the performance and accuracy of the new algorithms on three recent x86 microarchitectures and demonstrate that they compare favorably to previously published research and in some cases even outperform vendor-optimized libraries.

Keywords: Elementary functions, SIMD, Fused multiply-add

Engineering Nonlinear Pseudorandom Number Generators

Samuel Neves, Filipe Araujo
CISUC, Dept. of Informatics Engineering
University of Coimbra, Portugal
sneves@dei.uc.pt, filipius@uc.pt

In the era of multi- and many-core processors, computer simulations increasingly require parallel, small and fast pseudorandom number generation (PRNG). Although linear generators lend themselves to a simpler evaluation that ensures favorable properties like guaranteed period, they may negatively affect the result of simulations or be extremely slow. Conversely, non-linear generators may provide apparently random sequences, but are difficult to analyze regarding their period.

In this paper we propose non-linear PRNGs that are state-of-the-art regarding the speed/period tradeoff. These algorithms are among the fastest we are aware of for the periods they ensure: the fastest ones need states of 128 bits and provide corresponding periods of 2^{127} . However these are average numbers that may be as small as 1 in some unfavorable cases. To improve this figure, we may use different forms of counters impacting either the state or the speed of the generator. We also introduce two generators based on elliptic curves that use 2×127 bits for periods of 2^{116} and 2^{125} and low to moderate computational costs.

Keywords: Nonlinear pseudorandom generator, Elliptic curves, Elliptic curve linear congruential generator

Extending the Generalized Fermat Prime Number Search Beyond One Million Digits Using GPUs

Iain Bethune¹, Michael Goetz²

¹EPCC, The University of Edinburgh, James Clerk Maxwell Building
The King's Buildings, Mayfield Road, Edinburgh, EH9 3JZ, UK

²PrimeGrid, USA

ibethune@epcc.ed.ac.uk, mgoetz@primegrid.com

Great strides have been made in recent years in the search for ever larger prime Generalized Fermat Numbers (GFN). We briefly review the history of the GFN prime search, and describe new implementations of the ‘Genefer’ software (now available as open source) using CUDA and optimised CPU assembler which have underpinned this unprecedented progress. The results of the ongoing search are used to extend Gallot and Dubner’s published tables comparing the theoretical predictions with actual distributions of primes, and we report on recent discoveries of GFN primes with over one million digits.

Keywords: Generalized Fermat Numbers, Primality Testing, Volunteer Computing, Computational Mathematics, GPU Computing, CUDA

Iterative solution of singular systems with applications

Radim Blaheta¹, Ondrej Jakl¹, Jiri Stary¹, Erhan Turan²

¹Institute of Geonics AS CR, Czech Academy of Sciences

Ostrava, Czech Republic

²Mechran, Istanbul, Turkey

jakl@ugn.cas.cz

This paper deals with efficient solution of singular symmetric positive semidefinite problems. Our motivation arises from the need to solve special problems of geotechnics, e.g. to perform upscaling analysis of geocomposites. In those and other applications we have to solve boundary problems with pure Neumann boundary conditions. We show that the stabilized PCG with various preconditioners is a good choice for systems resulting from the numerical solution of Neumann problems, or more generally problems with a known small dimensional null space.

At the same time, we make use of this scenario to compare the implementation of the corresponding solvers in an in-house finite element software GEM and the more general solvers included in the Trilinos framework. Our case studies on a parallel solution of a problem of 6 million degrees of freedom indicate that GEM is highly competitive with its recognized counterpart.

Keywords: singular system, symmetric positive semidefinite problem, stabilized preconditioned conjugate gradient method, GEM software, Trilinos framework

Statistical estimates for the conditioning of linear least squares problems

Marc Baboulin¹, Serge Gratton², Rémi Lacroix¹, Alan Laub³

¹INRIA and Université Paris-Sud, France

²ENSEEIH and CERFACS, France

³University of California Los Angeles, USA

{marc.baboulin, remi.lacroix}@inria.fr,
serge.gratton@enseeiht.fr, laub@ats.ucla.edu

In this paper we are interested in computing linear least squares (LLS) condition numbers to measure the numerical sensitivity of an LLS solution to perturbations in data. We propose a statistical estimate for the normwise condition number of an LLS solution where perturbations on data are measured using the Frobenius norm for matrices and the Euclidean norm for vectors. We also explain how condition numbers for the components of an LLS solution can be computed. We present numerical experiments that compare the statistical condition estimates with their corresponding exact values.

Keywords: Linear least squares, Condition number, Statistical condition estimation, Componentwise conditioning

Numerical treatment of a cross-diffusion model of biofilm exposure to antimicrobials

Kazi Rahman, Hermann Eberl
University of Guelph, Guelph ON, N1G2W1, Canada
{krahman, heberl}@uoguelph.ca

We present a numerical method for a highly nonlinear PDE model of biofilm response to antibiotics with three nonlinear diffusion effects: (i) porous medium degeneracy, (ii) super-diffusion singularity, (iii) nonlinear cross-diffusion. The scheme is based on a Finite Volume discretization in space and semi-implicit, non-local time integration. The resulting discretized system is implemented in Fortran and parallelized with OpenMP. The numerical method is validated in a simulation study.

Keywords: biofilm, cross-diffusion, numerical method

Performance analysis for stencil-based 3D MPDATA algorithm on hybrid CPU-GPU platform

Krzysztof Rojek, Lukasz Szustak, Roman Wyrzykowski
Institute of Computer and Information Science
Czestochowa University of Technology
Dabrowskiego 69, 42-201 Czestochowa, Poland
{krojek, lszustak, roman}@icis.pcz.pl

EULAG (Eulerian/semi-Lagrangian fluid solver) is an established computational model for simulating thermo-fluid flows across a wide range of scales and physical scenarios. The multidimensional positive defined advection transport algorithm (MPDATA) is among the most time-consuming components of EULAG.

The main contribution of our work is to design an efficient adaptation of the MPDATA algorithm to the hybrid of Intel CPU SandyBridge with NVIDIA GPU Kepler architecture. We focus on analysis of resources usage in the CPU-GPU platform and its influence on the performance results. In this paper, a performance model is proposed, which makes a deep analysis of the resources consumption including registers, shared, global and texture memory. The performance model allows us to identify bottlenecks of the algorithm and indicates the right direction of optimization.

The group of the most common bottlenecks is considered in this work, including data transfers between host memory and GPU global memory, GPU global memory and shared memory, latencies and serialization of instructions, as well as GPU occupancy. We put the emphasis on providing fixed memory access pattern, padding, reducing divergent branches and instructions latencies, as well as organizing computations in the MPDATA algorithm in order to efficient shared memory and register file reusing.

Keywords: GPGPU, CUDA, Hybrid programming, EULAG, stencil, MPDATA, geophysical flows, parallel programming

Elliptic solver performance evaluation on modern hardware architectures

Milosz Ciznicki¹, Piotr Kopta¹, Michal Kulczewski¹, Krzysztof Kurowski¹,
Pawel Gepner²

¹Poznan Supercomputing and Networking Center
Noskowskiego 10 Street, 61-704 Poznan, Poland

²Intel Corporation

michal.kulczewski@man.poznan.pl

The recent advent of novel multi- and many-core architectures forces application programmers to deal with hardware-specific implementation details and to be familiar with software optimisation techniques to benefit from new high-performance computing machines. An extra care must be taken for communication-intensive algorithms, which may be a bottleneck for forthcoming era of exascale computing. This paper aims to present performance evaluation of preliminary adaptation techniques to hybrid MPI+OpenMP parallelisation schemes we provided into the EULAG code. Various techniques are discussed, and the results will lead us toward efficient algorithms and methods to scale communication-intensive elliptic solver with preconditioner, including GPU architectures to be provided later in the future.

Keywords: Eulag, elliptic solver, high performance computing, petascale computing, hybrid parallelisation, Intel Xeon Phi

Parallel Geometric Multigrid Preconditioner for Three-Dimensional FEM in NuscaS Software Package

Tomasz Olas
Institute of Computer and Information Science
Czestochowa University of Technology
Dabrowskiego 69, 42-201 Czestochowa, Poland
olas@icis.pcz.pl

Multigrid methods are among the fastest numerical algorithms for solving large sparse systems of linear equations. The Conjugate Gradient (CG) method with Multigrid as a preconditioner (MGCG) features a good convergence even when the Multigrid solver itself is not efficient.

The parallel FEM package NuscaS allows us to solve adaptive FEM problems with 3D unstructured meshes on parallel computers such as PC-clusters. The parallel version of the library is based on the geometric decomposition applied for computing nodes of a parallel system; the distributed-memory architecture and message-passing model of parallel programming are assumed. In previous works we extend the NuscaS functionality by introducing parallel adaptation of tetrahedral FEM meshes and dynamic load balancing.

In this work we discuss key issues of efficient implementation of parallel geometric multigrid preconditioner for 3D FEM problems in the NuscaS package. Based on the geometric decomposition, for each level of multigrid meshes are partitioned and assigned to corresponding processors of a parallel architecture. Fine-grid levels are constructed by subdivision of mesh elements using the parallel 8-tetrahedra longest-edge refinement mesh algorithm. Every process keeps assigned part of mesh on each level of multigrid.

Keywords: Parallel algorithms, Multigrid methods, FEM

Scalable Parallel Generation of Very Large Sparse Benchmark Matrices

Daniel Langr¹, Ivan Šimeček¹, Pavel Tvrdík¹, Tomáš Dytrych²

¹Czech Technical University in Prague

Department of Computer Systems, Faculty of Information Technology
Thákurova 9, 160 00, Praha, Czech Republic

²Louisiana State University

Department of Physics and Astronomy
Baton Rouge, LA 70803, USA

langrd@fit.cvut.cz

We present a method and an accompanying algorithm for scalable parallel generation of sparse matrices intended primarily for benchmarking purposes, namely for evaluation of performance of generic massively parallel algorithms that involve sparse-matrix computations. The proposed method is based on enlargement of small input matrices, which are supposed to be obtained from public sparse matrix collections containing numerous matrices arising in different application domains and thus having different structural and numerical properties. The resulting matrices are distributed among processors of a parallel computer system. The enlargement process is designed so its users may easily control structural and numerical properties of resulting matrices as well as the distribution of their nonzero elements to particular processors.

Keywords: sparse matrix, benchmark matrix, enlargement, parallel algorithm, scalability

Co-operation Schemes for the Parallel Memetic Algorithm

Jakub Nalepa¹, Mirosław Blocho^{1,2}, Zbigniew J. Czech^{1,3}

¹Silesian University of Technology, Gliwice, Poland

²ABB ISDC, Kraków, Poland

³University of Silesia, Sosnowiec, Poland

{jakub.nalepa, zbigniew.czech}@polsl.pl

miroslaw.blocho@pl.abb.com

This paper presents a study of co-operation schemes for the parallel memetic algorithm to solve the vehicle routing problem with time windows. In the parallel co-operative search algorithms the processes communicate with each other in order to exchange the up-to-date solutions, which may guide the search and improve the results. The interactions between processes are defined by the content of the exchanged data, timing, connectivity and mode. We present how co-operation schemes influence the speed of search convergence and solutions quality. The quality of a feasible solution is defined as its proximity to the best, currently-known one. We present the experimental study for the well-known Gehring and Homberger's benchmark tests. The new world's best solutions obtained in the study confirm that the choice of a co-operation scheme has a strong impact on the quality of final solutions.

Keywords: parallel memetic algorithm, parallel processes co-operation schemes, genetic and local search algorithms, vehicle routing problem with time windows

Optimal diffusion for load balancing in heterogeneous networks

Katerina Dimitrakopoulou, Nikolaos Missirlis
Department of Informatics and Telecommunications
University of Athens
Panepistimiopolis, 157 84, Athens, Greece
{kdim, nmis}@di.uoa.gr

In [11] we studied the local Extrapolated Diffusion (EDF) method for the load balancing problem in case of homogeneous torus networks. The present paper develops the convergence theory of the local EDF for heterogeneous torus networks. In particular, we determine its quasi-optimal iteration parameters and the corresponding quasi-optimal convergence factor using local mode analysis. As a result dynamic load balancing becomes an efficient procedure since the parameters of local EDF are computed via a closed form formulae resulting in the maximization of its rate of convergence. Moreover, it is shown how the convergence factor depends upon the communication edge weights and the processor speeds of the network.

Keywords: heterogeneous, Fourier, Diffusion

Parallel bounded model checking of security protocols

Olga Siedlecka-Lamch, Mirosław Kurkowski, Sabina Szymoniak, Henryk Piech
Institute of Computer and Information Science
Czestochowa University of Technology
Dabrowskiego 69, 42-201 Czestochowa, Poland
{mkurkowski, olga.siedlecka}@icis.pcz.pl

The verification of security protocols is a difficult process taking into consideration a concept and computations. The difficulties start just during the protocol specification and during studying its properties. In case of the computation connected with constructing and searching of the modeling structures of protocol execution and scattered knowledge of the users the problems are the sizes of those structures. For small values of parameters such as numbers of sessions, users, or encryption keys the proper models are usually not very big and searching them is not a problem, however in case of increasing the values of above mentioned parameters the models are sometimes too big and there is no possibility to construct them or search properly. In order to increase the values of studying protocol parameters and necessary increase of the computation effectiveness, the appropriate solutions must be introduced. In the article the solutions which enable full and effective parallelization of the computations during automatic verification of security protocols are introduced. The suitable experimental results are also presented.

Keywords: Parallel model checking, Security protocols, Protocols' verification

Efficient Parallel Selection

Christian Siebert
RWTH Aachen University
Department of Computer Science
Laboratory for Parallel Programming
`christian.siebert@rwth-aachen.de`

Selection algorithms find the k -th smallest element from a set of elements. Although there are optimal parallel selection algorithms available for theoretical machines, these algorithms are not only difficult to implement but also inefficient in practice. Consequently, parallel applications usually use only the special cases minimum and maximum, where efficient implementations exist as easy-to-use functionality. We present a general parallel selection solution that scales even on today's largest supercomputers. Our approach is based on an efficient, unbiased median approximation method, recently introduced as median-of-3 reduction, and Hoare's original QuickSelect idea from 1961. We explain implementation details, provide a statistical and theoretical analysis of the median approximation method and show performance results using more than 400,000 processor cores. Furthermore, we discuss applications in the area of parallel performance analysis and modeling, which can benefit from such an efficient parallel selection solution.

Keywords: Selection, QuickSelect, Median, Parallel Algorithms, MPI

Development of Domain-Specific Solutions within the Polish Infrastructure for Advanced Scientific Research

J. Kitowski^{1,2}, P. Bała^{5,6}, M. Borcz⁵, A. Czyzewski⁴, Ł. Dutka², R. Kluszczynski⁵, J. Kotus⁴,
P. Kustra³, N. Meyer⁷, A. Milenin³, Z. Mosurska², R. Pajak², Ł. Rauch³, M. Sterzel²,
D. Stokłosa⁷, T. Szepieniec²

¹AGH University, Department of Computer Science, Kraków, Poland

²AGH University, ACC Cyfronet AGH, Kraków, Poland

³AGH University, Department of Applied Computer Science and Modelling
Kraków, Poland

⁴Gdansk University of Technology, ETI Faculty
Multimedia Systems Department, Gdansk, Poland

⁵University of Warsaw, Interdisciplinary Centre for Mathematical
and Computational Modelling, Warsaw, Poland

⁶Nicolaus Copernicus University, Department of Mathematics
and Computer Science, Torun, Poland

⁷Poznan Supercomputing and Networking Center, Poznan, Poland

kito@agh.edu.pl, bala@mat.umk.pl, joseph@ssound.eti.pg.gda.pl,
meyer@man.poznan.pl, z.mosurska@cyfronet.pl

The Polish Grid computing infrastructure was established during the PL-Grid project (2009-2012). The main purpose of this Project was to provide the Polish scientists with an IT basic platform, allowing them to conduct interdisciplinary research on a national scale, and giving them transparent access to international grid resources via international grid infrastructures. Currently, the infrastructure is maintained and extended within a follow-up PLGrid Plus project (2011-2014). Its main objective is to increase the potential of the Polish Science by providing necessary IT services for research teams in Poland, in line with European solutions. The paper presents several examples of the domain-specific computational environments, developed within the Project. For particular environments, specialized IT solutions are prepared, i.e. dedicated software implementation and infrastructure adaptation, suited for particular researchers' demands.

Keywords: PL-Grid, PLGrid Plus, domain-specific solutions, computing infrastructure, computational environment

Cost Optimization of Execution of Multi-level Deadline-constrained Scientific Workflows on Clouds

Maciej Malawski¹, Kamil Figiela¹, Marian Bubak^{1,2}, Ewa Deelman³,
Jarek Nabrzyski⁴

¹Department of Computer Science AGH,
al. Mickiewicza 30, 30-059 Kraków, Poland

²ACC CYFRONET AGH, ul. Nawojki 11, 30-950 Kraków, Poland

³USC Information Sciences Institute,
4676 Admiralty Way, Marina del Rey, CA, USA

⁴University of Notre Dame, Center for Research Computing, Notre Dame, IN, USA
{malawski,bubak}@agh.edu.pl, fkamil@student.agh.edu.pl,
deelman@isi.edu, naber@nd.edu

This paper introduces a cost optimization model for scientific workflows on IaaS clouds such as Amazon EC2 or RackSpace. We assume multiple IaaS clouds with heterogeneous VM instances, with limited number of instances per cloud and hourly billing. Input and output data are stored on a Cloud Object Store such as Amazon S3. Applications are scientific workflows modeled as DAGs as in the Pegasus Workflow Management System. We assume that tasks in the workflows are grouped into levels of identical tasks. Our model is specified in AMPL modeling language and allows to minimize the cost of workflow execution under deadline constraints. We present results obtained using our model and the benchmark workflows representing real scientific applications such as Montage, Epigenomics, LIGO. We indicate how this model can be used for scenarios that require resource planning for scientific workflows and their ensembles.

Keywords: AMPL Optimization, cloud computing, scientific workflows

Parallel Computations in the Volunteer based Comcute System

Paweł Czarnul, Jarosław Kuchta, Mariusz Matuszek
Faculty of Electronics, Telecommunications and Informatics
Gdansk University of Technology, Poland
{pczarnul,qhta,mrm}@eti.pg.gda.pl

Volunteer computing plays an important role in supplying the computational power demanded by science. Harnessing the power of personal computers connected to the Internet requires dedicated systems, which distribute computations and collect results. Usually such systems require the user of a personal computer to install and run a dedicated client software, which often presents a difficulty to less technical-savvy users. Comcute was designed to let users donate the computing power of their PCs in a simplified manner, requiring only pointing their web browser at a specific web address and clicking a mouse. Additionally, the server side allows execution that can survive failures of individual computers and definition of redundancy of desired order. This paper presents results of scalability experiments carried on the Comcute system.

Keywords: Volunteer computing, parallel computations, scalability, reliability

Secure storage and processing of confidential data on public clouds

Jan Meizner¹, Marian Bubak^{2,3}, Maciej Malawski², Piotr Nowakowski¹

¹ACC Cyfronet AGH,

AGH University of Science and Technology, Kraków, Poland

²Department of Computer Science,

AGH University of Science and Technology, Kraków, Poland

³Informatics Institute,

University of Amsterdam, The Netherlands

{bubak, malawski}@agh.edu.pl

{jan.meizner, p.nowakowski}@cyfronet.pl

The goal of this paper is to describe problems associated with storage and processing of confidential data in public clouds, and to suggest relevant mitigation strategies. It is clear that many types of data in the commercial and scientific worlds require special handling. This issue affects highly valuable data, such as trade secrets and financial information, as well as personal data including medical records used in scientific research. In addition to discussing situations which require special care, we propose a set of applicable solutions and describe feasibility studies based on tests performed using popular cryptographic software (OpenSSL).

Keywords: security, data, clouds, hybrid, encryption, AES, 3DES, OpenSSL

Efficient service delivery in complex heterogeneous and distributed environment

Jan Kwiatkowski, Mariusz Fras
Wrocław University of Technology
Poland
{jan.kwiatkowski, mariusz.fras}@pwr.wroc.pl

The problem of providing support for quality of service (QoS) guarantees is studied in many areas of information technologies. In recent years the evolution of software architectures led to the rising prominence of the Service Oriented Architecture (SOA) concept. For Web-based systems there are three attributes that directly relate to everyday perception of the QoS for the end user: availability, usability, and performance. The paper focuses on performance issues of service delivery. The architecture of Virtual Service Delivery System (VSDS), a tool to serve requests for synchronized services is presented. It is proposed suitable monitoring technique used for estimation of values of service parameters and allocation of communication and execution resources by means of service distribution. The paper also presents results of experiments performed in real environment that show effectiveness of proposed solution.

Keywords: service request distribution, service monitoring, quality of services

Domain-driven Visual Query Formulation over RDF Data Sets

Bartosz Balis¹, Tomasz Grabiec², Marian Bubak^{1,2}

¹AGH University of Science and Technology

Department of Computer Science

²AGH University of Science and Technology

ACC Cyfronet AGH, Krakow, Poland

{balis,bubak}@agh.edu.pl

Semantic Web technologies, such as RDF, SPARQL and OWL, are increasingly used for data representation and information retrieval in real-world applications including those from the e-Science domain. Visual query formulation has been recognized as a useful approach facilitating information retrieval for domain-experts not familiar with RDF query languages. We propose a visual query approach over RDF data sets based on an abstract domain-driven (conceptual) query language. The basis for the query model are ontologies describing the RDF data sets. We have built the QUaTRO2 tool which implements this query approach and provides a high usability graphical user interface to assist domain experts in constructing complex queries and browsing their results. The concepts and their implementation are validated by applying the QUaTRO2 tool to query the UniProt protein database.

Keywords: Semantic Web, RDF, SPARQL, visual query formulation, OWL, ontologies

Towards High Reliability of a Multi-Agent System Designed for Intrusion Detection in MANET

Leila Mechtri, Fatiha Djemili Tolba, Salim Ghanemi
Badji-Mokhtar University, Computer Science Dept.
P. O. 12, 23000 Annaba, Algeria
mechteri@lrs-annaba.net, fatiha.djemili@univ-annaba.org,
ghanemisalim@yahoo.com

Replication is widely used to enhance both reliability and performance in distributed systems. In this paper, we present MASID-R, a replication framework for replicating MASID (Multi-Agent System for Intrusion Detection) agents, in which two varieties of agent replication are combined, namely replication at system initialization and replication on demand. The key difference between standard replication approaches and our approach is that the former depends on the notion of replicate groups, while MASID-R allows more flexibility through the introduction of a new type of agents called: generator of replicas. The main goal of our work is, then, to develop a new system for intrusion detection (IDS) highly available, reliable and at the same time lightweight. Extensive simulations have been performed to study the feasibility and prove the optimality of the proposed approach.

Keywords: Agent replication, Fault tolerance, Intrusion detection, Ad-hoc networks

Distributed program execution control based on application global states monitoring in PEGASUS DA framework

Marek Tudruj^{1,2}, Damian Kopanski¹, Lukasz Masko²

¹Polish-Japanese Institute of Information Technology, Warsaw, Poland

²Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland
tudruj@pjwstk.edu.pl

This paper presents control implementation methods for an original distributed program design framework PEGASUS DA (Program Execution Governed by Asynchronous SUPervision of States in Distributed Applications) which provides automated design of distributed program execution control based on program global states monitoring. The framework includes a built in support for handling local and global application states as well as automatic construction and use of strongly consistent application global states for program execution control. In particular, the paper presents methods used to implement distributed program control inside the PEGASUS DA framework run on clusters of contemporary multicore processors based on multithreading. The program execution design methods are illustrated on some typical applications such as a distributed Branch and Bound algorithm and load balancing in distributed multicore applications.

Keywords: distributed program execution control, distributed program design, global application states, strongly consistent global states, program design tools

New Scalable SIMD-based Ray Caster Implementation for Virtual Machining

Torsten Welsch, Alexander Leutgeb, Michael Hava

RISC Software GmbH

Softwarepark 35

A-4232 Hagenberg, Austria

{alexander.leutgeb, torsten.welsch}@risc-software.at

michael.hava@risc-software.at

We present a highly efficient implementation of a ray casting system for the visualization of subtractive manufacturing, combining state-of-the-art results in a wide range of active research fields. Beside popular techniques like acceleration structures, coherent traversal and frustum culling, our object-space based implementation integrates the novel surface cell evaluation (SCE) algorithm. Thus, our ray caster allows both a non-approximate and an interactive visualization of tens of thousands of Boolean subtraction operations between a stock workpiece and arbitrary triangular swept volumes. Compared to image-space based approaches for virtual machining, such as z-maps, dexels or layered depth images, our scalable SIMD implementation offers a view-independent visualization as well as a higher rendering performance. Hence, it is perfectly suited for both simulation and verification applications in computer-aided manufacturing (CAM).

Keywords: Ray casting, SIMD, Boolean subtraction operations, Subtractive manufacturing, CAM visualization, Material processing, Multi-axis milling

Parallelization of Permuting Schema-less XML Compressors

Tyler Corbin¹, Tomasz Muldner¹, Jan Krzysztof Miziołek²

¹Jodrey School of Computer Science,
Acadia University, Wolfville, B4P 2A9 NS, Canada

²University of Warsaw, Warsaw, Poland

{094568c, tomasz.muldner}@acadiau.ca, jkm@al.uw.edu.pl

The verbose nature of XML results in large overheads in storage and network transfers, but existing research on designing XML compressors have not attempted to take advantage of parallel computing. This paper studies efficiency of four versions of complete schema-less single-pass parallel XML compressors, based on XSAQCT, an existing XML compressor. We provide results of our tests implemented for multi-core machines for all these versions, using a test suite designed to incorporate XML documents with various characteristics. Finally, these results are analyzed to determine the theoretical upper bounds given by Amdahl's law, the actual speedup, and the compression ratios.

Keywords: XML, XML compression, parallelization, Java

Parallel Processing Model for Syntactic Pattern Recognition-Based Electrical Load Forecast

Mariusz Flasiński, Janusz Jurek, Tomasz Peszek
Information Technology Systems Department, Jagiellonian University
prof. St. Łojasiewicza 4, Kraków 30-348, Poland
{m.flasinski, j.jurek}@mfpartners.pl

A model of a recognition of distorted/fuzzy patterns for a electrical load forecast is presented in the paper. The model is based on a syntactic pattern recognition approach. Since a system implemented on the basis of the model is to perform in a real-time mode, it is parallelized. An architecture for parallel processing and a method of tasks distribution is proposed and experimental results are discussed.

Keywords: syntactic pattern recognition, distorted/fuzzy patterns, GDPLL(k) grammar, parallel parser, electrical load forecast

Parallel event-driven simulation based on application global state monitoring

Lukasz Masko¹, Marek Tudruj^{1,2}

¹Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

²Polish-Japanese Institute of Information Technology, Warsaw, Poland

masko@ipipan.waw.pl, tudruj@pjwstk.edu.pl

Discrete event simulation is a well known technique used for modelling and simulating complex parallel systems. In this approach, each simulated entity reacts to the events, which are addressed to it, and as a result produces new events, which may be sent as messages to other system components. Parallel simulation introduces multiple event queues working in parallel. Multiple queues introduce parallelism which is limited by the synchronization of local clocks supervising parallel queues. A proper synchronization between parallel queues must be introduced. Global state monitoring is a natural way to organize global simulation state monitoring and control. Every queue process reports its progress state, being the timestamp of the most recently processed event, to a global synchronizer. Reporting is done asynchronously and has no influence on the simulation process. A global simulation state can be defined as the vector containing timestamps of the most recently processed event in every queue. The paper presents the principles of parallel simulation designed by the use of a system infrastructure for global states monitoring. Comparison to existing parallel simulation methods is provided.

Keywords: parallel discrete event simulation, distributed application global states, distributed programs design, distributed program design tools

It's Not a Bug, It's a Feature.
Wait-free Asynchronous Cellular Genetic Algorithm

Frédéric Pinel¹, Bernabé Dorronsoro², Pascal Bouvry¹, Samee U. Khan³

¹FSTC/CSC/ILIAS, University of Luxembourg, Luxembourg

²University of Lille, France

³Department of Electrical and Computer Engineering
North Dakota State University, USA

{frederic.pinel, pascal.bouvry}@uni.lu,
bernabe.dorronsoro diaz@inria.fr, samee.khan@ndsu.edu

In this paper, we simplify a Parallel Asynchronous Cellular Genetic Algorithm, by removing thread locks for shared memory access. This deliberate error aims to accelerate the algorithm, while preserving its search capability. Experiments with three benchmark problems show an acceleration, and even a slight improvement in search capability, with statistical significance.

Keywords: Cellular Genetic Algorithm, Parallelism, wait-free

Evolutionary algorithms for abstract planning

Jarosław Skaruz¹, Artur Niewiadomski¹, Wojciech Penczek^{1,2}

¹Siedlce University of Natural Sciences and Humanities

3-Maja 54, 08-110 Siedlce, Poland

²Institute of Computer Science, Polish Academy of Sciences

Jana Kazimierza 5, 01-248, Warsaw, Poland

jaroslaw.skaruz@uph.edu.pl, artur@ii.uph.edu.pl,

penczek@ipipan.waw.pl

The paper presents a new approach based on evolutionary algorithms to an abstract planning problem, which is the first stage of the web service composition problem. An abstract plan is defined as an equivalence class of sequences of service types that satisfy a user query. Two sequences are equivalent if they are composed of the same service types, but not necessarily occurring in the same order. The objective of our genetic algorithm (GA) is to return representatives of abstract plans without generating all the equivalent sequences. Experimental results are presented and compared with these obtained using an SMT-solver, showing that GA finds solutions for very large sets of service types in a reasonable and shorter time.

Keywords: genetic algorithms, web service composition, abstract planning

Solution of the Inverse Continuous Casting Problem with the Aid of Modified Harmony Search Algorithm

Edyta Hetmaniok, Damian Słota, Adam Zielonka
Institute of Mathematics
Silesian University of Technology
Kaszubska 23, 44-100 Gliwice, Poland
{edyta.hetmaniok, damian.slota, adam.zielonka}@polsl.pl

In the paper a description of procedure for solving the inverse problem of continuous casting is given. The problem consists in reconstruction of the cooling conditions of solidified ingot and is based on minimization of the appropriate functional by using the modified Harmony Search algorithm – the algorithm of artificial intelligence inspired by process of composing the jazz music.

Keywords: Artificial Intelligence, Harmony Search Algorithm, Inverse Continuous Casting Problem

Influence of a Topology of a Spring Network on its Ability to Learn Mechanical Behaviour

Maja Czoków¹, Jacek Miękiś²

¹Faculty of Mathematics and Computer Science

Nicolaus Copernicus University

Toruń, Poland

²Institute of Applied Mathematics and Mechanics

University of Warsaw

Warsaw, Poland

maja@mat.umk.pl, miekisz@mimuw.edu.pl

We discuss how the topology of the spring system/network affects its ability to learn a desired mechanical behaviour. To ensure such a behaviour, physical parameters of springs of the system are adjusted by an appropriate gradient descent learning algorithm. We find the betweenness centrality measure particularly convenient to describe topology of the spring system structure with the best mechanical properties. We apply our results to refine an algorithm generating the structure of a spring network. We also present numerical results confirming our statements.

Keywords: spring systems, learning systems, betweenness, rigid graphs

Evaluation of autparallelization toolkits for commodity graphics hardware

David Williams¹, Valeriu Codreanu¹, Po Yang², Baoquan Liu², Feng Dong², Burhan Yasar³, Babak Mahdian⁴, Alessandro Chiarini⁵, Xia Zhao⁶, Jos Roerdink¹

¹University of Groningen, The Netherlands

²University of Bedfordshire, United Kingdom

³RotaSoft, Turkey

⁴ImageMetry, Czech Republic

⁵Super Computing Solutions, Italy

⁶AnSmart, United Kingdom

d.p.williams@rug.nl, vali.codreanu@arh.pub.ro

In this paper we evaluate the performance of the OpenACC and Mint toolkits against C and CUDA implementations of the standard PolyBench test suite. Our analysis reveals that performance is similar in many cases, but that a certain set of code constructs impede the ability of Mint to generate optimal code. We then present some small improvements which we integrate into our own GPSME toolkit (which is derived from Mint) and show that our toolkit now out-performs OpenACC in the majority of tests.

Keywords: Evaluation, Autoparallelization, GPU Computing

Real-Time Multiview Human Body Tracking using GPU-accelerated PSO

Bogusław Rymut¹, Bogdan Kwolek²

¹Rzeszów University of Technology

W. Pola 2, 35-959 Rzeszów, Poland

²AGH University of Science and Technology

30 Mickiewicza Av., 30-059 Kraków, Poland

brymut@prz.edu.pl, bkw@agh.edu.pl

This paper presents our approach to 3D model-based human motion tracking using a GPU-accelerated particle swarm optimization. The tracking involves configuring the 3D human model in the pose described by each particle and then rendering it in each particle's 2D plane. In our implementation, we launch one independent thread for each column of each 2D plane. Such a parallel algorithm exhibits the level of parallelism that allows us to effectively utilize the GPU resources. Owing to such task decomposition the tracking of the full human body can be performed at rates of 15 frames per second. The GPU achieves an average speedup of 7.5 over the CPU. The speedup that achieves the GPU over CPU grows with the number of the particles. For marker-less motion capture system consisting of four calibrated and synchronized cameras, the computations were conducted on four CPU cores and four GTX GPUs on two cards.

Keywords: GPGPU, Dynamic optimization, Real-time image processing and analysis

Implementation of a heterogeneous image reconstruction system for clinical Magnetic Resonance

Grzegorz Kowalik¹, Jennifer Steeden¹, David Atkinson², Andrew Taylor¹,
Vivek Muthurangu¹

¹Institute of Cardiovascular Science
UCL Centre for Cardiovascular Imaging, London, UK

²University College London, UK

kowgrzegorz@gmail.com

This paper describes development of a novel online, heterogeneous image reconstruction system. The system integrates an external computer equipped with a GPU card into an MR scanner's image reconstruction pipeline. The system promotes fast online reconstruction for computationally intensive algorithms making them feasible in a busy clinical service. An on-going work on analyse and improvement of execution time of the SENSE algorithm as well as networking framework are presented. The imaging algorithm was broken down into distinctive steps for execution time profiling. Also, steps to achieve overlapping of execution and transmission are described. The GPU version reduced the algorithm's bottleneck from 81% to 40% of iteration's time. This allowed a full overlap between an acquisition and a reconstruction making a delay time between finishing of the acquisition and that of the reconstruction to be constant. The system was successfully used in research and clinical studies requiring high data throughput.

Keywords: GPGPU, CUDA, Medical imaging

X-ray laser imaging of biomolecules using multiple GPUs

Stefan Engblom, Jing Liu

Uppsala University, Sweden

{stefan.engblom, jing.liu}@it.uu.se

Extremely bright X-ray lasers are becoming a promising tool for 3D imaging of biomolecules. By hitting a beam of streaming particles with a very short burst of a high energy X-ray and collecting the resulting scattering pattern, the 3D structure of the particles can be deduced. The computational complexity associated with transforming the data thus collected into a 3D intensity map is very high and calls for efficient data-parallel implementations.

We present ongoing work in accelerating this application using multiple GPU nodes. In particular, we look at the scaling properties of the application and give predictions as to the computational viability of this imaging technique.

Keywords: GPU cluster, CUDA/MPI, single molecule imaging, X-ray laser

Out-of-Core Solution of Eigenproblems for Macromolecular Simulations on GPUs

José I. Aliaga¹, Davor Davidovic², Enrique S. Quintana-Ortí¹

¹Dpto. de Ingeniería y Ciencia de Computadores, Universidad Jaume I
12.071–Castellón, Spain

²Institut Ruder Bošković, Centar za Informatiku i Računarstvo - CIR
10000–Zagreb, Croatia

{aliaga, quintana}@icc.uji.es, ddavid@irb.hr

We consider the solution of large-scale eigenvalue problems that appear in the motion simulation of complex macromolecules on desktop platforms. To tackle the dimension of the matrices that are involved in these problems, we formulate out-of-core (OOC) variants of the two selected eigensolvers, that basically decouple the performance of the solver from the storage capacity. Furthermore, we contend with the high computational complexity of the solvers by off-loading the arithmetically-intensive parts of the algorithms to a hardware graphics accelerator.

Keywords: Macromolecular motion simulation, eigenvalue problems, out-of-core computing, multicore processors, GPUs

GPU implementation of the Monte-Carlo simulations of the extended Ginzburg–Landau mode

Piotr Bialas^{1,2}, Jakub Kowal¹, Adam Strzelecki¹

¹Faculty of Physics, Astronomy and Applied Computer Science
Jagiellonian University, ul. Reymonta 4, 30-059 Kraków, Poland

²Mark Kac Complex Systems Research Centre

Faculty of Physics, Astronomy and Applied Computer Science
Jagiellonian University, Reymonta 4, 30-059 Kraków, Poland

pbialas@th.if.uj.edu.pl, jakub.kowal@uj.edu.pl

adam.strzelecki@uj.edu.pl

In this contribution we describe an efficient GPU implementation of the Monte-Carlo simulation of the Ginzburg–Landau model. We achieve the performance close to 50% of the peak performance of the used GPU. We compare this performance with a parallelized and vectorized CPU code and discuss the observed differences.

Keywords: GPU computing, Monte-Carlo simulations, vectorisation, Ginzburg-Landau model

Using GPUs for parallel stencil computations in relativistic hydrodynamic simulation

Sebastian Cygert¹, Daniel Kikoła³, Joanna Porter-Sobieraj¹, Jan Sikorski²,
Marcin Słodkowski²

¹Faculty of Mathematics and Information Science
Warsaw University of Technology, Poland

²Faculty of Physics, Warsaw University of Technology
Koszykowa 75, 00-662 Warsaw, Poland

³Department of Physics, Purdue University
525 Northwestern Ave., West Lafayette, IN 47907, United States

cygerts@student.mini.pw.edu.pl, j.porter@mini.pw.edu.pl
{kikola, slodkow}@if.pw.edu.pl

This paper explores the possibilities of using a GPGPU for complex 3D finite difference computation. We propose a new approach to this topic using surface memory and compare it with 3D stencil computations carried out via shared memory, which is currently considered to be the best approach. The case study was performed for the extensive computation of collisions between heavy nuclei in terms of relativistic hydrodynamics. To provide a more detailed comparison between our approach and the one using shared memory we present tests for complex and simplified versions of the algorithm.

Keywords: GPU computing, CUDA, Parallel algorithms, Finite difference, MUSTA-FORCE algorithm, Riemann solver

PDNOC: An Efficient Partially Diagonal Network-on-Chip Design

Thomas Canhao Xu, Ville Leppänen, Pasi Liljeberg, Juha Plosila, Hannu Tenhunen
Department of Information Technology, University of Turku
20014, Turku, Finland
canxu@utu.fi

With the constantly increasing of number of cores in multicore processors, more emphasis should be paid to the on-chip interconnect. Performance and power consumption of an on-chip interconnect are directly affected by the network topology. The efficiency can also be optimized by proper mapping of applications. Therefore in this paper, we propose a novel Partially Diagonal Network-on-Chip (PDNOC) design that takes advantage of both heterogeneous network topology and congestion-aware application mapping. We analyse the partially diagonal network in terms of area usage, power consumption, routing algorithm and implementation complexity. The key insight that enables the PDNOC is that most communication patterns in real-world applications are hot-spot and bursty. We implement a full system simulation environment using SPLASH-2 benchmarks. Evaluation results shown that, the proposed PDNOC provides up to 25% improvement in execution time over concentrated mesh, and 3.6x better energy delay product over fully connected diagonal network.

Keywords: Multicore, Heterogeneous, Network-on-Chip, Network topology, Mapping

Adaptive Fork-Heuristics for Software Thread-Level Speculation

Zhen Cao, Clark Verbrugge
School of Computer Science, McGill University
Montréal, Québec, Canada
zhen.cao@mail.mcgill.ca, clump@cs.mcgill.ca

Fork-heuristics play a key role in software Thread-Level Speculation (TLS). Current fork-heuristics either lack real parallel execution environment information to accurately evaluate fork points and/or focus on hardware-TLS implementation which cannot be directly applied to software TLS. This paper proposes adaptive fork-heuristics as well as a feedback-based selection technique to overcome the problems. Adaptive fork-heuristics insert and speculate on all potential fork/join points and purely rely on the runtime system to disable inappropriate ones. Feedback-based selection produces parallelized programs with ideal speedups using log files generated by adaptive heuristics. Experiments of three scientific computing benchmarks on a 64-core machine show that feedback-based selection and adaptive heuristics achieve more than 88% and 50% speedups of the manual-parallel version, respectively. For the Barnes-Hut benchmark, feedback-based selection is 49% faster than the manual-parallel version.

Keywords: software thread-level speculation, fork heuristics, automatic parallelization, performance tuning

Inexact sparse matrix vector multiplication in Krylov subspace methods: An application-oriented reduction method

Ahmad Mansour, Jürgen Götze
Information Processing Lab
TU Dortmund
Otto-Hahn-Str. 4, Dortmund 44227, Germany
ahmad.mansour@tu-dortmund.de

Iterative solvers based on Krylov subspace method proved to be robust in the presence of well monitored inexact matrix vector products. In this paper, we show that the iterative solver performs well while gradually reducing the number of nonzero elements of the matrix throughout the iterations. We benefit from this robustness in reducing the computational effort and the communication volume when implementing sparse matrix vector multiplication (SpMV) on Network-on-Chip (NoC).

Keywords: Krylov subspace method, Network-on-Chip, Sparse matrix vector multiplication

The regular expression matching algorithm for the energy efficient reconfigurable SoC

Paweł Russek, Kazimierz Wiatr
AGH University of Science and Technology
al. Mickiewicza 30, 30-059, Kraków, Poland
{russek, wiatr}@agh.edu.pl

The paper presents an algorithm for a regular expressions pattern matching system. The goal was to achieve an attractive performance and a low energy consumption. The proposed scheme is particularly useful when a big set of complex regular expression patterns must be inspected in parallel (e.g. in computer malware and anti-virus systems). The idea of the algorithm derives from a concept of the Bloom filter algorithm. The Bloom filter operation is used to inspect an incoming data to find static sub-patterns of regular expressions. When the Bloom filter reports a match, a closer inspection is performed. The Bloom filtering is done by a hardware dedicated co-processor. The regular expressions' wildcard matching part is executed by a CPU. Above concept was implemented and tested on Xilinx ZynqTM-7000 All Programmable SoC platform. The results and performance for regular expressions patterns from ClamAV virus database is given.

Keywords: regular expressions, matching energy-efficient systems, custom architectures, HW-SW Co-Design

Transparent application acceleration by intelligent scheduling of shared library calls on heterogeneous systems

João Colaço, Adrian Matoga, Aleksandar Ilić, Nuno Roma, Pedro Tomás,
Ricardo Chaves
INESC-ID / IST
Rua Alves Redol, 9, 1000-029 Lisboa, Portugal
ilic@sips.inesc-id.pt
{nuno.roma, pedro.tomas, Ricardo.Chaves}@inesc-id.pt

Transparent application acceleration in heterogeneous systems can be performed by automatically intercepting shared libraries calls and by efficiently orchestrating the execution across all processing devices. To fully exploit the available computing power, the intercepted calls must be replaced with faster accelerator-based implementations and intelligent scheduling algorithms must be incorporated. When compared with previous approaches, the framework proposed herein does not only transparently intercepts and redirects the library calls, but it also incorporates state-of-art scheduling algorithms, for both divisible and indivisible applications. When compared with highly optimized implementations for multi-core CPUs (e.g., MKL and FFTW), the obtained experimental results demonstrate that, by applying appropriate light-weight scheduling and load-balancing mechanisms, performance speedups as high as 7.86 (matrix multiplication) and 4.6 (FFT) can be obtained.

Keywords: Transparent acceleration, Heterogeneous computing, Automatic scheduling, Load balancing

Improving Parallel I/O Performance Using Multithreaded Two-Phase I/O with Processor Affinity Management

Yuichi Tsujita^{1,2}, Kazumi Yoshinaga^{1,2}, Atsushi Hori^{1,2}, Mikiko Sato^{3,2},
Mitaro Namiki^{3,2}, Yutaka Ishikawa⁴

¹Advanced Institute for Computational Science, RIKEN, Japan

²JST CREST, Japan

³Tokyo University of Agriculture and Technology, Japan

⁴The University of Tokyo, Japan

yuichi_tsujita@fw.ipsj.or.jp

I/O operation is one of the performance bottlenecks in parallel computing, especially in data-intensive one. I/O performance improvement by using parallel I/O APIs is one of the effective approaches in order to improve parallel computing performance. MPI-IO, which is a parallel I/O interface in the MPI standard, supports a variety of parallel I/O access patterns. An MPI-IO implementation named ROMIO has several performance optimization mechanisms such as two-phase I/O for parallel I/O with non-contiguous access patterns. Although advanced application users are likely to incorporate MPI-IO in their MPI programs, it is still difficult for most of standard users. Application-oriented portable I/O APIs such as HDF5 have been proposed and used, where such I/O interfaces hide complexity of MPI-IO handling from users. Since their parallel I/O interface libraries are built on an MPI-IO interface, performance degradation in an MPI-IO implementation affects performance of an upper layer interface. We have implemented our optimized two-phase I/O in the ROMIO library used in an HDF5 library. Furthermore we have extended it to be able to control CPU core binding from users' MPI processes as a preliminary implementation. We show performance advantages of the optimized scheme compared with the original implementation.

Keywords: MPI-IO, HDF5, two-phase I/O, multithreaded I/O, processor affinity management

Storage Systems for Organizationally Distributed Environments - PLGrid PLUS Case Study

Renata Słota¹, Łukasz Dutka², Bartosz Kryza², Darin Nikolow¹, Dariusz Król²,
Michał Wrzeszcz², Jacek Kitowski^{1,2}

¹AGH University of Science and Technology

Faculty of Computer Science, Electronics and Telecommunications

Department of Computer Science

al. Mickiewicza 30, 30-059 Kraków, Poland

²AGH University of Science and Technology, ACC Cyfronet AGH

ul. Nawojki 11, 30-950 Kraków, Poland

{rena,dutka,bkryza,darin,dkrol,wrzeszcz,kito}@agh.edu.pl

With the increasing amount of data, the research community is facing problems with methods of effectively accessing, storing, and processing data in large scale and geographically distributed environments. This paper addresses major data management issues, in particular use cases and scenarios (on the basis of Polish research community organized around the PLGrid PLUS Project) and discusses architectures of data storage systems available in both PL-Grid and other similar federated environments. On that basis, a concept of a new meta storage system, named VeilFS, is presented. The proposed system unifies file access methods for geographically distributed large scale systems and hides complexity of data access and management in such environments.

Keywords: storage systems, organizationally distributed environments, data access and management

The High Performance Internet of Things: using GVirtuS for gluing cloud computing and ubiquitous connected devices

Raffaele Montella¹, Giuliano Laccetti²

¹Department of Applied Science

University of Napoli Parthenope, Italy

²Department of Mathematics and Applications

University of Napoli Federico II, Italy

raffaele.montella@uniparthenope.it, giuliano.laccetti@unina.it

The availability of computing resources and the need for high quality services are rapidly evolving the vision about the acceleration of knowledge development, improvement and dissemination. During the so-called web 1.0, the Internet of Hyperlinks, HPC and grid computing technologies rose. According with the common shared knowledge, we are living the final part of the web 2.0, the Internet of Social, powered by the elastic cloud computing technology. The Internet of Things is growing up: the developers and then the users have the power to integrate computation with real stuff control. The high performance cloud computing is behind the scene powering “the next big thing”. Using the GVirtuS general purpose virtualisation service a model for internet of things developing architecture is proposed demonstrating different scenarios and application test beds.

Keywords: HPC, Cloud Computing, Internet of Thinks, Virtualisation, GVirtuS

Performance Evaluation of Sparse Matrix Multiplication Kernels on Intel Xeon Phi

Erik Saule¹, Kamer Kaya¹, Ümit Çatalyürek^{1,2}

¹The Ohio State University, Dept. Biomedical Informatics, USA

²The Ohio State University, Dept. Electrical and Computer Engineering, USA

{esaule, kamer, umit}@bmi.osu.edu

Intel Xeon Phi is a recently released high-performance coprocessor which features 61 cores each supporting 4 hardware threads with 512-bit wide SIMD registers achieving a peak theoretical performance of 1Tflop/s in double precision. Its design differs from classical modern processors; it comes with a large number of cores, the 4-way hyperthreading capability allows many applications to saturate the massive memory bandwidth, and its large SIMD capabilities allow to reach high computation throughput. The core of many scientific applications involves the multiplication of a large, sparse matrix with a single or multiple dense vectors which are not compute-bound but memory-bound. In this paper, we investigate the performance of the Xeon Phi coprocessor for these sparse linear algebra kernels. We highlight the important hardware details and show that Xeon Phi's sparse kernel performance is very promising and even better than that of cutting-edge CPUs and GPUs.

Keywords: Intel Xeon Phi, sparse matrix vector multiplication, sparse matrix matrix multiplication, many-core architectures, experimental evaluation

Portable HPC Programming on Intel Many-Integrated-Core Hardware with MAGMA Port to Xeon Phi

Jack Dongarra^{1,2,3}, Mark Gates¹, Azzam Haidar¹, Yulu Jia¹, Khairul Kabir¹,
Piotr Luszczek¹, Stan Tomov¹

¹University of Tennessee Knoxville, USA

²Oak Ridge National Laboratory, USA

³University of Manchester, UK

haidar@eecs.utk.edu

This paper presents the design and implementation of several fundamental dense linear algebra (DLA) algorithms for multicore with Intel Xeon Phi Coprocessors. In particular, we consider algorithms for solving linear systems. Further, we give an overview of the MAGMA MIC library, an open source, high performance library that incorporates the developments presented, and in general provides to heterogeneous architectures of multicore with coprocessors the DLA functionality of the popular LAPACK library. The LAPACK-compliance simplifies the use of the MAGMA MIC library in applications, while providing them with portably performant DLA. High performance is obtained through use of the high-performance BLAS, hardware-specific tuning, and a hybridization methodology where we split the algorithm into computational tasks of various granularities. Execution of those tasks is properly scheduled over the heterogeneous hardware components by minimizing data movements and mapping algorithmic requirements to the architectural strengths of the various heterogeneous hardware components. Our methodology and programming techniques are incorporated into the MAGMA MIC API, which abstracts the application developer from the specifics of the Xeon Phi architecture and is therefore applicable to algorithms beyond the scope of DLA.

Keywords: Xeon Phi, hybrid algorithms, heterogeneous programming, high performance computing, dense linear algebra, accelerators techniques

Accelerating a Massively Parallel Numerical Simulation in Electromagnetism using a Cluster of GPUs

Cédric Augonnet, David Goudin, Agnès Pujols, Muriel Sesques
CEA/CESTA, 15 Avenue des Sablières, CS 600001, 33116 Le Barp, France
`cedric.augonnet@cea.fr`

We have accelerated a legacy massively parallel code solving 3D Maxwell's equations on a hybrid cluster enhanced with GPUs. To minimize the impact on our existing code, we combine its original Full- MPI approach with task parallelism to design an efficient accelerated LL' solver that efficiently shares the same GPUs between different processes and relies on an optimized communication patterns. On 180 nodes of the Tera100 cluster, our GPU-accelerated LL' decomposition reaches 66.8 TFlop/s on a problem with 363388 unknowns, whereas the sustained machine's CPU and GPU peaks are respectively 7.2 and 76.3 TFlop/s.

Keywords: GPU, Dense linear algebra, Cluster computing, Application, Electromagnetism

Multidimensional Monte Carlo Integration on Clusters with Hybrid GPU-accelerated Nodes

Dominik Szalkowski, Przemysław Stpoczyński
Institute of Mathematics, Maria Curie–Skłodowska University
Pl. Marii Curie–Skłodowskiej 1, 20-031 Lublin, Poland
{dominisz,przem}@hektor.umcs.lublin.pl

The aim of this paper is to show that the multidimensional Monte Carlo integration can be efficiently implemented on clusters with hybrid GPU-accelerated nodes using recently developed parallel versions of LCG and LFG pseudorandom number generators. We explain how to utilize multiple GPUs and all available cores of CPUs within a single node and how to extend computations on all available nodes of a cluster using MPI. The results of experiments performed on a Tesla-based GPU cluster are also presented and discussed.

Keywords: multidimensional integration, Monte Carlo methods, parallelized pseudorandom number generators, GPU clusters

Efficient Execution of Erasure Codes on AMD APU Architecture

Roman Wyrzykowski, Marcin Woźniak, Lukasz Kuczynski
Institute of Computer and Information Science
Czestochowa University of Technology
Dabrowskiego 69, 42-201 Czestochowa, Poland
{roman, marcin.wozniak, lkucz}@icis.pcz.pl

Erasure codes such as Reed-Solomon codes can improve the availability of distributed storage in comparison with replication systems. In previous studies we investigated implementation of these codes on multi/many-core architectures, such as Cell/B.E. and GPUs. In particular, it was shown that bandwidth of PCIe bus is a bottleneck for the implementation on GPUs.

In this paper, we focus on investigation how to map systematically the Reed-Solomon and Cauchy Reed-Solomon erasure codes onto the AMD Accelerated Processing Unit (APU), a new heterogeneous multi/many-core architecture. This architecture combines CPU and GPU in a single chip, eliminating costly transfers between them through the PCI bus. Moreover, APU processors combine some features of Cell/B.E. processors and many-core GPUs, allowing both vectorization and SIMT processing simultaneously.

Based on the previous works the method for the systematic mapping of computation kernels of Reed-Solomon and Cauchy Reed-Solomon algorithms onto the AMD APU architecture is proposed, taking into account properties of the architecture on all the levels of its parallel processing hierarchy.

Keywords: Erasure codes, Reed-Solomon codes, multicore architectures, GPU, APU

AVX acceleration of DD arithmetic between a sparse matrix and vector

Toshiaki Hishinuma¹, Akihiro Fujii¹, Teruo Tanaka¹, Hidehiko Hasegawa²

¹Major of Informatics, Kogakuin University, Tokyo, Japan

²Faculty of Lib., Info. & Media Sci., University of Tsukuba, Tsukuba, Japan

em13015@ns.kogakuin.ac.jp

High precision arithmetic can improve the convergence of Krylov subspace methods; however, it is very costly. One system of high precision arithmetic is double-double (DD) arithmetic, which uses more than 20 double precision operations for one DD operation. We accelerated DD arithmetic using AVX SIMD instructions. The performances of vector operations in 4 threads are 51-59% of peak performance in a cache and bounded by the memory access speed out of the cache. For SpMV, we used a double precision sparse matrix A and DD vector x to reduce memory access and achieved performances of 19-46% of peak performance using padding in execution. We also achieved performances that were 9-33% of peak performance for a transposed SpMV. For these cases, the performances were not bounded by memory access.

Keywords: Double-double arithmetic, AVX, SpMV, high precision

Using Quadruple Precision Arithmetic to Accelerate Krylov Subspace Methods on GPUs

Daichi Mukunoki^{1,3}, Daisuke Takahashi²

¹Graduate School of Systems and Information Engineering
University of Tsukuba, Japan

²Faculty of Engineering, Information and Systems, University of Tsukuba
1-1-1 Tennodai, Tsukuba, Ibaraki 305-8573, Japan

³Research Fellow of the Japan Society for the Promotion of Science
mukunoki@hpcs.cs.tsukuba.ac.jp, daisuke@cs.tsukuba.ac.jp

The convergence of the Krylov subspace methods is affected by round-off errors. The number of iterations until convergence may be decreased by reducing round-off errors through the use of quadruple precision arithmetic instead of double precision. We implemented the CG and BiCGStab methods using quadruple precision arithmetic and compared the performance with the standard double precision implementations on an NVIDIA Tesla K20X GPU. Our results show that in some cases our implementations using quadruple precision arithmetic outperform the double precision versions. We will show that quadruple precision arithmetic is not costly for the CG and BiCGStab methods on GPUs and the use of quadruple precision arithmetic may be a more effective alternative to the use of preconditioning.

Keywords: Krylov subspace method, CG method, BiCGStab method, Quadruple precision, GPU

Effectiveness of sparse data structure for double-double and quad-double arithmetics

Tsubasa Saito¹, Satoko Kikkawa¹, Emiko Ishiwata², Hidehiko Hasegawa³

¹Graduate School of Science, Tokyo University of Science, Japan

²Tokyo University of Science, Japan

³University of Tsukuba, Japan

hasegawa@slis.tsukuba.ac.jp

Double-double and quad-double arithmetics are effective tools to reduce the round-off errors in floating-point arithmetic. However, the dense data structure for high-precision numbers in MuPAT/Scilab requires large amounts of memory and a great deal of the computation time. We implemented sparse data types ddsp and qdsp for double-double and quad-double numbers. We showed that sparse data structure for high-precision arithmetic is practically useful for solving a system of ill-conditioned linear equation to improve the convergence and obtain the accurate result in smaller computation time.

Keywords: ill-conditioned matrix problem, sparse matrix, multiple precisions

Efficient heuristic adaptive quadrature on GPUs: Design and Evaluation

Daniel Thuerck¹, Sven Widmer², Arjan Kuijper^{1,3}, Michael Goesele¹

¹TU Darmstadt, Germany

²Graduate School Computational Engineering, TU Darmstadt, Germany

³Fraunhofer IGD, Germany

daniel.thuerck@googlemail.com, arjan.kuijper@igd.fraunhofer.de,
michael.goesele@gris.informatik.tu-darmstadt.de

Numerical integration is a common sub-problem in many applications. It can be solved easily in CPU-based applications using adaptive quadrature such as the adaptive Simpson's rule. These algorithms rely, however, on error estimation yielding a significant computational overhead. In addition, they require recursive function evaluations, which are not well suited for parallel computation on graphics processing units (GPUs) due to warp divergence issues. In this paper, we introduce heuristic forward quadrature as an alternative that is not only more efficient than traditional methods, but also better suited for accelerated massively-parallel calculation on GPUs. Additionally, we will give an error estimate for our method and demonstrate performance results for 1D and 2D integral applications which show that the algorithm leverages quadrature for the efficient implementation on GPUs.

Keywords: Quadrature, Numerical Algorithms, GPGPU, Massively-Parallel Algorithms

Square Block Code for Positive Definite Symmetric Cholesky Band Routines

Fred G. Gustavson¹, José R. Herrero², Enric Morancho²

¹IBM T.J. Watson Research Center, Emeritus, and Umeå University

²Universitat Politècnica de Catalunya, BarcelonaTech

Barcelona, Spain

fg2935@hotmail.com, {josepr,enricm}@ac.upc.edu

We describe a minor data format change for storing a symmetric band matrix AB using the same array space specified by LAPACK [1]. In LAPACK lower band codes $AB_{i,j}$ is referenced in its code as $AB_{i-j+1,j}$. This makes the code less readable than it could be as one would like to reference the (i,j) element of a matrix AB as $AB_{i,j}$. Furthermore, the layout of lower AB in the LAPACK's user Guide, page 142 of [1] shows the user a rectangular matrix with the diagonal of AB residing in the first row. Clearly, a layout description where the diagonal of AB resides on the main diagonal of AB , see again page 142 of [1], is more suggestive and other things being equal is preferable.

This contribution describes a Square Block, SB, implementation of LAPACK PD symmetric lower band code. This is possible by rearranging “in place” LAPACK Band Layout to become a SB layout. The floating point operations we perform are identical to the current LAPACK lower band code operations. Hence, this new code produces identical results with the current LAPACK band code. In the new code the (i,j) element of AB is represented as $AB_{i,j}$. Also, in the layout description of AB the diagonal of AB is depicted as laying on the diagonal of AB . The new layout for uplo = ‘L’ consists of two geometric figures; a parallelogram and a lower isosceles triangle T of side equal to kd . The parallelogram is partitioned into slabs of width nb . Each slab of P is also a parallelogram P_i of size $kd + 1$ by nb . P_i consists of two isosceles triangles of sizes nb and $nb - 1$ and a rectangle R_i of size $kd + 1 - nb$ by nb . Now the two triangles concatenate to form a SB of order nb . Hence, P_i also consists of just a SB and R_i . By transposing R_i in-place R_i becomes $(kd + 1)/nb - 1$ SB's plus a leftover rectangular block. We note that transposing AB gives an uplo = ‘U’ LAPACK implementation starting from the uplo = ‘L’ implementation. Thus, to get our SB formulation we follow this procedure. Triangle T now becomes an upper isosceles triangle. We also map T into upper blocked packed format [2] so it becomes “compatible” with the transposed parallelogram P .

This contribution proves the correctness of the SB implementation of AB fac-

torization. It also gives preliminary performance results.

Keywords: New data storage format, Banded Cholesky factorization, Square Block format

Exploiting Data Sparsity in Parallel Matrix Powers Computations

Nicholas Knight, Erin Carson, James Demmel
University of California, Berkeley, USA
{knight, ecc2z, demmel}@cs.berkeley.edu

We derive a new parallel communication-avoiding matrix powers algorithm for matrices of the form $A = D + USV^H$, where D is sparse and USV^H has low rank and is possibly dense. We demonstrate that, with respect to the cost of computing k sparse matrix-vector multiplications, our algorithm asymptotically reduces the parallel latency by a factor of $O(k)$ for small additional bandwidth and computation costs. Using problems from real-world applications, our performance model predicts up to $24\times$ speedups on petascale machines.

Keywords: communication-avoiding, matrix powers, power-law graphs, hierarchical matrices, blocking covers

Communication Avoiding ILU0 Preconditioner

Laura Grigori, Sophie Moufawad
INRIA France
`laura.grigori@inria.fr`

Keywords:

Parallel Design and Performance of Nested Filtering Factorization Preconditioner

Laura Grigori¹, Frédéric Nataf², Long Qu¹

¹INRIA, France

²Université Pierre et Marie Curie, Paris, France

`laura.grigori@inria.fr`

Keywords:

Hiding global communication latency in the GMRES algorithm on massively parallel machines

Wim Vanroose¹, Pieter Ghysels², Karl Meerbergen³, Tom Ashby⁴

¹Universiteit Antwerpen, Belgium

²Lawrence Berkeley National Laboratory, USA

³Katholieke Universiteit Leuven, Belgium

⁴IMEC, Belgium

Keywords:

A Domain Decomposition Method for Discretization of Multiscale Elliptic Problems by Discontinuous Galerkin Method

Max Dryja
Univeristy of Warsaw, Department of Mathematics
Banacha 2, 02-097 Warsaw, Poland
dryja@mimuw.edu.pl

In this talk boundary value problems for second order elliptic equations with highly discontinuous coefficients are considered. The problems are discretized by a discontinuous Galerkin (DG) with finite element method (FEM) on triangular elements using piecewise linear functions.

The goal is to design and analyze a parallel algorithm for solving the discrete problem whose rate of convergence is independent of the jumps of the coefficient. The method discussed is an additive Schwarz method (ASM) which belongs to a class of domain decomposition methods and is one of the most efficient parallel algorithm for solving discretizations of PDEs.

It turns out that the convergence of the method presented here is almost optimal and only weakly depends on the jumps of coefficients. The suggested method is very well suited for parallel computations.

Keywords: interior penalty method, discontinuous Galerkin method, elliptic equations with discontinuous coefficients, finite element method, additive Schwarz method

Parallel preconditioner for Finite Volume Element discretization of elliptic problem

Leszek Marcinkowski¹, Talal Rahman²

¹Faculty of Mathematics, University of Warsaw

Banacha 2, 02-097 Warszawa, Poland

²Faculty of Engineering, Bergen University College

Nygårdsgaten 112, 5020 Bergen, Norway

`l.marcinkowski@mimuw.edu.pl`, `talal.rahman@hib.no`

In this paper we present a parallel preconditioner for standard Finite Volume Element (FVE) discretization of elliptic problem with standard FEM continuous piecewise linear function space. The proposed preconditioner is constructed using an abstract framework of Additive Schwarz Method and is fully parallel. The convergence rate of GMRES method with our preconditioner is almost optimal i.e. it depends poly-logarithmically on the mesh coefficients.

Keywords: Finite Volume Element, Parallel Preconditioner, Domain Decomposition Method, Additive Schwarz Method

Abstract Schwarz Method for Nonsymmetric Local Discontinuous Galerkin Discretization of Elliptic Problem

Filip Klawe

Institute of Applied Mathematics and Mechanics, University of Warsaw

Banacha 2, 02-097 Warszawa, Poland

fzklawe@mimuw.edu.pl

The subject of our research is to show the result of numerical test regarding two-level additive Schwarz method for non-symmetric, elliptic problem in two dimensions with the use of discretization by local discontinuous Galerkin method (LDG). To construct the preconditioner, we use the domain decomposition method. Condition of the preconditioned system does not depend on the size of fine and coarse meshes, but only on the ratio of the coarse mesh size H and the overlap measure δ .

Keywords: Parallel preconditioner, local discontinuous Galerkin, convection-diffusion

Fast Numerical Method for 2D Initial-Boundary Value Problems for the Boltzmann Equation

Alexei Heintz¹, Piotr Kowalczyk²

¹Department of Mathematics, Chalmers University of Technology
and Göteborg University, SE-412 96, Göteborg, Sweden

²Institute of Applied Mathematics and Mechanics, University of Warsaw
Banacha 2, 02-097 Warszawa, Poland
heintz@chalmers.se, pkowal@mimuw.edu.pl

We present a new numerical scheme for the initial-boundary value problem for the Boltzmann equation in two-dimensional physical space. It is based on a splitting procedure in which the collision equation is solved using the adaptive algorithm for the computation of the full three-dimensional Boltzmann collision operator on non-uniform velocity grids introduced in the previous paper by the authors. The computation of the collision operator is performed in parallel for every physical grid cell. For the two-dimensional transport equation we use a second order finite volume method. The numerical example showing the effectiveness of our method is given.

Keywords: Boltzmann equation, numerical methods, non-uniform grids

Simulating phase transition dynamics on nontrivial domains

Maria Gokieli, Łukasz Bolikowski
Interdisciplinary Centre for Mathematical and Computational Modelling
University of Warsaw
Prosta 69, 00-838 Warsaw, Poland
{l.bolikowski, m.gokieli}@icm.edu.pl

Our goal is to investigate the influence of the geometry and topology of the domain on the solutions of the phase transition and other diffusion-driven phenomena in domain, modeled e.g. by the Allen-Cahn, Cahn-Hilliard, reaction-diffusion equations. We present FEM numerical schemes for the Allen-Cahn and Cahn-Hilliard equation based on the Eyre's algorithm and present some numerical results on split and dumbbell domains.

Keywords: diffusion, steady states, stability, free element, gradient system

Variable block multilevel iterative solution of general sparse linear systems

Bruno Carpentieri¹, Jia Liao¹, Masha Sosonkina²

¹Institute of Mathematics and Computing Science

University of Groningen, 9747 AG Groningen, The Netherlands

²Department of Modeling, Simulation & Visualization Engineering

Old Dominion University, Norfolk, VA 23529, USA

{b.carpentieri, j.liao}@rug.nl, msosonki@odu.edu

We present numerical results with a variable block multilevel incomplete LU factorization preconditioners for solving sparse linear systems arising, e.g., from the discretization of 2D and 3D partial differential equations on unstructured meshes. The proposed method automatically detects and exploits any available block structure in the matrix to maximize computational efficiency. Both sequential and parallel experiments are shown on selected matrix problems in different application areas, also against other standard preconditioners.

Keywords: linear systems, sparse matrices, Krylov methods, algebraic preconditioners, multilevel incomplete LU factorization, graph compression.

An automatic way of finding optimal elimination trees for sequential and parallel multi-frontal direct solver for adaptive finite element method

Hassan AbouEisha¹, Piotr Gurgul², Anna Paszyńska³, Maciej Paszyński², Mikhail Moshkov¹, Krzysztof Kuźnik²

¹King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

²AGH University of Science and Technology, Krakow, Poland

³Jagiellonian University, Krakow, Poland

paszynsk@agh.edu.pl

In this paper we present a dynamic programming algorithm for finding optimal elimination trees for the multi-frontal direct solver algorithm executed over two dimensional meshes with point singularities. The elimination tree found by the optimization algorithm results in a linear computational cost of sequential direct solver. Based on the optimal elimination tree found by the optimization algorithm we construct heuristic sequential multi-frontal direct solver algorithm resulting in a linear computational cost as well as heuristic parallel multi-frontal direct solver algorithm resulting in a logarithmic computational cost. The resulting parallel algorithm is implemented on NVIDIA CUDA GPU architecture based on our graph-grammar based approach.

Keywords: parallel multi-frontal direct solver, elimination tree, dynamic programming, adaptive finite element method, graph grammar

Parallel efficiency of an adaptive, dynamically balanced flow solver

Stanisław Gepner, Jerzy Majewski, Jacek Rokicki
Warsaw University of Technology
The Institute of Aeronautics and Applied Mechanics
Nowowiejska 24, 00-665 Warsaw, Poland
{sgepner, jmajewski, jack}@meil.pw.edu.pl

Computations in Fluid Dynamics require minimisation of time in which the result could be obtained. While parallel techniques allow for handling of large problems, it is the adaptivity that ensures that computational effort is focused on interesting regions in time and space. Parallel efficiency, in a domain decomposition based approach, strongly depends on partitioning quality. For adaptive simulation partitioning quality is lost due to the dynamic modification of the computational mesh. Maintaining high efficiency of parallelization requires rebalancing of the numerical load. This paper presents performance results of an adaptive and dynamically balanced in-house flow solver. The results prove that the rebalancing technique is a remedy to the adverse effects of adaptivity on overall performance.

Keywords: adaptation, dynamic load balancing, mesh refinement, parallel efficiency

Modification of the Newton's method for the simulations of gallium nitride semiconductor devices

Konrad Sakowski^{1,2}, Leszek Marcinkowski², Stanislaw Krukowski^{1,3}

¹Institute of High Pressure Physics, Polish Academy of Sciences
ul. Sokolowska 29/37, 01-142 Warsaw, Poland

²Faculty of Mathematics, University of Warsaw
Banacha 2, 02-097 Warsaw, Poland

³Interdisciplinary Centre for Mathematical and Computational Modelling
Warsaw University, ul. Pawinskiego 5a, 02-106 Warsaw, Poland
konrad@unipress.waw.pl

In this paper we present an application of the Newton's method for simulation of gallium nitride semiconductor devices in the steady state. Our simulations cover basic characteristics of devices, which are difficult to measure physically, such as distribution of the electron and hole concentrations, electrostatic potential, quasi-Fermi levels, recombination rates, etc., as well as the current-voltage characteristics, which can be easily measured and compared with. The drift-diffusion model of carrier transport in the semiconductor material is used. It consists of three nonlinear elliptic differential equations. In this paper we present a back-tracking strategy for the coupled Newton's method, which takes into account the specific nature of the drift-diffusion equations and improves convergence of the method.

Keywords: drift-diffusion, van Roosbroeck equations, gallium nitride, coupled Newton method

A project of numerical realization of the one-dimensional model of burning methanol

Krzysztof Moszyński
University of Warsaw, Poland
kmoszyns@mimuw.edu.pl

This project is based mainly on the so called time-splitting method and will be realized on a cluster. Our goal is to make this program as fast as possible.

Keywords: evolution partial differential-integral equation, time splitting method, linear transport, diffusion, kinetics of chemical reaction, mixers

Scheduling Bag-of-Tasks Applications to Optimize Computation Time and Cost

Anastasia Grekoti, Natalia V. Shakhlevich
School of Computing, University of Leeds
Leeds LS2 9JT, UK
{scag, n.shakhlevich}@leeds.ac.uk

Bag-of-tasks applications consist of independent tasks that can be performed in parallel. Although such problems are well known in classical scheduling theory, the distinctive feature of Grid and cloud applications is the importance of the cost factor: in addition to the traditional scheduling criterion of minimizing computation time, in Grids and clouds it is also important to minimize the cost of using resources. We study the structural properties of the time/cost model and explore how the existing scheduling techniques can be extended to handle the additional cost criterion. Due to the dynamic nature of distributed systems, one of the major requirements to scheduling algorithms is related to their speed. The heuristics we propose are fast and, as we show in our experiments, they compare favourably with the existing Grid scheduling algorithms.

Keywords: scheduling, bag-of-tasks applications, heuristics

Scheduling Moldable Tasks with Precedence Constraints and Arbitrary Speedup Functions on Multiprocessors

Sascha Hunold
Research Group Parallel Computing
Vienna University of Technology, Austria
hunold@par.tuwien.ac.at

Due to the increasing number of cores of current machines, the question arises to which cores parallel tasks should be mapped. Thus, parallel task scheduling is now more relevant than ever, especially under the moldable task model, in which tasks are allocated a fixed number of processors before execution. A common assumption for scheduling algorithms is that the speedup function of moldable tasks is either non-decreasing, sub-linear or concave. However, in practice the resulting speedup of parallel programs on current hardware with deep memory hierarchies is most often neither non-decreasing nor concave. We present a new algorithm for the problem of scheduling moldable tasks with precedence constraints for the makespan objective and for arbitrary speedup functions. We show through simulation that the algorithm not only creates competitive schedules for moldable tasks with arbitrary speedup functions, but also outperforms other published heuristics and approximation algorithms for non-decreasing speedup functions.

Keywords: scheduling, heuristics vs. approximation algorithms, moldable task model, directed acyclic graphs

OStrich: Fair Scheduling for Multiple Submissions

Joseph Emeras¹, Vinicius Pinheiro², Krzysztof Rządca³, Denis Trystram⁴

¹Laboratoire d'Informatique de Grenoble CEA - CNRS, France

²Lab. for Parallel and Distributed Computing University of São Paulo, Brasil

³Institute of Informatics, University of Warsaw, Poland

⁴Grenoble Institute of Technology and Institut Universitaire de France

joseph.emeras@imag.fr, vinicius.pinheiro@ime.usp.br,
krzadca@mimuw.edu.pl, trystram@imag.fr

Campaign Scheduling is characterized by multiple job submissions issued from multiple users over time. This model perfectly suits today's systems since most available parallel environments have multiple users sharing a common infrastructure. When scheduling individually the jobs submitted by various users, one crucial issue is to ensure fairness. This work presents a new fair scheduling algorithm called OStrich whose principle is to maintain a virtual time-sharing schedule in which the same amount of processors is assigned to each user. The completion times in the virtual schedule determine the execution order on the physical processors. Then, the campaigns are interleaved in a fair way by OStrich. For independent sequential jobs, we show that OStrich guarantees the stretch of a campaign to be proportional to campaign's size and the total number of users. The stretch is used for measuring by what factor a workload is slowed down relative to the time it takes on an unloaded system. The theoretical performance of our solution is assessed by simulating OStrich compared to the classical FCFS algorithm, issued from synthetic workload traces generated by two different user profiles. This is done to demonstrate how OStrich benefits both types of users, in contrast to FCFS.

Keywords: Job scheduling, Fairness, Job campaigns, Multi-User, Workload Traces

Fair share is not enough: measuring fairness in scheduling with cooperative game theory

Piotr Skowron, Krzysztof Rządca
Faculty of Mathematics, Informatics and Mechanics
University of Warsaw, Poland
{p.skowron, krzadca}@mimuw.edu.pl

We consider the problem of fair scheduling in a multi-organizational system in which organizations are contributing their own resources to the global pool and the jobs to be processed on the common resources. We consider on-line, non-clairvoyant scheduling of sequential jobs without preemption. The organizations must agree on the order of executing the jobs, i.e. on the scheduling algorithm. To ensure that the organizations are willing to cooperate the scheduling algorithm must be fair. To characterize fairness, we use cooperative game theory approach. The contribution of an organization is computed based on how this organization influences the utility (which can be any metric, e.g., flow time, turnaround, resource allocation, etc.) of all organizations. Formally, the contribution of the organization is its Shapley value in the cooperative game. The scheduling algorithm should ensure that the contributions of the organizations are close to their utilities. Our previous work proves that this problem is NP-hard and hard to approximate. In this paper we propose a heuristic scheduling algorithm for the fair scheduling problem. We experimentally evaluate the heuristic and compare its fairness to other algorithms (fair share, round robin) and the exact exponential algorithm. Our results show that fairness of the heuristic algorithm is close to the optimal. The difference between our heuristic and the fair share algorithm is more visible on longer traces with more organizations. We believe that these results prove that fair share might not be an optimal solution for multi-organizational systems.

Keywords: scheduling, fairness, game theory, Shapley value

Setting up clusters of computing units to process several data streams efficiently

Daniel Millot, Christian Parrot
Telecom sudParis, Institut Mines-Telecom
France
{daniel.millot, christian.parrot}@mines-telecom.fr

Let us consider an upper bounded number of data streams to be processed by a Divisible Load application. The total workload is unknown and the available speeds for communicating and computing can be poorly a priori estimated. This paper presents a resource selection method that aims at maximizing the throughput of this processing. From a set of processing units linked by a network, this method consists in forming an optimal set of master-workers clusters. Results of simulations are presented to assess the efficiency of this method experimentally. Before focusing on the proposed resource selection method, the paper reminds the underlying adaptive scheduling method on which it relies.

Keywords: Parallel Computing, Scheduling, Divisible Load, Resource Selection, Adaptation

A New Multi-Criteria based Divisible Load Scheduling Algorithm

Shamsollah Ghanbari¹, Mohamed Othman^{1,2}, Wah June Leong¹,
Mohd Rizam Abu Bakar¹

¹Laboratory of Computational Science and Mathematical Physics

Institute For Mathematical Research

Universiti Putra Malaysia, 43400 UPM Serdang, Selangor D.E., Malaysia

²Dept of Communication Tech and Network

Faculty of Computer Science and Information Technology

Universiti Putra Malaysia, 43400 UPM Serdang, Selangor D.E., Malaysia

myrshg@gmail.com, mothman@fsktm.upm.my

Divisible load theory has become a popular area of researches during the two past decades. Based on divisible load theory the computations and communications can be divided into some arbitrarily independent parts and each part can be processed independently by a processor. The goal of existing divisible load scheduling algorithms is to reduce finish time and they do not consider any priority in processing fraction of loads. In some situation it is necessary to be processed the loads based on some priorities. In this paper we propose a new multi-criteria divisible load scheduling algorithm. Experimental result indicates the proposed algorithm can handle the priority of processors.

Keywords: Divisible load scheduling, priority, multi criteria, Analytical Hierarchy Process

Bridging the Gap: From Cellular Automata to Differential Equation Models for Pedestrian Dynamics

Felix Dietrich¹, Gerta Koester², Michael Seitz², Isabella von Sivers²

¹Technical University of Munich, Boltzmannstr. 3, 85747 Garching, Germany

²Munich University of Applied Sciences, Lothstr. 64, 80335 München, Germany

`felix.dietrich@tum.de`

`{koester, m.seitz, isabella.von_sivers}@hm.edu`

Cellular automata (CA) and ordinary differential equation (ODE) based models compete for dominance in microscopic pedestrian dynamics. Both are inspired by the idea that pedestrians are subject to forces. However, there are two major differences: In a CA, movement is restricted to a coarse grid. Navigation is achieved directly by pointing the movement in the direction of the forces. In force based ODE models, navigation is computed indirectly through the acceleration vector. We present two models emanating from the CA and ODE approaches that remove these two differences: the Optimal Steps Model and the Gradient Navigation Model. Both models are very robust and produce trajectories similar to each other, bridging the gap between the older models. Both approaches are grid-free and free of oscillations, giving cause to the hypothesis that the two major differences are also the two major weaknesses of the older models.

Keywords: Cellular Automata, Ordinary Differential Equation, Pedestrian Dynamics, Optimal Step Model, Gradient Navigation Model

Cellular model of pedestrian dynamics with adaptive time span

Marek Bukáček, Pavel Hrabák, Milan Krbálek
Faculty of Nuclear Sciences and Physical Engineering
Czech Technical University in Prague
Czech Republic
{bukacma2, pavel.hrabak, milan.krbalek}@fjfi.cvut.cz

A cellular model of pedestrian dynamics based on the Floor Field model is introduced. Contrary to the parallel update in Floor Field, the concept of adaptive time span is introduced. This concept, together with the concept of bounds, supports the spontaneous line formation and chaotic queue in front of the bottleneck. Model simulations are compared to the experiment “passing through”, from which a phase transition from low to high density is observed.

Keywords: pedestrian dynamics, experimental study of phase transition, adaptive time span

The use of GPGPU in continuous and discrete models of crowd dynamics

Hubert Mróz, Jarosław Wąs, Paweł Topa
AGH University of Science and Technology,
al. Mickiewicza 30, 30-059 Kraków, Poland
hubertmroz@o2.pl, {jarek, topa}@agh.edu.pl

The aim of the study is twofold: firstly to compare the possibilities of using GPGPU (General-Purpose Computing on Graphics Processing Units) in continuous and discrete crowd dynamics simulation, secondly to draw conclusions on the applicability of GPUs in engines of professional crowd simulations. For this purpose the authors have implemented two models of pedestrian dynamics: continuous - Social Forces model and discrete, cellular automata based - Social Distances model. The presented simulations refer to outdoor, large area pedestrian movement.

Keywords: Cellular Automata, GPGPU, Crowd modeling

Modeling Behavioral Traits of Employees in a Workplace with Cellular Automata

Petros Saravakos, Georgios Ch. Sirakoulis
Department of Electrical and Computer Engineering
Democritus University of Thrace
67100 Xanthi, Greece
{psaravak, gsirak}@ee.duth.gr

The aim of this paper is to examine a parameterized working environment on the basis of behavioral traits of employees in an organization. Firstly we define several behavioral traits of the employees, including the employee's attitude in the workplace, the influence radius and her/his reluctance to adapt to organizational norms, stated as insistence. The combination of these traits allows us to model employee interactions to a satisfactory extent for a realistic model of the working environment. Secondly, we define two metrics illustrating the policies adopted by the organization either to restrain unwanted or impose desirable behavioral patterns. Finally, the corresponding Cellular Automaton (CA) model enables us to utilize the aforementioned parameters and to simulate the under study workplace. The presented simulation results can be used as a complementary tool for managerial decisions illustrating workplace dynamics and forecast future trends.

Keywords: Behavioral Traits, Cellular Automata, Working Environment, Simulation

Probabilistic Pharmaceutical Modelling: A Comparison Between Synchronous and Asynchronous Cellular Automata

Marija Bezbradica, Heather J. Ruskin, Martin Crane
Centre for Scientific Research and Complex Systems Modelling (Sci-Sym)
School of Computing, Dublin City University
Dublin, Ireland
`{mbezbradica,hruskin,mcrane}@computing.dcu.ie`

In recent years, the field of pharmaceutical modelling has benefited from using probabilistic methods based on cellular automata, which seek to overcome some of the limitations of differential equation based models. By modelling discrete structural element interactions instead, this approach is able to provide data quality adequate for the early design phases in drug modelling. In the literature, both synchronous (CA) and asynchronous (ACA) types of automata have been used previously, but analysis of their comparative impact on model output is lacking. In this paper, we compare several variants of probabilistic CA and ACA algorithms for building models of complex systems used in controlled drug delivery, and analyse advantages and disadvantages of different choices. The appropriate update mechanism, besides having an impact on perceived realism of the simulation, also determines the applicability and performance of different parallelisation algorithms when used in large-scale simulations.

Keywords: discrete systems, complex modelling, controlled drug delivery systems, parallel algorithms

Coupling Lattice Boltzmann Gas and Level Set Method for Simulating Free Surface Flow in GPU/CUDA Environment

Tomir Kryza, Witold Dzwinel
Institute of Computer Science
AGH University of Science and Technology
Kraków, Poland
{tomir, dzwinel}@agh.edu.pl

We present here a proof-of-concept of a novel, efficient method for modeling of liquid/gas interface dynamics. Our approach consists in coupling the lattice Boltzmann gas (LBG) and the level set (LS) methods. The inherent parallel character of LBG accelerated by level sets is the principal advantage of our approach over similar particle based solvers. Consequently, this property allows for efficient use of our solver in GPU/CUDA environment. We demonstrate preliminary results and GPU/CPU speedups simulating two standard free surface fluid scenarios: the falling droplet and the breaking dam problems.

Keywords: free surface flow, lattice Boltzmann gas, level sets, CUDA, GPGPU

Creation of Agent's Vision of Social Network through Episodic Memory

Michał Wrzeszcz¹, Jacek Kitowski^{1,2}

¹AGH University of Science and Technology, Faculty of Computer Science
Electronics and Telecommunications, Department of Computer Science
al. Mickiewicza 30, 30-059 Kraków, Poland

²AGH University of Science and Technology, ACC Cyfronet AGH
ul. Nawojki 11, 30-950 Kraków, Poland
{wrzeszcz,kito}@agh.edu.p

Human societies appear in many types of simulations. One of the most important and the most difficult society elements to be modelled is the social context. In this paper we show how social context can be provided using agents that are equipped with internal visions of social relations between others. Internal vision is a representation of social relations from the agent's point of view so, being subjective, it may be inconsistent with the reality. We introduce the agent model and the mechanism of rebuilding the agent's internal vision that is similar to that used by humans. An experimental proof of concepts is also presented.

Keywords: social networks, behaviour modelling, simulation of human societies, multi-agent systems, social context

The influence of multi-agent cooperation on the efficiency of taxi dispatching

Michał Maciejewski^{1,2}, Kai Nagel²

¹Institute of Machines and Motor Vehicles, Faculty of Machines and Transportation
Poznan University of Technology, ul. Piotrowo 3, 60-965 Poznan, Poland

²Transport Systems Planning (VSP), Institute for Land and Sea Transport Systems
TU Berlin, Salzufer 17-19 Sekr. SG12, D-10587 Berlin, Germany

{maciejewski, nagel}@vsp.tu-berlin.de

The paper deals with the problem of the optimal collaboration scheme in taxi dispatching between customers, taxi drivers and the dispatcher. The authors propose three strategies that differ by the amount of information exchanged between agents and the intensity of cooperation between taxi drivers and the dispatcher. The strategies are evaluated by means of a microscopic multi-agent transport simulator (MATSim) coupled with a dynamic vehicle routing optimizer (the DVR-POptimizer), which allows to realistically simulate dynamic taxi services as one of several different transport means, all embedded into a realistic environment. The evaluation is carried out on a scenario of the Polish city of Mielec. The results obtained prove that the cooperation between the dispatcher and taxi drivers is of the utmost importance, while the customer-dispatcher communication may be reduced to minimum and compensated by the use of more sophisticated dispatching strategies, thereby not affecting the quality of service.

Keywords: dynamic taxi dispatching, dynamic vehicle routing, on-line optimization, multi-agent simulation, MATSim, traffic flow simulation, simulation-based optimization

Basic Endogenous-Money Economy: an Agent-Based Approach

Ivan Blečić, Arnaldo Cecchini, Giuseppe A. Trunfio
DADU, Department of Architecture, Design and Urban Planning
University of Sassari, Alghero, Italy
{ivan, cecchini, trunfio}@uniss.it

We present an agent-based model of a simple endogenous-money economy. The model simulates agents representing individual persons who can work, consume, invent new products and related production technologies, apply for a loan from the bank and start up a business. Through the interaction of persons with the firms, we simulate the production of goods, consumption and labour market. This setting allows us to explore how an endogenous-money economy may build up from scratch, as an emergent property of actions and interactions among heterogeneous agents, once the money is being injected into a non-monetary self-production (or barter) economy. We provide and discuss the results of several computational experiments under three scenarios: (1) with just one firm, (2) with a limited number of firms and abundant workforce, (3) and with unlimited number of firms.

Keywords: agent-based computational economics, endogenous-money economy, heterogeneous agents

A shaving method for interval linear systems of equations

Milan Hladík, Jaroslav Horáček
Charles University, Faculty of Mathematics and Physics
Department of Applied Mathematics
Malostranské nám. 25, 118 00, Prague, Czech Republic
{hladik, horacek}@kam.mff.cuni.cz

We propose an iterative improvement method for an enclosure of the solution set of a system of interval linear equations. The method sequentially cuts off (shaves) parts of a given enclosure that contain no solution, yielding thus tighter enclosures. Since shaving can be done independently in the coordinates, the procedure is easily parallelized. Our approach is convenient for problems with wide input intervals, where traditional methods give poor enclosures. Finally, we present a limited computational study.

Keywords: Interval linear systems, interval matrix, reliable computing, parallelization

Inner Estimation of Linear Parametric AE-Solution Sets

Evgenija Popova
Institute of Mathematics and Informatics
Bulgarian Academy of Sciences
Acad. G. Bonchev str., bldg. 8, 1113 Sofia, Bulgaria
epopova@math.bas.bg, epopova@bio.bas.bg

We consider a linear algebraic systems $A(p)x = b(p)$, where the elements of the matrix and the right-hand side vector are linear functions of uncertain parameters varying within given intervals. For the so-called parametric AE-solution set, which is defined by two disjoint sets of existentially and universally quantified parameters, an interval inner estimation is sought. In the special case when the parametric AE-solution set has linear shape, we give parametric generalization of the so-called centered approach, which was developed so far for nonparametric and parametric tolerable solution sets. The small computational complexity of the method providing guaranteed numerical results makes it attractive for solving large scale problems. Numerical examples illustrate the applicability of the method for controllability analysis of linear dynamical systems and to linear systems arising in finite element analysis of uncertain mechanical structures.

Keywords: interval linear equations, dependent data, AE-solution sets, inner inclusion

Finding Enclosures for Linear Systems using Interval Matrix Multiplication in CUDA

Alexander Dallmann, Philip-Daniel Beck
University of Würzburg, Chair of Computer Science II
Am Hubland, D 97074 Würzburg, Germany
dallmann@informatik.uni-wuerzburg.de

In this paper we present CUDA kernels that compute an interval matrix product. Starting from a naive implementation we investigate possible speedups using commonly known techniques from standard matrix multiplication. We also evaluate the achieved speedup when our kernels are used to accelerate a variant of an existing algorithm that finds an enclosure for the solution of a linear system. Moreover the quality of our enclosure is discussed.

Keywords: GPGPU, Interval arithmetic, Linear algebra, Parallel computing

GPU accelerated metaheuristics for solving large scale parametric interval algebraic systems

Iwona Skalna and Jerzy Duda
AGH University of Science and Technology
Kraków, Poland
skalna@agh.edu.pl, jduda@zarz.agh.edu.pl

A study of GPU accelerated solution of parametric interval linear systems is presented. Two population-based metaheuristics are the main computational tool. Several illustrative examples from structural mechanics are employed to show that the proposed approach can significantly reduce time complexity. The stress is put on large uncertainties which are hard to be dealt with less time-consuming methods.

Keywords: GPU computing, parametric linear systems, metaheuristics

Parallel approach to Monte Carlo simulation for Option Price Sensitivities using the Adjoint and Interval Analysis

Grzegorz Kozikowski¹, Bartłomiej Kubica²

¹School of Computer Science

University of Manchester, United Kingdom

²Institute of Control and Computation Engineering

Warsaw University of Technology, Poland

grzegorz.kozikowski@manchester.ac.uk

bkubica@elka.pw.edu.pl

This paper concerns a powerful approach to evaluation of Option Price sensitivities using the Monte Carlo simulation. The idea is to apply the parallel GPU architecture and Automatic Differentiation methods. In order to study rounding errors, the interval arithmetic is applied. Considerations are based on a sequential and parallel authors' implementation. For efficient differentiation, the Adjoint method is employed. Computational experiments include analysis of performance, uncertainty error and rounding error and consider Black-Scholes and Heston models.

Keywords: option pricing, the Greeks, Automatic Differentiation, the Adjoint, interval computations, GPGPU, CUDA

Subsquares Approach - Simple Scheme for Solving Overdetermined Interval Linear Systems

Jaroslav Horáček, Milan Hladík
Charles University, Faculty of Mathematics and Physics
Department of Applied Mathematics
Malostranské nám. 25, 118 00, Prague, Czech Republic
{horacek, hladik}@kam.mff.cuni.cz

In this work we present a new simple but efficient scheme - Subsquares approach - for development of algorithms for enclosing the solution set of overdetermined interval linear systems. We are going to show two algorithms based on this scheme and discuss their features. We start with a simple algorithm as a motivation, then we continue with a sequential algorithm. Both algorithms can be easily parallelized. The features of both algorithms will be discussed and numerically tested.

Keywords: interval linear systems, interval enclosure, overdetermined systems, parallel computing

Using quadratic approximations in an interval method of solving underdetermined and well-determined nonlinear systems

Bartłomiej Jacek Kubica
Institute of Control and Computation Engineering
Warsaw University of Technology, Poland
bkubica@elka.pw.edu.pl

This paper considers quadratic approximation as a narrowing tool in an interval branch-and-prune method. A quadratic equation with interval parameters is solved. A heuristic to decide, when to use the developed operator, is proposed. Numerical results for some benchmark problems are presented and analyzed.

Keywords: nonlinear equations systems, interval computations, quadratic approximation, interval quadratic equation, heuristic

Interval Finite Difference Method for Solving the Problem of Bioheat Transfer between Blood Vessel and Tissue

Malgorzata A. Jankowska

Poznan University of Technology, Institute of Applied Mechanics

Jana Pawla II 24, 60-965 Poznan, Poland

malgorzata.jankowska@put.poznan.pl

The paper concerns the problem of bioheat transfer between a single large blood vessel and a surrounding tissue. On the basis of the conventional finite difference scheme with the appropriate truncation error terms included, the interval finite difference method is proposed. The interval values that contain the local truncation error of the conventional scheme can be found directly for a limited number of problems or just approximated in the way described.

Keywords: interval finite difference method and interval arithmetic, vessel model of bioheat transfer, uncertain values of parameters

Chosen Interval Methods For Solving An Interval Linear System of Equations

Barbara Szyszka
Poznan University of Technology
Piotrowo 3A, 60-965 Poznan, Poland
`barbara.szyszka@put.poznan.pl`

The paper is devoted to chosen interval methods for solving an interval linear system of equations. This kind of the interval linear system of equations from the central difference interval method of the second order for solving the one dimensional wave equation is obtained. A matrix of coefficients of this system is a special type: sparse and five-diagonals, and can be with interval parameter. A comparison of four interval algorithms in floating point interval arithmetic for solving this linear system will be presented during conference.

Keywords: interval methods, wave equation, difference method, interval linear systems

Numerical reproducibility in HPC: the interval point of view

Nathalie Revol, Philippe Theveny
INRIA - LIP, Universite de Lyon
Lyon, France
`{nathalie.revol, philippe.theveny}@ens-lyon.fr`

What is called “numerical reproducibility” is the problem of getting the same result, when the scientific computation is run several times, either on the same machine (and this is called “repeatability”) or on different machines, with different numbers of processing units, types, execution environments, computational loads etc. This problem is especially stringent for HPC results. For interval computations, numerical reproducibility is of course an issue for testing and debugging purposes. However, as long as the computed result encloses the exact and unknown result, the inclusion property, which is the main property of interval arithmetic, is satisfied and getting bit for bit identical results may not be crucial. However, implementation issues may invalidate the inclusion property, in particular if the rounding modes set by the user are modified during the execution. We will present several ways to circumvent these issues, on the example of the product of matrices with interval coefficients.

Keywords: interval arithmetic, numerical reproducibility, HPC

A parallel solver for the time-periodic Navier-Stokes equations

Peter Arbenz¹, Daniel Hupp¹, Dominik Obrist²

¹ETH Zürich, Computer Science Department, Zürich, Switzerland

²ETH Zürich, Institute of Fluid Dynamics, Zürich, Switzerland

arbenz@inf.ethz.ch

We investigate parallel algorithms for the solution of the Navier-Stokes equations in space-time. For periodic solutions, the discretized problem can be written as a large non-linear system of equations. This system of equations is solved by a Newton iteration. The Newton correction is computed using a preconditioned GMRES solver. The parallel performance of the algorithm is illustrated.

Keywords: Navier-Stokes equations, time-periodic solution, space-time parallelism

Parallel Numerical Algorithms for Simulation of Rectangular Waveguides by Using GPU

Raimondas Čiegis¹, Andrej Bugajev¹, Žilvinas Kancleris², Gediminas Šlekas²

¹Vilnius Gediminas Technical University

Saulėtekis av. 11, LT10223 Vilnius, Lithuania

²Semiconductor Physics Institute

A. Goštauto 11, LT01108 Vilnius, Lithuania

{rc, andrej.bugajev}@vgtu.lt

kancleris@pfi.lt gediminas.slekas@gmail.com

In this article we consider parallel numerical algorithms to solve the mathematical model of Maxwell equations, that describes wave propagation in rectangular waveguide. The main goal is to formulate and analyze the simple algorithmic template to solve this problem using CUDA platform. The template is based on explicit finite difference schemes obtained after approximation of systems of differential equations on the staggered grid. The geometry is assumed to be a cube in three dimensional space, the classical stencils of discrete schemes are used. The parallelization of the discrete algorithm is based on the domain decomposition method. The theoretical complexity model is derived and the scalability of the parallel algorithm is investigated. Some results of numerical simulations are presented and the efficiency of the proposed parallel algorithm is investigated.

Keywords: parallel algorithms, numerical simulation, wave radiation, CUDA, GPU, scalability analysis

OpenACC Parallelisation For Diffusion Problems, Applied To Temperature Distribution On A Honeycomb Around the Bee Brood: A Worked Example Using BiCGSTAB

Hermann Eberl, Rangarajan Sudarsan
University of Guelph
Guelph ON, N1G2W1, Canada
{heberl, rsudarsa}@uoguelph.ca

We discuss a simple OpenACC implementation of BiCGSTAB, a Krylov subspace solver for linear systems that is particularly suitable for non-symmetric problems. Problems of this type arise in the numerical solution of diffusion-reaction problems, where the linear solver constitutes the most computationally expensive component of the simulation (80% of time spend) and therefore has often been the primary target for parallelization. We deploy and test this method on a desktop workstation with two supported GPU accelerators, one targeted for high performance computing, one a consumer level GPU, to compute the temperature distribution on a honeycomb around the bee brood. The paper is written from a user's, not from a GPU computing expert's perspective and aims to fill a gap we noticed between real world application problems and the simple problems solved in introductory OpenACC tutorials. Due to space limitations, the test case in the accompanying manuscript is a simple steady state linear problem. In the conference we will discuss the application to a more involved transient quasi-linear problem.

Keywords: diffusion equation, GPU, OpenACC, Krylov-suspace methods, honeycomb, bee brood

Application of CUDA for Acceleration of Calculations in Boundary Value Problems Solving Using PIES

Andrzej Kuźelewski, Eugeniusz Zieniuk, Agnieszka Bołtuć
University of Białystok, Institute of Computer Science
Sosnowa 64, 15-887 Białystok
{akuzel, ezieniuk, aboltuc}@ii.uwb.edu.pl

The main purpose of this paper is examination of an application of modern parallel computing solutions to speed up the calculation in the numerical solution of parametric integral equations systems (PIES). Solving boundary value problems by PIES sometimes requires large computing time, particularly in more complex problems. This paper presents use of graphics cards programming in general-purpose applications (GPGPU). The boundary value problems modelled by 3D Navier-Lamé equations are solved using PIES and NVidia® CUDA technology. The testing example shows that the use of GPGPU significantly increases speed of calculations in PIES.

Keywords: parametric integral equations systems (PIES), CUDA, boundary value problems

Modeling and simulations of beam stabilization in edge-emitting broad area semiconductor devices

Mindaugas Radziunas¹, Raimondas Čiegis¹

¹Weierstrass Institute, 10117 Berlin, Germany

²Vilnius Gediminas Technical University

Saulėtekio al. 11, LT-10223 Vilnius, Lithuania

radziuna@wias-berlin.de

A 2+1 dimensional PDE traveling wave model describing spatial-lateral dynamics of edge-emitting broad area semiconductor devices is considered. A numerical scheme based on a split-step Fourier method is presented and implemented on a parallel compute cluster. Simulations of the model equations are used for optimizing of existing devices with respect to the emitted beam quality, as well as for creating and testing of novel device design concepts.

Keywords: Broad area device, traveling wave model, numerical scheme, simulations, beam improvement

Concurrent nomadic and bundle search: A class of parallel algorithms for local optimization

Costas Voglis¹, Dimitrios Papageorgiou², Isaac Lagaris¹

¹Department of Computer Science

University of Ioannina, P.O. BOX 1186

GR-45110, Ioannina, Greece

²Department of Materials Science and Engineering

University of Ioannina, P.O. BOX 1186

GR-45110, Ioannina, Greece

voglis@cs.uoi.gr

We present a family of algorithms for local optimization that exploit the parallel architectures of contemporary computing systems to accomplish significant performance enhancements. This capability is important for demanding real time applications, as well as, for problems with timeconsuming objective functions. The proposed concurrent schemes namely nomadic and bundle search are based upon well established techniques such as quasi-Newton updates and line searches. The parallelization strategy consists of (a) distributed computation of an approximation to the Hessian matrix and (b) parallel deployment of line searches on different directions (bundles) and from different starting points (nomads). Preliminary results showed that the new parallel algorithms can solve problems in less iterations than their serial rivals.

Keywords: parallel local optimization, quasi-Newton, SR1 update, line search, nomadic search, bundle search

Parallel Multi-Objective Memetic Algorithm for Competitive Facility Location

Algirdas Lančinskas, Julius Žilinskas
Institute of Mathematics and Informatics, Vilnius University
Akademijos 4, 08663, Vilnius, Lithuania
{algirdas.lancinskas,julius.zilinskas}@mii.vu.lt

A hybrid genetic algorithm for global multi-objective optimization is parallelized and applied to solve competitive facility location problems. The impact of usage of local search on the performance of parallel algorithm has been investigated. An asynchronous version of parallel genetic algorithm with local search has been proposed and investigated by solving competitive facility location problem utilizing hybrid distributed and shared memory parallel programming model on high performance computing system.

Keywords: Facility Location, Multi-Objective Optimization, Memetic Algorithms

Parallelization of Encryption Algorithm Based on Chaos System and Neural Networks

Dariusz Burak
Faculty of Computer Science and Information Technology
West Pomeranian University of Technology
ul.Żołnierska 49, 71-210 Szczecin, Poland
dburak@wi.zut.edu.pl

In this paper, the results of parallelizing of encryption algorithm based on a chaos system and neural networks are presented. A data dependence analysis of loops was applied in order to parallelize the algorithm. The parallelism of the algorithm is demonstrated in accordance with the OpenMP standard. As a result of my study, it was stated that the most time-consuming loops of the algorithm are suitable for parallelization. The efficiency measurement of a parallel program is showed.

Keywords: encryption algorithm, chaotic system, neural networks, parallelization, OpenMP

Towards Standardization of Measuring the Usability of Parallel Languages

Ami Marowka
Department of Computer Science
Bar-Ilan University, Israel
amimar2@yahoo.com

The efforts of the research community and the software industry to make the art of parallel programming easier continue. Measuring the usability of contemporary parallel programming languages and libraries by empirical studies is the key to understanding how programmers are thinking, designing, coding, and debugging parallel programs. In this paper we take apart into their component ingredients the empirical experiments done in the recent years. By analyzing each component separately we can better understand what is missing in these experiments and thereby improve the outcome of future studies. The result of this work is a set of recommendations that aims to make usability studies more convincing so that parallel language designers will take them seriously.

Keywords: Usability, Parallel Language, Empirical Study, Productivity

Experiences with Implementing Task Pools in Chapel and X10

Claudia Fohry¹, Jens Breitbart²

¹Research Group Programming Languages / Methodologies
Universität Kassel, Germany

²Institut for Applied and Numerical Mathematics IV
Karlsruhe Institute of Technology, Germany

amimar2@yahoo.com, jens.breitbart@kit.edu

The Partitioned Global Address Space (PGAS) model is a promising approach to combine programmability and performance in an architecture-independent way. Well-known representatives of PGAS languages include Chapel and X10. Both languages incorporate object orientation, but fundamentally differ in their way of accessing remote memory, synchronization constructs, and other issues of language design.

From a user's point of view, this paper evaluates and compares the languages with focus on programmability. We concentrate on the interplay between object orientation and parallelism/ distribution, the deployment of synchronization constructs, and other issues of coding task parallelism. In particular, we discuss the realization of patterns such as objects that internally contain distributed arrays, and suggest improvements such as support for activity-local and place-local data, as well as scalar variable-based reduction. Our study is based on Unbalanced Tree Search (UTS), a well-known benchmark that uses task pools.

Keywords: PGAS languages Chapel X10 Programmability User experience

Parampl: A simple approach for parallel and distributed execution of AMPL programs

Artur Olszak¹, Andrzej Karbowski^{2,3}

¹Institute of Computer Science

Warsaw University of Technology, Poland

²Institute of Control and Computation Engineering

Warsaw University of Technology, Poland

³NASK, Research and Academic Computer Network, Warsaw, Poland

a.olszak@ii.pw.edu.pl, a.karbowski@elka.pw.edu.pl

Due to the physical processor frequency scaling constraint, current computer systems are equipped with more and more processing units. Therefore, parallel computing has become an important paradigm in the recent years. AMPL is a comprehensive algebraic modeling language for formulating optimization problems. However, AMPL itself does not support defining tasks to be executed in parallel. Although in last years the parallelism is often provided by solvers, which take advantage of multiple processing units, in many cases it is more efficient to formulate the problem in a decomposed way and apply various problem specific enhancements. Moreover, when the number of cores is permanently growing, it is possible to use both types of parallelism.

This paper presents the design of Parampl - a simple tool for parallel and distributed execution of AMPL programs. Parampl introduces explicit asynchronous execution of AMPL subproblems from within the program code. Such an extension implies a new view on AMPL programs, where a programmer is able to define a complex, parallelized optimization task.

Keywords: AMPL paralel distributed optimization modeling languages

Optimization of an OpenCL-Based Multi-Swarm PSO Algorithm on an APU

Wayne Franz, Parimala Thulasiraman, Ruppa Thulasiram
University of Manitoba, Canada
{umfranzw, thulasir, tulsi}@cs.umanitoba.ca

The multi-swarm particle swarm optimization (MPSO) algorithm incorporates multiple independent PSO swarms that cooperate by periodically exchanging information. This algorithm is embarrassingly parallel, making it an ideal candidate for execution on data-parallel GPUs. Unfortunately, MPSO is also memory bound, limiting its performance. Recently, heterogeneous multi-core architectures such as the AMD Accelerated Processing Unit (APU) architecture have fused the CPU and GPU together on a single die, eliminating the traditional PCIe bottleneck between them. In this paper, we provide our experiences developing an OpenCL-based MPSO algorithm for the task scheduling problem on the APU architecture. We use the AMD A8-3530MX APU that packs four x86 computing cores and 80 four-way processing elements. We make effective use of hardware features such as the hierarchical memory structure on the APU, the four-way very long instruction word (VLIW) feature for vectorization, and global-to-local memory DMA transfer features. We observe a 29% decrease in overall execution time on the APU.

Keywords: OpenCL, Particle Swarm Optimization, Accelerated Processing Unit, Optimization

Algorithms for In-Place Matrix Transposition

Fred G. Gustavson¹, David W. Walker²

¹IBM T. J. Watson Research Center, Emeritus, and Umeå University

²School of Computer Science and Informatics, Cardiff University, UK

walkerdw@cardiff.ac.uk

This paper presents an implementation of an in-place swap-based algorithm for transposing rectangular matrices, and a proof of correctness is also sketched. The implementation is based on an algorithm described by Tretyakov and Tyrtshnikov, but we have introduced a number of variations. In particular, we show how the original algorithm can be modified to require constant additional memory. We also identify opportunities for exploiting parallelism.

Keywords: Matrix transposition, Parallel computing, In-place algorithms

FooPar: A Functional Object Oriented Parallel Framework in Scala

Felix P. Hargreaves, Daniel Merkle
Department of Mathematics and Computer Science
University of Southern Denmark
Odense, Denmark
{daniel, felhar07}@imada.sdu.dk

We present FooPar, an extension for highly efficient Parallel Computing in the multi-paradigm programming language Scala. Scala offers concise and clean syntax and integrates functional programming features. Our framework FooPar combines these features with parallel computing techniques. FooPar is designed modular and supports easy access to different communication backends for distributed memory architectures as well as high performance math libraries. In this article we use it to parallelize matrix-matrix multiplication and show its scalability by a isoefficiency analysis. In addition, results based on an empirical analysis on two supercomputers are given. We achieve close-to-optimal performance wrt. theoretical peak performance. Based on this result we conclude that FooPar allows to fully access Scala's design features without suffering from performance drops when compared to implementations purely based on C and MPI.

Keywords: Functional Programming, Scala language extension, Isoefficiency

Prototyping framework for parallel numerical computations

Ondřej Meca, Stanislav Böhm, Marek Běhálék, Martin Šurkovský
Department of Computer Science
FEI VŠB Technical University of Ostrava
Ostrava, Czech Republic
{ondrej.meca, stanislav.bohm, marek.behalek}@vsb.cz
martin.surkovsky@vsb.cz

Our research is focused on the simplification of parallel programming for distributed memory systems. Our goal is to build a unifying framework for creating, debugging, profiling and verifying parallel applications. The result of this effort is an open source tool Kaira. In this paper, we focus on prototyping parallel applications. We have extended Kaira by the ability to generate parallel libraries. More precisely, we present a framework for fast prototyping of parallel numerical computations. We demonstrate our idea on a combination of parallel libraries generated by our tool Kaira and GNU Octave. Hence, a user can verify the idea in a short time, create a real running program and verify its performance and scalability.

Keywords: prototyping, parallel computing, visual programming, libraries

Core Allocation Strategies on Multicore Platforms to Accelerate Forest Fire Spread Predictions

Tomàs Artés, Andrés Cencerrado, Ana Cortés, Tomàs Margalef
Computer Architecture and Operating Systems Department
Universitat Autònoma de Barcelona
Campus UAB, Edifici Q, 08193 Bellaterra, Spain
{tomas.artes, andres.cencerrado, ana.cortes, tomas.margalef}@uab.cat

In order to predict the forest fire spread software simulators are developed. Simulators require a several input parameters which usually are very difficult to know accurately. The input data uncertainty can provoke a mismatch between the predicted forest fire spread and the actual evolution. To overcome this uncertainty a two stage prediction methodology based on input parameters calibration is used. In the first stage a genetic algorithm is applied to find the best input parameter set. Afterwards, the prediction is carried out using the calibrated input parameter set. This method improves the prediction error, but increments the execution time in a context with hard time constraints. We present a first approach to speed up the two stage prediction methodology by exploiting multicore architectures. A hybrid MPI-OpenMP application has been developed and different allocation policies have been tested to accelerate the forest fire prediction with an efficient use of the available resources.

Keywords: forest fire, simulation, data uncertainty, hybrid MPI-OpenMP, evolutionary computation, resource assignment, multicore architecture

Effects of Segmented Finite Difference Time Domain on GPU

Jose Mijares^{1,2}, Parimala Thulasiraman¹, Ruppa Thulasiram¹, Gagan Battoo¹

¹Department of Computer Science

University of Manitoba, MB, Canada

²Department of Electrical and Computer Engineering

University of Manitoba, MB, Canada

thulasir@cs.umanitoba.ca

Finite Difference Time Domain (FDTD) is the most popular method in computational electromagnetics. In acoustics, FDTD is often used as a numerical analysis technique to model mechanical wave and acoustics. FDTD in general is computationally expensive in terms of time due to its large number of time steps for accurate precision and is data parallel in nature. However, it is also memory bounded. Although previous work on FDTD has studied the effect of parallelizing FDTD on accelerators to reduce computational cost, the memory bounded problem has not been studied. In this work we consider the segmented FDTD (SFDTD) algorithm that divides the problem space into segments to reduce computational redundancy and also reduce memory. We exploit the memory hierarchy of the GPU to efficiently implement the SFDTD algorithm. To the best of our knowledge, this is the first work that studies the implementation of the SFDTD algorithm on GPU and its effect on memory.

Keywords: GPU, CUDA, Segmented FDTD, Acoustics

Resolving Load Balancing Issues in BWA on NUMA Multicore Architectures

Charlotte Herzeel^{1,4}, Thomas J. Ashby^{2,4}, Pascal Costanza^{3,4}, Wolfgang De Meuter¹

¹Software Languages Lab, Vrije Universiteit Brussel

Pleinlaan 2, B-1050 Brussel, Belgium

²IMEC, Kapeldreef 75, B-3001 Leuven, Belgium

³Intel, Veldkant 31, B-2550 Kontich, Belgium

⁴Intel ExaScience Lab

Kapeldreef 75, B-3001 Leuven, Belgium

{caherzee, wdmeuter}@vub.ac.be, ashby@imec.be

pascal.costanza@intel.com

Running BWA in multithreaded mode on a multi-socket server results in poor scaling behaviour. This is because the current parallelisation strategy does not take into account the load imbalance that is inherent to the properties of the data being aligned, e.g. varying read lengths and numbers of mutations. Additional load imbalance is also caused by the BWA code not anticipating certain hardware characteristics of multi-socket multicore, such as the non-uniform memory access time of the different cores. We show that rewriting the parallel section using Cilk removes the load imbalance, resulting in a factor two performance improvement over the original BWA.

Keywords: BWA, multithreading, NUMA, load balancing, Cilk

K-mulus: Strategies for BLAST in the Cloud

Christopher Hill¹, Carl Albach¹, Sebastian Angel², Mihai Pop¹

¹Center for Bioinformatics and Computational Biology

University of Maryland, College Park, MD, USA

²University of Texas, Austin, TX, USA

{fcmhill, calbach, mpopg}@umd.edu, sebs@cs.utexas.edu

With the increased availability of next-generation sequencing technologies, researchers are gathering more data than they are able to process and analyze. One of the most widely performed analyses is identifying regions of similarity between DNA or protein sequences using the Basic Local Alignment Search Tool, or BLAST. Due to the large amount of sequencing data produced, parallel implementations of BLAST are needed to process the data in a timely manner. While these implementations have been designed for those researchers with access to computing grids, recent web-based services, such as Amazon's Elastic Compute Cloud, now offer scalable, pay-as-you-go computing. In this paper, we present K-mulus, an application that performs distributed BLAST queries via Hadoop MapReduce using a collection of established parallelization strategies. In addition, we provide a method to speedup BLAST by clustering the sequence database to reduce the search space for a given query. Our results show that users must take into account the size of the BLAST database and memory of the underlying hardware to efficiently carry out the BLAST queries in parallel. Finally, we show that while our database clustering and indexing approach offers a significant theoretical speedup, in practice the distribution of protein sequences prevents this potential from being realized.

Keywords: Bioinformatics, Cloud computing, Sequence alignment, Hadoop

Faster GPU-accelerated Smith-Waterman Algorithm with Alignment Backtracking for Short DNA Sequences

Yongchao Liu, Bertil Schmidt

Institut für Informatik, Johannes Gutenberg Universität Mainz

55099 Mainz, Germany

{liuy, bertil.schmidt}@uni-mainz.de

In this paper, we present a GPU-accelerated Smith-Waterman (SW) algorithm with Alignment Backtracking, called GSWAB, for short DNA sequences. This algorithm performs all-to-all pairwise alignments and retrieves optimal local alignments on CUDA-enabled GPUs. To facilitate fast alignment backtracking, we have investigated a tile-based SW implementation using the CUDA programming model. This tiled computing pattern enables us to more deeply explore the powerful compute capability of GPUs. We have evaluated the performance of GSWAB on a Kepler-based GeForce GTX Titan graphics card. The results show that GSWAB can achieve a performance of up to 56.8 GCUPS on large-scale datasets. Furthermore, our algorithm yields a speedup of up to 53.4 and 10.9 over MSA-CUDA (the first stage) and gpu-pairAlign on the same hardware configurations.

Keywords: Smith-Waterman, Alignment backtracking, GPU, CUDA

Accelerating string matching on MIC architecture for motif extraction

Solon Pissis^{1,2}, Christian Goll², Pavlos Pavlidis³, Alexandros Stamatakis²

¹Florida Museum of Natural History, University of Florida, USA

²Heidelberg Institute for Theoretical Studies, Germany

³Foundation for Research & Technology – Hellas, Greece

{solon.pissis, alexandros.stamatakis}@h-its.org

Identifying repeated factors that occur in a string of letters or common factors that occur in a set of strings represents an important task in computer science and biology. Such patterns are called *motifs*, and the process of identifying them is called *motif extraction*. In biology, motifs may correspond to functional elements in DNA, RNA, or protein molecules. In this article, we orchestrate MoTeX, a high-performance computing tool for MoTif eXtraction from large-scale datasets, on Many Integrated Core (MIC) architecture. MoTeX uses state-of-the-art algorithms for solving the fixed-length approximate string-matching problem. It comes in three flavors: a standard CPU version; an OpenMP version; and an MPI version. We compare the performance of our MIC implementation to the corresponding CPU version of MoTeX. Our MIC implementation accelerates the computations by a factor of ten compared to the CPU version. We also compare the performance of our MIC implementation to the corresponding OpenMP version of MoTeX running on modern Multicore architecture. Our MIC implementation accelerates the computations by a factor of two compared to the OpenMP version.

Keywords: motif extraction, high-performance computing, MIC architecture

A Parallel, Distributed-memory Framework for Comparative Motif Discovery

Dieter De Witte¹, Michiel Van Bel^{2,3}, Pieter Audenaert¹, Piet Demeester¹,
Bart Dhoedt¹, Klaas Vandepoele^{2,3}, Jan Fostier¹

¹Department of Information Technology (INTEC), Ghent University - iMinds
Gaston Crommenlaan 8 bus 201, Ghent, Belgium

²Department of Plant Systems Biology, VIB, Technologiepark 927, Ghent, Belgium

³Department of Plant Biotechnology and Bioinformatics, Ghent University
Technologiepark 927, Ghent, Belgium

{dieter.dewitte,pieter.audenaert}@intec.ugent.be

{piet.demeester,bart.dhoedt,jan.fostier}@intec.ugent.be

{michiel.vanbel,klaas.vandepoele}@psb.vib-ugent.be

The increasing number of sequenced organisms has opened new possibilities for the computational discovery of cis-regulatory elements ('motifs') based on phylogenetic footprinting. Word-based, exhaustive approaches are among the best performing algorithms, however, they pose significant computational challenges as the number of candidate motifs to evaluate is very high. In this contribution, we describe a parallel, distributed-memory framework for de novo comparative motif discovery. Within this framework, two approaches for phylogenetic footprinting are implemented: an alignment-based and an alignment-free method. The framework is able to statistically evaluate the conservation of motifs in a search space containing over 160 million candidate motifs using a distributed-memory cluster with 200 CPU cores in a few hours. Software available from <http://bioinformatics.intec.ugent.be/blsspeller/>

Keywords: Motif discovery, phylogenetic footprinting, parallel computing, distributed-memory

Parallel seed-based approach to protein structure similarity detection

Guillaume Chapuis¹, Mathilde Le Boudic - Jamin¹, Rumen Andonov¹, Hristo Djidjev²,
Dominique Lavenier¹

¹INRIA/ IRISA Rennes, GenScale, Rennes, France

²Los Alamos National Laboratory, Los Alamos NM, USA

{guillaume.chapuis, mathilde.le_boudic-jamin}@irisa.fr

{rumen.andonov, dominique.lavenier}@irisa.fr

djidjev@lanl.gov

Finding similarities between protein structures is a crucial task in molecular biology. Many tools exist for finding an optimal alignment between two proteins. These tools, however, only find one alignment even when multiple similar regions exist. We propose a new parallel heuristic-based approach to structural similarity detection between proteins that discovers multiple pairs of similar regions. We prove that returned alignments have RMSDc and RMSDd lower than a given threshold. Computational complexity is addressed by taking advantage of both fine- and coarse-grain parallelism.

Keywords: protein structure comparison, parallel computing, seed-based, heuristic, alignment graph

Monitoring Performance and Power for Application Characterization with Cache-aware Roofline Model

Leonel Sousa
Technical University of Lisbon
Lisbon, Portugal
las@inesc-id.pt

Accurate on-the-fly characterization of application behavior requires assessing a set of execution-related parameters at runtime. These parameters do not only refer to performance, but also to power and energy consumption, which can be obtained by relying on built-in hardware measurement facilities in modern multi-core architectures, such as performance and energy counters. However, current OSs do not provide the means to directly obtain this characterization data, and the user still needs to rely on complex custom-built libraries with limited capabilities, which might introduce significant execution and measurement overheads. In this keynote, two different techniques are proposed for efficient performance, power and energy monitoring of multi-core systems. The proposed monitoring facilities allow capturing the run-time behavior of a wide range of applications at two different system levels, i.e., one facility at the user-space level and the other which is directly integrated into the OS kernel, thus examining the real application execution from completely different aspects. Although the importance of the proposed monitoring facilities is patent for many purposes, this presentation is focused on their employment for application characterization by relying on the recently proposed Cache-aware Roofline model.

Keywords:

Energy and Deadline Constrained Robust Stochastic Static Resource Allocation

Mark Oxley¹, Sudeep Pasricha^{1,2}, Howard Siegel^{1,2}, Anthony Maciejewski¹

¹Department of Electrical and Computer Engineering

Colorado State University, Fort Collins, Colorado 80523, USA

²Department of Computer Science

Colorado State University, Fort Collins, Colorado 80523, USA

{mark.oxley, sudeep,hj,aam}@colostate.edu

In this paper, we study the problem of energy and deadline constrained static resource allocation where a collection of independent tasks (“bag-of-tasks”) is assigned to a heterogeneous computing system. Computing systems often operate in environments where task execution times vary (e.g., due to data dependent execution times), therefore we model the execution time of tasks stochastically. This research focuses on the design of energy-constrained resource allocation heuristics that maximize robustness against the uncertainties in task execution times. We design and evaluate a new resource allocation heuristic based on Tabu Search that employs dynamic voltage and frequency scaling (DVFS) and exploits heterogeneity by incorporating novel local search techniques.

Keywords: heterogeneous computing, static resource allocation, power-aware computing, DVFS, robustness

Performance and Energy Analysis of the Iterative Solution of Sparse Linear Systems on Multicore and Manycore Architectures

José I. Aliaga, Maribel Castillo, Juan C. Fernández, Germán León, Joaquín Pérez, Enrique S. Quintana-Ortí

Dpto. de Ingeniería y Ciencia de Computadores, Universidad Jaume I
12.071–Castellón, Spain

{aliaga,castillo,jfernand,leon,al001566,quintana}@uji.es

In this paper we investigate the performance-energy balance of a collection of concurrent architectures, from general-purpose and digital signal multicore systems to graphics processors (GPUs), representative of current technology. This analysis employs the conjugate gradient method, an important algorithm for the iterative solution of linear systems that is basically composed of the sparse matrix-vector product and other (minor) vector kernels. To allow a fair comparison, and mimic the common case for many scientific applications, we leverage simple implementations of the numerical methods and underlying kernels, and thus rely only on those optimizations applied by the compiler.

Keywords: Energy efficiency, high-performance computing, sparse linear algebra, multicore processors, low-power processors, GPUs

Energy and Power Consumption Trends in HPC

Piotr Luszczek
University of Tennessee, USA

Over the past few years, the energy consumption and power constraints have started playing a major role in the use of HPC systems. They have also become one of the major design constraints for building the future hardware for the hundred Peta- and the Exa-flop/s supercomputers. This talk will track some of the technological trends, that were the effect of this transition and are likely to continue into the future. Adapting to this dramatic shift in hardware constraints is believed to be one of the major issues for porting the essential applications of computational science and have them run efficiently on the upcoming systems. The examples in the talk will provide examples of software solution for better utilization of energy and adaptation of the power constraints.

Keywords:

Measuring the Sensitivity of Graph Metrics to Missing Data

Anita Zakrzewska, David A. Bader
Georgia Institute of Technology
azakrzewska3@gatech.edu

The increasing energy consumption of high performance computing has resulted in rising financial and environment costs. On embedded systems, power supply is a major limiting factor in performance. As a result, reducing the energy consumption of computation is an emerging area of interest. We use an approach of data filtering to reduce the memory accesses and energy costs of sparse graph algorithms. Such a filtering approach clearly causes errors in resulting graph metrics and can only be used with tolerable error levels. Therefore, in order to be a feasible method, the effects of missing or filtered data on error must be determined. We measure this effect by performing sensitivity analysis for various graph algorithms and analyzing the trade-off between energy consumption reduction and graph metric error.

Keywords: graph algorithms, sensitivity analysis, missing data, energy consumption, power

The Energy/Frequency Convexity Rule: Modeling and Experimental Validation on Mobile Devices

Karel De Vogeleer¹, Gerard Memmi¹, Pierre Jouvelot², Fabien Coelho²

¹TELECOM ParisTech – INFRES – CNRS LTCI - UMR 5141 – Paris, France

²MINES ParisTech – Centre de recherche en informatique – Fontainebleau, France

{karel.devogeleer, gerard.memmi}@telecom-paristech.fr

{pierre.jouvelot, fabien.coelho}@mines-paristech.fr

This paper provides both theoretical and experimental evidence for the existence of an Energy/Frequency Convexity Rule, which relates energy consumption and CPU frequency on mobile devices. We monitored a typical smartphone running a specific computing-intensive kernel of multiple nested loops written in C using a high-resolution power gauge. Data gathered during a week-long acquisition campaign suggest that energy consumed per input element is strongly correlated with CPU frequency, and, more interestingly, the curve exhibits a clear minimum over a 0.2GHz to 1.6GHz window. We provide and motivate an analytical model for this behavior, which fits well with the data. Our work should be of clear interest to researchers focusing on energy usage and minimization for mobile devices, and provide new insights for optimization opportunities.

Keywords: DVFS, energy measurements, power measurements, energy optimization

Simulations of the adsorption behavior of dendrimers

Jaroslav Kłos^{1,3}, Jens -Uwe Sommer^{1,2}

¹Leibniz Institute of Polymer Research Dresden e. V.

01069 Dresden, Germany

²Institute for Theoretical Physics, Technische Universität Dresden

01069 Dresden, Germany

³Faculty of Physics, A. Mickiewicz University

Umultowska 85, 61-614 Poznań, Poland

klosj@o2.pl

Using Monte Carlo simulations we study adsorption of dendrimers with flexible spacers onto a flat surface in a wide range of molecular weight, N , generation number, G , spacer length, S , and the monomer-surface interaction strength parameter, τ . Our calculations indicate that for large values of N the dendrimers exist in three τ -dependent regions referred to as non-adsorbed, critical and adsorbed. Slightly below the critical point of adsorption, τ_c , a weakly adsorbed state is approached in which the molecules stick to the surface and are spherical in shape. By further lowering τ below a spacer-length dependent value, $\tau^*(S) < \tau_c$, a jump-like transition into a strongly adsorbed state occurs. Here, the dendrimers become flat and their lateral size is described by a 2D mean-field model.

Keywords: polymer, dendrimer, adsorption, simulation

An optimized Lattice Boltzmann code for BlueGene/Q

Marcello Pivanti¹, Filippo Mantovani², Sebastiano Fabio Schifano¹,
Raffaele Tripiccione¹, Luca Zenesini¹

¹Università di Ferrara and INFN, Ferrara, Italy

²Facultät für Physik, Universität Regensburg, Germany

`schifano@fe.infn.it`

In this paper we describe an optimized implementation of a Lattice Boltzmann (LB) code on the BlueGene/Q system, the latest generation massively parallel system of the BlueGene family. We consider a state-of-art LB code, that accurately reproduces the thermo-hydrodynamics of a 2D-fluid obeying the equations of state of a perfect gas. The regular structure of LB algorithms offers several levels of algorithmic parallelism that can be matched by a massively parallel computer architecture. However the complex memory access patterns associated to our LB model make it not trivial to efficiently exploit all available parallelism. We describe our implementation strategies, based on previous experience made on clusters of many-core processors and GPUs, present results and analyze and compare performances.

Keywords: Lattice Boltzmann, High performance applications, Massively parallel architectures, Performance analysis

A Parallel and Scalable Iterative Solver for Sequences of Dense Eigenproblems Arising in FLAPW

Mario Berljafa¹, Edoardo Angelo Di Napoli^{2,3}

¹Department of Mathematics, Faculty of Science, University of Zagreb
Bijenička cesta 30, 10000–Zagreb, Republic of Croatia

²Jülich Supercomputing Centre, Forschungszentrum Jülich
Wilhelm-Johnen straße, 52425–Jülich, Germany

³Aachen Institute for Advance Study in Computational Engineering Science
Schinkelstraße 2, 52072 Aachen, Germany

mberljaf@student.math.hr, e.di.napoli@fz-juelich.de

In one of the most important methods in Density Functional Theory – the Full-Potential Linearized Augmented Plane Wave (FLAPW) method – dense generalized eigenproblems are organized in long sequences. Moreover each eigenproblem is strongly correlated to the next one in the sequence. We propose a novel approach which exploits such correlation through the use of an eigensolver based on subspace iteration and accelerated with Chebyshev polynomials. The resulting solver, parallelized using the Elemental library framework, achieves excellent scalability and is competitive with current dense parallel eigensolvers.

Keywords: Chebyshev polynomials, Subspace Iteration, Eigenproblem Sequence, Density Functional Theory, Elemental

Sequential Monte Carlo in Bayesian assessment of contaminant source localization based on the sensors concentration measurements

Anna Wawrzynczak-Szaban^{1,2}, Piotr Kopka^{1,3}, Mieczysław Borysiewicz¹

¹National Centre for Nuclear Research, Świerk-Otwock, Poland

²Institute of Computer Science, Siedlce University, Poland

³Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland

awawrzynczak@ap.siedlce.pl

Accidental atmospheric releases of hazardous material pose great risks to human health and the environment. In this context it is valuable to develop the emergency action support system, which can quickly identify probable location and characteristics of the release source based on the measurement of dangerous substance concentration by the sensors network. In this context Bayesian approach occurs as a powerful tool being able to combine observed data along with prior knowledge to gain a current understanding of unknown model parameters. We have applied the methodology combining Bayesian inference with Sequential Monte Carlo (SMC) to the problem of the atmospheric contaminant source localization. The algorithm input data are the on-line arriving concentrations of given substance registered by the downwind distributed sensor's network.

We have proposed the different version of the Hybrid SMC along with Markov Chain Monte Carlo (MCMC) algorithms and examined its effectiveness to estimate the probabilistic distributions of atmospheric release parameters. The proposed algorithms scan 5-dimensional parameters' space searching for the contaminant source coordinates, release strength and atmospheric transport dispersion coefficients.

Keywords: Bayesian inference, stochastic reconstruction, MCMC, SMC

Effective parallelization of quantum simulations: nanomagnetic molecular rings

Piotr Kozłowski¹, Grzegorz Musiał¹, Michał Antkowiak¹, Dante Gatteschi²

¹Faculty of Physics, Adam Mickiewicz University

ul. Umultowska 85, PL-61-614, Poznań, Poland

²Department of Chemistry, INSTM, and Università di Firenze

via della Lastruccia 3, I-50019 Sesto Fiorentino, Italy

gmusial@amu.edu.pl

The effective parallelization of processing exploiting the MPI library for the numerically exact quantum transfer matrix (QTM) and exact diagonalization (ED) deterministic simulations of chromium-based rings is proposed. In the QTM technique we have exploited parallelization of summation in the partition function. The efficiency of the QTM calculations is above 80% up to about 1000 processes. With our test programs we calculated low temperature torque, specific heat and entropy for the chromium ring Cr₈ exploiting realistic Hamiltonian with single-ion anisotropy and the alternation of the nearest neighbor exchange couplings. Our parallelized ED technique makes use of the self-scheduling scheme and the longest processing time algorithm to distribute and diagonalize separate blocks of a Hamiltonian matrix by slave processes. Its parallel processing scales very well, with efficiency above 90% up to about 10 processes only. This scheme is improved by processing more input data sets in one job which leads to very good scalability up to arbitrary number of processes. The scaling is improved for both techniques when larger systems are considered.

Keywords: parallelization of processing, MPI, numerical simulations, nanomagnetic rings, Heisenberg model

DFT study of the Cr₈ molecular magnet within chain-model approximations

V. Bellini^{1,2}, D. M. Tomecka³, Bartosz Brzostowski⁴, M. Wojciechowski⁴,
F. Troiani¹, F. Manghi¹, M. Affronte¹

¹CNR-INFM-National Research Center on nanoStructures and bioSystems
at Surfaces (S3), Via Campi 213/A, 41100 Modena, Italy

²Istituto di Struttura della Materia (ISM) - Consiglio Nazionale delle Ricerche
(CNR), I-34149, Trieste, Italy

³Faculty of Physics, A. Mickiewicz University
ul. Umultowska 85, 61-614 Poznań, Poland

⁴Institute of Physics, University of Zielona Góra
ul. Prof. Szafrana 4a, 65-516 Zielona Góra, Poland
brzost@proton.if.uz.zgora.pl

We present a density functional theory (DFT) study of the electronic and magnetic properties of the Cr₈ molecular ring. The all-electron linearized augmented plane wave method (LAPW) implemented in the Wien2k package and pseudopotential method implemented in SIESTA package are used to calculate the electronic states, exchange coupling parameters of an infinite chain model system of Cr₈. We demonstrate how, under opportune modifications to the ring cycle structure, different one-dimensional chain models can be devised, with the capability of mimicking with good approximation the electronic and magnetic properties of the original Cr₈ molecule. Such models offer an unique opportunity, in virtue of the reduced computational effort, to carry out extensive investigations of a whole set of molecules belonging to the Cr-based molecular rings family.

Keywords: density functional theory, Cr antiferromagnetic rings, molecular nanomagnets

The effect of parallelization on a tetrahedral mesh optimization method

Domingo Benítez, Eduardo Rodríguez, Jose M. Escobar, Rafael Montenegro
University of Las Palmas de Gran Canaria
SIANI Institute, Spain
{dbenitez, erodriguez, jmescobar, rmontenegro}@siani.es

In this paper a new parallel algorithm for simultaneous untangling and smoothing of tetrahedral meshes is proposed. This algorithm is derived from a mathematical mesh optimization method. We provide a detailed analysis of parallel scalability, load balancing, parallelism bottlenecks, and influence of graph coloring algorithms on the performance of the parallel algorithm for six benchmark meshes with a wide range of sizes. We demonstrate that this algorithm is highly scalable when run on a high-performance shared-memory many-core computer with up to 128 Itanium 2 processors. However, some parallel deterioration is observed. Here, we analyze the causes of this parallel deterioration using experimental results obtained with performance counter hardware. Finally, we compare the impact of three previously published graph coloring algorithms on the performance of our parallel algorithm.

Keywords: graph meshes, numerical methods, optimization, geometrical problems and computations, performance evaluation, parallel computing

Analysis of Partitioning Models and Metrics in Parallel Sparse Matrix-Vector Multiplication

Kamer Kaya¹, Bora Uçar¹, Ümit Çatalyürek^{1,3}

¹Dept. of Biomedical Informatics, The Ohio State University, USA

²CNRS and LIP, ENS Lyon, France

³Dept. of Electrical & Computer Engineering, The Ohio State University, USA

{kamer, umit}@bmi.osu.edu, bora.ucar@ens-lyon.fr

Graph/hypergraph partitioning models and methods have been successfully used to minimize the communication among processors in several parallel computing applications. Parallel sparse matrix-vector multiplication (SpMxV) is one of the representative applications that renders these models and methods indispensable in many scientific computing contexts. We investigate the interplay of the partitioning metrics and execution times of SpMxV implementations in three libraries: Trilinos, PETSc, and an in-house one. We carry out experiments with up to 512 processors and investigate the results with regression analysis. Our experiments show that the partitioning metrics influence the performance greatly in a distributed memory setting. The regression analyses demonstrate which metric is the most influential for the execution time of the libraries.

Keywords: parallel sparse-matrix vector multiplication, hypergraph partitioning, regression analysis

Achieving Memory Scalability in the Gysela Code to Fit Exascale Constraints

Guillaume Latu, Jean Roman, Fabien Rozar

¹IRFM, CEA Cadarache, FR-13108 Saint-Paul-les-Durance, France

²Maison de la simulation, CEA Saclay, FR-91191 Gif sur Yvette, France

³INRIA, Université de Bordeaux, CNRS, FR-33405 Talence, France

`fabien.rozar@cea.fr`

Gyrokinetic simulations lead to huge computational needs. Up to now, the semi-Lagrangian code Gysela performed large simulations using a few thousands cores (65k cores). But to understand more accurately the nature of the plasma turbulence, finer resolutions are wished which make Gysela a good candidate to exploit the computational power of future Exascale machines. Among the Exascale challenges, the less memory per core issue is one of the most critical. This paper deals with memory management in order to reduce the memory peak, and presents an approach to understand the memory behaviour of the application when dealing with very large meshes. This enables us to extrapolate the behaviour of Gysela for expected capabilities of Exascale machine.

Keywords: Exascale, Memory scalability, Memory footprint reduction, Plasma Physics

Probabilistic analysis of barrier eliminating method applied to load-imbalanced parallel application

Naoki Yonezawa¹, Ken'ichi Katou¹, Issei Kino¹, Koichi Wada²

¹Kanagawa University, Hiratsuka, Kanagawa, Japan

²University of Tsukuba, Tsukuba, Ibaraki, Japan

yonezawa@info.kanagawa-u.ac.jp

In order to reduce the overhead of barrier synchronization, we have proposed an algorithm which eliminates barrier synchronizations and evaluated its validity experimentally in our previous study. As a result, we have found that the algorithm is more effective to the load-imbalanced program than load-balanced program. However, the degree of the load balance has not been discussed quantitatively. In this paper, we model the behavior of parallel programs. In our model, the execution time of a phase contained in a parallel program is represented as a random variable. To investigate how the degree of the load balance influences the performance of our algorithm, we varied the coefficient of variation of probability distribution which the random variable follows. Using the model, we evaluated the execution time of parallel programs and found that theoretical results are consistent with experimental ones.

Keywords: barrier elimination, probabilistic analysis, time reduction

Multi-GPU parallel memetic algorithm for capacitated vehicle routing problem

Mieczysław Wodecki, Wojciech Bożejko, Michał Karpiński, Maciej Pacut
Institute of Computer Science, University of Wrocław
Joliot-Curie 15, 50-383 Wrocław, Poland
{mwd, michal.karpinski, maciej.pacut}@ii.uni.wroc.pl

The goal of this paper is to propose and test a new memetic algorithm for the capacitated vehicle routing problem in parallel computing environment. In this paper we consider a simple variation of the vehicle routing problem in which the only parameter is the capacity of the vehicle and each client only needs one package. We analyse the efficiency of the algorithm using the hierarchical Parallel Random Access Machine (PRAM) model and run experiments with code written in CUDA.

Keywords: Multi-GPU, Vehicle Routing Problem, Metaheuristics

Parallel Applications Performance Evaluation using the Concept of Granularity

Jan Kwiatkowski
Wrocław University of Technology
Wrocław, Poland
`jan.kwiatkowski@pwr.wroc.pl`

With the advent of multi-core processors and the growing popularity of local clusters installations, better understanding of parallel applications behavior becomes a necessity. It causes the necessity of evaluation the parallel application performance. It can be argued that the raising popularity of parallelization results in the dare need of methods and tools capable of automatic analysis and prediction of parallel applications efficiency. On the other hand performance evaluation can be employed to raise the usefulness of all the parallel and distributed processing environments. The results of the evaluation are suitable for resource allocation, load balancing, on-demand reservation of dynamic resources, quality of service negotiations, runtime and wait time predictions, etc. In traditional methods of performance evaluation the results are based on wall-time measurements. This approach requires consecutive application executions or, when the detailed application profile is created, involves a time-consuming data analysis. In the paper an alternative approach is proposed. Utilizing the decomposition of execution time, a separate analysis of the computation time and overheads related to parallel execution can be used to calculate the granularity of application and then determine the analyzed application efficiency. The accuracy of techniques based on internal and external measurements is evaluated and compared to the results achieved with the classical analysis.

Keywords: parallel processing, parallel applications, performance evaluation

Improving Perfect Parallelism

Lars Karlsson, Carl Christian Kjelgaard Mikkelsen, Bo Kågström
Department of Computing Science
Umeå University, Sweden
{larsk, spock, bokg}@cs.umu.se

We reconsider the familiar problem of executing a perfectly parallel workload consisting of N independent tasks on a parallel computer with $P \ll N$ processors. We show that there are memory-bound problems for which the runtime can be reduced by the forced parallelization of individual tasks across a small number of cores. Specific examples include solving differential equations, performing sparse matrix–vector multiplications, and sorting integer keys.

Keywords: perfectly parallel problem, resource contention, forced parallelization

Parallel One-Sided Jacobi SVD Algorithm with Variable Blocking Factor

Martin Bečka, Gabriel Okša
Institute of Mathematics, Slovak Academy of Sciences
Bratislava, Slovak Republic
{martin.becka, gabriel.oksa}@savba.sk

Parallel one-sided block-Jacobi algorithm for the matrix singular value decomposition (SVD) requires an efficient computation of symmetric Gram matrices, their eigenvalue decompositions (EVDs) and an update of matrix columns and right singular vectors by matrix multiplication. In our recent parallel implementation with p processors and blocking factor $\ell = 2p$, these tasks are computed serially in each processor in a given parallel iteration step because each processor contains exactly two block columns of an input matrix A . However, as shown in our previous work, with increasing p (hence, with increasing blocking factor) the number of parallel iteration steps needed for the convergence of the whole algorithm increases linearly but faster than proportionally to p , so that it is hard to achieve good speedup. To achieve the better parallel performance, we propose to break the tight relation $\ell = 2p$ and to use a small blocking factor $\ell = p/k$ for some integer k that divides p , ℓ even. In this case the algorithm works with pairs of logical block columns that are distributed among processors so that all computations inside a parallel iteration step are themselves parallel, i.e., they are implemented using some block cyclic matrix distribution on some processor grid. We discuss the optimal data distribution for parallel subproblems in the one-sided block-Jacobi algorithm and analyze its computational complexity. Working with small blocking factors is advantageous in modern parallel systems that are constructed hierarchically from individual CPUs and computational nodes, whereby the amount of memory per node is relatively small. First results of numerical experiments are presented that confirm our theoretical analysis.

Keywords: parallel one-sided Jacobi SVD algorithm, variable blocking factor, distributed parallel computer

Using Intel Xeon Phi coprocessor to accelerate computations in MPDATA algorithm

Lukasz Szustak¹, Krzysztof Rojek¹, Pawel Gepner²

¹Institute of Computer and Information Science

Czestochowa University of Technology

Dabrowskiego 69, 42-201 Czestochowa, Poland

² Intel Corporation

{krojek, lszustak}@icis.pcz.pl, pawel.gepner@intel.com

The multidimensional positive definite advection transport algorithm (MPDATA) is one of the major part of the dynamic core of EULAG model. The MPDATA algorithm belongs to the group of nonoscillatory forward in time algorithms, and performs a sequence of stencil computations.

Keywords: EULAG model, MPDATA advection algorithm, Intel Xeon Phi coprocessor, multi-/manycore programming, OpenMP, adaptation

Genetic Programming in Automatic Discovery of Relationships in Computer System Monitoring Data

Włodzimierz Funika^{1,2}, Paweł Koperek¹

¹AGH University of Science and Technology

Faculty of Computer Science, Electronics and Telecommunications

Department of Computer Science

al. Mickiewicza 30, 30-059 Kraków, Poland

¹ACC CYFRONET AGH, ul. Nawojki 11, 30-950 Kraków, Poland

funika@agh.edu.pl, pkoperek@gmail.com

Modern computer systems have become very complex. Analyzing and modifying them requires substantial experience and knowledge. To make administrative tasks easier, automated methods for discovering relationships between system components are required. In this paper we discuss the use of genetic programming as a method for identification of meaningful relationships between computer system components. We present our implementation of evolutionary computations environment and compare it with an already existing solution. Next we analyze results of a sample experiment and share our conclusions. The final section provides directions for future work.

Keywords: distributed systems, automatic modeling, monitoring, genetic programming, automatic system management

Genetic Algorithms Execution Control Under a Global Application State Monitoring Infrastructure

Adam Smyk, Marek Tudruj
Polish-Japanese Institute of Information Technology
Warsaw, Poland
{asmyk, tudruj}@pjwstk.edu.pl

We present a new approach to the design of parallel applications genetic algorithms executed in multicore systems with core clustering. The proposed system is based on the use of global parallel program control functions and asynchronous process/thread internal execution control. Global application states monitoring is provided by generalized synchronization processes called synchronizers which collect local states of program elements, construct consistent global states, control predicates and send control signals to program computational elements. Such control/synchronization mechanisms are provided as ready to use program infrastructure which is structurally decoupled from computational code and can implemented hierarchically for threads, processes and groups of threads or processes. Construction and management of global program states for the design of the global program execution control at both thread and process levels is less troublesome and less error prone. A parallel genetic algorithm designed based on the global control infrastructure is discussed. The algorithm has been used for optimal core mesh partitioning for solving the FDTD numerical problem. To increase an efficiency of the genetic algorithm, its execution will be performed and controlled simultaneously with the FDTD computations.

Keywords: distributed program design paradigms, application global states monitoring, genetic algorithms

Non-perturbative methods in phenomenological simulations of ring-shape molecular nanomagnets

Piotr Kozłowski¹, Grzegorz Musiał¹, Monika Haglauer¹, Wojciech Florek¹,
Michał Antkowiak¹, Filippo Esposito², Dante Gatteschi³

¹Faculty of Physics, Adam Mickiewicz University
ul. Umultowska 85, PL-61-614, Poznań, Poland

²Dipartimento di Scienze Fisiche, Università di Napoli
Piazzale V. Tecchio 80, I-80125 Napoli, Italy

³Department of Chemistry, INSTM and Università di Firenze
via della Lastruccia 3, I-50019 Sesto Fiorentino, Italy
gmusial@amu.edu.pl

Two non-perturbative numerically exact methods: exact diagonalization and quantum transfer matrix are applied to computationally complex Heisenberg-like spin models of ring shaped molecular nanomagnets and implemented in the high performance computing environment. These methods are applicable to the wide class of ring-shaped nanomagnets. For the hypothetical antiferromagnetic nanomagnet Ni_{12} the influence of single-ion anisotropy on the ground states is investigated. For Cr_8 it is demonstrated that the alternation of the nearest-neighbor bilinear exchange couplings leads to small changes in the magnetic torque with respect to the uniformly coupled system. Specific heat and entropy for Cr_8 are showed to be good indicators of crossing fields. The applicability of the Lande rule to both systems is checked.

Keywords: molecular nanomagnet, quantum transfer matrix, exact diagonalization, Heisenberg Hamiltonian, magnetic torque, high performance computing

Non-uniform quantum spin chains: static and dynamic properties

Artur Barasiński¹, Bartosz Brzostowski¹, Ryszard Matysiak², Pawel Sobczak³, Dariusz Woźniak¹

¹Institute of Physics, University of Zielona Góra
Prof. Z. Szafrana 4A, 65-516 Zielona Góra, Poland

²Institute of Engineering and Computer Education, University of Zielona Góra,
Prof. Z. Szafrana 4, 65-516 Zielona Góra, Poland

³Faculty of Physics, A. Mickiewicz University
ul. Umultowska 85, PL-61-614 Poznań, Poland
b.brzostowski@proton.if.uz.zgora.pl

Since magnetic materials are often composed of magnetically isolated chains, their magnetic properties can be modeled by the one-dimensional quantum Heisenberg model. The quantum transfer matrix (QTM) method based on a checkerboard structure has been applied. In order to increase the length of the transfer matrix in Trotter direction we apply the density-matrix renormalization technique and check the efficiency of parallelization for a part of the code: the construction of the transfer matrix. Moreover, using the Matrix Product States representation the time evolution of the ground-state magnetization has been performed after the sudden change in applied field.

Keywords: static and dynamic magnetic properties, Heisenberg model, numerical simulations, parallel processing

Neighborhood selection and rules identification for cellular automata: a rough sets approach

Bartłomiej Placzek
Institute of Computer Science
University of Silesia
Będzińska 39, 41-200 Sosnowiec, Poland
placzek.bartlomiej@gmail.com

In this paper a method is proposed which uses data mining techniques based on rough sets theory to select neighborhood and determine update rule for cellular automata (CA). According to the proposed approach, neighborhood is detected by reducts calculations and a rule-learning algorithm is applied to induce a set of decision rules that define the evolution of CA. Experiments were performed with use of synthetic as well as real-world data sets. The results show that the introduced method allows identification of both deterministic and probabilistic CA-based models of real-world phenomena.

Keywords: rough sets, cellular automata, model identification

The Graph of Cellular Automata Applied for Modelling Tumour Induced Angiogenesis

Paweł Topa
Institute of Computer Sciences
AGH University of Science and Technology,
al. Mickiewicza 30, 30-059 Kraków, Poland
topa@agh.edu.pl

Angiogenesis is the process of formation of vascular network. Blocking tumour induced angiogenesis is one of the treatments applied in oncology. Research involving computer simulations looking for the rules influencing the structure of vascular network and its functionality. This paper summarizes the applications of Graph of Cellular Automata modelling tool, developed by the Author, for modelling Tumour Induced Angiogenesis. Vascular network which is modelled by the graph interacts with surrounding tissue represented by the lattice of automata. The network develops and reorganizes accordingly to locally acting factors (stimulators and inhibitors). The model includes blood flow calculations in a modelled vascular network.

Keywords: cellular automata, tumour inducing angiogenesis, computer modelling

Preconditioning iterative substructuring methods using inexact local solvers

Piotr Krzyżanowski
Institute of Applied Mathematics, University of Warsaw
Banacha 2, 02-097 Warszawa, Poland
`piotr.krzyzanowski@mimuw.edu.pl`

We consider block preconditioning for iterative substructuring algorithms with inexact subdomain solvers, including incomplete Cholesky and V-cycle multigrid. Numerical results show that block triangular preconditioners are very competitive and in many cases outperform presently used preconditioners based on full block triangular decomposition.

Keywords: preconditioning, iterative substructuring, domain decomposition, block preconditioner

An Efficient Representation on GPU for Transition Rate Matrices for Markov Chains

Jaroslaw Bylina, Beata Bylina, Marek Karwacki
Institute of Mathematics
Marie Curie-Skłodowska University
Lublin, Poland
`jmbylina@hektor.umcs.lublin.pl`

In the paper the authors present a novel modification of the HYB format — known from the CUSP library. The new format is suitable for sparse Markovian transition rate matrices. The improvement is based on the analysis of the matrix. The aim of this new format is to enable bigger matrices to be processed on one GPU, but improving computation performance at the same time — on one and two GPUs in comparison to CPU. Particularly, the SpMV operation — that is the multiplication of a sparse matrix by a vector — is analyzed for this format on one GPU and two GPUs. Numerical experiments for transition rate matrices of Markov chains show that the proposed format allows to process matrices of sizes about 3.6×10^7 rows with the use of one GPU (3GB RAM) and gives the speedup up to 13 times in comparison to multithreaded CPU with the use of the CSR format (SpMV from the MKL library on CPU).

Keywords: GPU, SpMV, Markov chains, transition rate matrix, CUSP, sparse matrices

Eigen-G: GPU-based eigenvalue solver for real-symmetric dense matrices

Toshiyuki Imamura^{1,3}, Susumu Yamada^{2,3}, Masahiko Machida^{2,3}

¹RIKEN Advanced Institute for Computational Science, 7-1-26

Minatojima-minami-machi, Chuo-ku, Kobe-shi, Hyogo 650-0047, Japan

²CCSE, Japan Atomic Energy Agency, Kashiwa-shi, Chiba 277-8587, Japan

³CREST, Japan Science and Technology Agency, Japan

imamura.toshiyuki@riken.jp

This paper presents a performance review of Eigen-G, which is a GPU-based eigenvalue solver for real-symmetric real matrices. We examine that Eigen-G outperforms other GPU-based eigensolver implementation such as `magma_dsyevd` included in the MAGMA library. Selection of well-optimized CUDA BLAS libraries and CPU-GPU hybrid scheduling yield performance improvement. For example, Eigen-G diagonalizes an eight thousand dimensional matrix by approximately 18 seconds.

Keywords: GPU-based eigenvalue solver, GPGPU, multicore, CPU-GPU hybrid scheduling

A study on adaptive algorithms for numerical quadrature on hybrid GPU and multicore based systems

Giuliano Laccetti¹, Marco Lapegna¹, Valeria Mele¹, Diego Romano²

¹Department of Mathematics and Applications, University of Naples Federico II
Complesso Universitario Monte S. Angelo, Via Cintia, Naples, Italy

²ICAR-CNR, Via P. Castellino 111, Naples, Italy

{giuliano.laccetti, marco.lapegna, valeria.mele}@unina.it
diego.romano@na.icar.cnr.it

In this work, a parallel adaptive algorithm for the computation of a multidimensional integral on hybrid GPU and multicore based systems is described. Two different strategies have been combined together in the algorithm: the first procedure is responsible for the load balancing among the threads on the multicore CPU, while the second one is responsible for an efficient execution of the computational kernel on GPU. The performance is analyzed and experimental results have been presented for a system with a quad-core CPU and two GPUs.

Keywords: Hierarchical parallelism, hybrid algorithms, adaptive algorithms, multidimensional integration

The definition of interval-valued intuitionistic fuzzy sets in the framework of Dempster-Shafer theory

Ludmila Dymova, Pavel Sevastjanov
Institute of Computer and Information Science
Czestochowa University of Technology
Dabrowskiego 69, 42-201 Czestochowa, Poland
{dymova, sevast}@icis.pcz.pl

In this report, a critical analysis of conventional operations on interval-valued intuitionistic fuzzy values (IVIFVs) and their applicability to the solution of multiple criteria decision making (MCDM) problems in the interval-valued intuitionistic fuzzy setting are presented. It is shown that the classical definition of Atanassov's interval-valued intuitionistic fuzzy set (A-IV IFS) may lead to controversial results. Therefore, a new more constructive definition of A-IV IFS is proposed. It is shown that this new definitions makes it possible to present IVIFVs in the framework of interval-extended Dempster-Shafer theory of evidence (DST) as belief intervals with bounds presented by belief intervals.

Keywords: Atanassov's interval-valued intuitionistic fuzzy set, Interval-valued intuitionistic fuzzy values Interval extended zero method, Interval-extended Dempster-Shafer theory