Preface

Eduard Mehofer*

Research Group Scientific Computing,
University of Vienna,
Waehringer Strasse 29/6.20, 1090, Vienna, Austria
E-mail: eduard.mehofer@univie.ac.at
*Corresponding author

Markus Schordan

Institute of Computer Science,
UAS Technikum Wien,
Hoechstaedtplatz 5, 1200, Vienna, Austria
E-mail: schordan@technikum-wien.at

Dan Quinlan

Lawrence Livermore National Laboratory,
Box 808, L-422, Livermore, CA 94551-0808, USA
E-mail: quinlan1@llnl.gov

Beniamino Di Martino

Department of Industrial and Information Engineering,
Second University of Naples,
Real Casa dell'Annunziata,
via Roma n. 29 – 81031, Aversa (CE), Italy
E-mail: Beniamino.dimartino@unina.it

Biographical notes: Eduard Mehofer is an Associate Professor at the University of Vienna, Austria. He conducts research in parallel computing with focus on parallel programming models and languages, parallelising compilers, and runtime techniques to adapt programmes to hardware and application characteristics. Currently, his main research interests are programming support and optimisation techniques for multi-core and accelerator-based architectures. He is the Vice-Director of the research group Scientific Computing at the University of Vienna. Prior to joining the University of Vienna, he was the Project Leader at the Alcatel Research Center. In 2001, he was a Visiting Professor at the University of Arizona, AZ, and in 2005 he was a Guest Professor at the Indian Institute of Science (IISc) Bangalore, India. He served as the Chairman, member of programme committees, Guest Editor, editorial board member, and referee for numerous international journals, conferences and workshops.

Markus Schordan is the Deputy Programme Director of Multimedia and Software Engineering as well as Game Engineering and Simulation at the University of Applied Sciences Technikum Wien. He previously held positions at University Klagenfurt (1997 to 2001), Lawrence Livermore National Laboratory (2001 to 2003), and TU Vienna (2003 to 2008). In 2009, he received an R&D 100 AWARD (ROSE) and in 2011 a Best Paper Award at the 11th IEEE Source-Code and Manipulation Conference (SCAM 2011). His research interests include programme analysis, programming languages, high-performance computing, and software analysis tools. He received his PhD degree from University Klagenfurt in 2001 (with distinction). He served as a Co-Chair of several workshops and as member in Programme Committees of 15+ conferences and workshops.

Dan Quinlan is the Project Leader for the ROSE Compiler Project at Lawrence Livermore National Laboratory. His research is in numerous areas that intersect computer science and mathematics. His research interests include: source code and binary analysis, high performance computing, parallel computation, semantics-based source code transformations and optimisations, language specific compiler tools/infrastructure/design, cache-based optimisations, object oriented abstractions, parallel load balancing algorithms, object-oriented numerical frameworks, parallel adaptive mesh refinement, and parallel multigrid algorithms. In 2009, he received an R&D 100 Award (ROSE). He was previously a Research Staff at Los Alamos National Laboratory (1993 to 1998). He earned his PhD in Computational Mathematics from University of Colorado in 1993.

Beniamino Di Martino is a Full Professor at the Second University of Naples, Italy. He is the author of seven international books and more than 200 publications in international journals and conferences. He is the Project Coordinator of the EU funded FP7-ICT Project 'mOSAIC'. He is a member of the IEEE Working Group on Cloud Interoperability, of the Cloud Standards Customer Council, of the Cloud Computing Experts' Group of European Commission. He served as a General and Programme Chairman, and member in Programme Committees, of several international conferences, and as Guest Editor for several journals' special issues. He is Editor, editorial board member and Chair of international journals. His research interests include: cloud computing, high performance computing and architectures, knowledge discovery and management, semantic web services, mobile and intelligent agents and mobile computing, reverse engineering, automated programme analysis and transformation, and programme comprehension.

GPUs are cost-effective platforms for computational intensive applications providing tremendous peak performance. However, it is a major challenge to deliver the intrinsic performance of such architectures to end applications. Although GPU programming has been simplified with the appearance of CUDA and OpenCL, it is still a difficult task to develop efficient GPU programmes. To achieve a high occupancy level of GPUs many hardware details like memory size, number of threads, or number of registers have to be considered by programmers. Therefore, programming support and optimisation techniques are of key importance on such systems.

This special issue collects extended versions of the best papers presented to the International Workshop on GPUs and Scientific Applications (GPUScA 2011) held in Galveston, Texas, in October 2011, in conjunction with PACT 2011, the Annual International Conference on Parallel Architectures and Compilation Techniques. The main topics of the workshop were programming approaches and key techniques to leverage the computing power of GPUs.

The paper 'OpenMPC: extended OpenMP for efficient programming and tuning on GPUs', by Seyong Lee and Rudolf Eigenmann, Purdue University (USA), proposes a directive-based, high-level programming model which addresses programmability and tunability issues on GPGPUs. OpenMPC is an OpenMP-to-CUDA translation and optimisation framework which comprises a standard OpenMP API plus a new set of directives and environment variables to control the translation process.

The paper 'Fast and memory-efficient minimum spanning tree on the GPU', by Scott Rostrup, Shweta Srivastava, and Kishore Singhal, Synopsis Inc. (Canada), presents a data-parallel implementation of Kruskal's algorithm for finding the minimum spanning tree. The approach exemplifies the difficulties of mapping irregular algorithms to GPUs.

The paper 'Comparing the performance of stochastic simulation on GPUs and OpenMP', by Weijun Xiao, Peng Li, and David J. Lilja, University of Minnesota (USA), introduces a GPU-based solution and an OpenMP-based solution for simulating stochastic computing and evaluates the performance behaviour. A stochastic implementation of the frame difference-based image segmentation algorithm is used as an example to conduct extensive experiments.

The paper 'ForOpenCL: transformations exploiting array syntax in Fortran for accelerator programming', by Matthew J. Sottile, Craig E. Rasmussen, Wayne N. Weseloh, Robert W. Robey, Dan Quinlan and Jeffrey Overbey, Lawrence Livermore National Laboratory, University of Illinois, Galois Inc., and Los Alamos National Laboratory (USA), presents a programming methodology that provides for Fortran programmers an easy access to GPU systems. An implicitly parallel subset of Fortran-like arrays using assignment statements and binary operators is taken and transformed to OpenCL kernels.

The paper 'Iterative statistical kernels on contemporary GPUs', by Thilina Gunarathne, Bimalee Salpitikorala, Arun Chauhan, and Geoffrey Fox, Indiana University (USA), presents an experimental evaluation of three important algorithms used in iterative statistical applications for large scale data processing on different GPU platforms using OpenCL. The algorithms are multi-dimensional scaling, PageRank, and K-means clustering.

We wish to thank the members of the programme committee of the workshop, and the additional reviewers, who helped us in selecting the best papers for this special issue.

Programme committee

Scott Baden	University of California, USA
Uday Bondhugula	Indian Institute of Science, IND
Zhigang Deng	University of Houston, USA
Thomas Fahringer	University of Innsbruck, AUT
Dieter Kranzlmüller	Ludwig Maximilian University of Munich, GER
Beniamino Di Martino	Second University of Naples, ITA
Naoya Maruyama	Tokyo Institute of Technology, JAP
Eduard Mehofer	University of Vienna, AUT
Nacho Navarro	Barcelona Supercomputing Center, SPA
Michael O'Boyle	University of Edinburgh, GBR
Craig Rasmussen	Los Alamos National Laboratory, USA

Bernhard Scholz	University of Sydney, AUS
Markus Schordan	University of Applied Sciences Technikum Wien, AUT
Michael Schwarz	Arizona State University, USA
Xipen Shen	College of William and Mary, USA
Matthew Sottile	Galois Inc., USA
Apan Qasem	Texas State University, USA
Dan Quinlan	Lawrence Livermore National Laboratory, USA
Richard Vuduc	Georgia Tech College of Computing, USA
Michael Wimmer	Vienna University of Technology, AUT
Qing Yi	University of Texas at San Antonio, USA

Additional reviewers

Javier Cabezas

Isaac Gelado

Klaus Kofler

Stephan Reiter