Preface

John T. Feo and Antonino Tumeo*

Pacific Northwest National Laboratory,
High Performance Computing,
902 Battelle Blvd, MSIN J4-30,
Richland, WA 99352, USA
E-mail: john.feo@pnnl.gov
E-mail: antonino.tumeo@pnnl.gov
*Corresponding author

Oreste Villa

NVIDIA Research,
2701 San Tomas Express Way, 95050 Santa Clara, CA, USA
E-mail: ovilla@nvidia.com

Simone Secchi

ARM,
110 Fulbourn Rd., Cambridge CB1 9NJ, UK
E-mail: simone.secchi@arm.com

Biographical notes: John T. Feo is the Principal Investigator of the High Performance Data Analytic Project and Deputy Division Leader at the Pacific Northwest Laboratory. He received his PhD in Computer Science from The University of Texas at Austin. He began his career at Lawrence Livermore National Laboratory, where he was a Principal Investigator of the Sisal Language Project. He then joined Tera Computer Company as Principal Engineer and Product Manager for the MTA. After a short 'sabbatical' at Microsoft, where he led a software group developing a next-generation virtual reality platform, he joined PNNL. His research interests are parallel programming, multithreaded architectures, functional languages, and performance studies.

Antonino Tumeo received his MSc in Informatic Engineering in 2005, and PhD in Computer Engineering in 2009, from Politecnico di Milano in Italy. Since 2011, he has been a Research Scientist at Pacific Northwest National Laboratory (PNNL). He joined PNNL in 2009 as a Post-Doctoral Research Associate. Previously, he was a Post-Doctoral Researcher at Politecnico di Milano. His research interests are modelling and simulation of high performance architectures, hardware-software codesign, FPGA prototyping and GPGPU computing.

Oreste Villa is a Senior Research Scientist at NVIDIA Research in the Architecture Research Group (ARG). Previously, from 2008 to 2013, he was a Research Scientist at Pacific Northwest National Lab (PNNL) in the high performance computing group. He received his PhD in Computer Science from Politecnico di Milano (Italy) in 2008 and MSc in Electronic Engineering in 2003 from the University of Cagliari (Italy). He also received his ME in Embedded Systems Design from the University of Lugano (Switzerland) in 2004.

Simone Secchi received his MSc in Electronic Engineering in 2007 and PhD in Electronic Engineering in 2011 from University of Cagliari, Italy. He is currently a Senior Application Engineer at ARM UK. Previously, he was a Post-Masters Research Associate from 2010 to 2011 and a Postdoctoral Research Associate until 2012 at the Pacific Northwest National Laboratory. His main research interests include modelling and parallel software simulation of high performance computing architectures, FPGA-based energy-aware emulation of multiprocessor systems and advanced network-on-chip architectures.

Many emerging applications, such as tools for analysing social, communication or security networks, bioinformatics, machine learning software packages, graph databases, and, in general, data-intensive knowledge discovery tools, are irregular. These applications employ pointer- or linked-list data structures, which usually cause fine-grained, unpredictable, data accesses. They often exhibit irregularity in the control flow and/or in communication, and present high synchronisation intensity. They may also have significant amounts of inherent dynamic parallelism, which

enables spawning many new concurrent activities, while the algorithms explore the data structures. However, they may have very large datasets, often unstructured and difficult to partition in a balanced way, which may not fit even on the most modern and bigger fat nodes, and which may make parallelism complex to manage.

For these reasons, these applications seem unfit for current processor and accelerator architectures, designed with high-flop ratings, complex cache hierarchies and growing numbers of cores to provide high performance and efficiency with regular, arithmetic intensive, and easily partitionable workloads. Current high performance computing clusters are further optimised for these types of workloads, typical of many scientific simulations, by integrating network interconnects that reach very high bandwidths only with large, batched data transfers, but perform very poorly with fine-grained transactions.

Addressing the class of irregular applications will become more and more important for next generation high performance systems, which will need to answer complex questions by discovering knowledge from exponentially growing datasets.

The solutions needed to address the challenges brought by these applications can only come by considering the problem from all perspectives: from micro-to-system-architectures, from compilers to languages, from libraries to runtimes, from algorithm design to data characteristics. Only collaborative efforts among researchers with different expertise, including end users, domain experts, and computer scientists, could lead to significant breakthroughs.

This special section presents three interesting papers that examine and propose solutions to some of the problems posed by irregular applications from three different points of view.

The first paper presents an interesting case study of irregular application: the clustering of real world graphs constructed out of biological data using parallel computers. The paper presents the design and the evaluation of two different parallel implementations of a popular serial graph clustering heuristic. The first implementation targets shared memory multicore platforms; the second is designed for distributed memory clusters using Hadoop MapReduce.

The second paper looks at approaches to accelerate irregular applications on reconfigurable devices. In particular, it presents CHAT, an extension of the ROCCC compiler, which generates custom hardware accelerators synthesisable on FPGA devices starting from C code. CHAT implements mechanisms to generate concurrent hardware threads for FPGAs, trying to mask the latency of the unpredictable data accesses typical of these applications.

Finally, the third paper presents a parallel algorithm for the optimal binary search tree problem. The authors introduce a coarse grain multicomputer (CGM) parallel algorithm, which reduces the number of broadcast transfers among processors with respect to previous parallel formulations of the algorithm.

We believe that this special section represents a first, important, step towards increasing awareness of the issues and challenges related to irregular applications. We hope that it will foster the research around these themes, strengthening the community. We are sure that you will enjoy reading it, and thank you for your attention.