# Editorial


## Gabriele Mencagli*

Department of Computer Science,
University of Pisa (UNIPI),
Largo B. Pontecorvo, 3,
I-56127, Italy
Email: mencagli@di.unipi.it
*Corresponding author

## Cristiana B. Bentes

Departamento de Engenharia de Sistemas e Computação,
Faculdade de Engenharia,
Universidade do Estado do Rio de Janeiro,
Rio de Janeiro, RJ, Brazil
Email: cris@eng.uerj.br

## Leandro A.J. Marzulo

Departamento de Informática e Ciência da Computação,
Instituto de Matemática e Estatística,
Universidade do Estado do Rio de Janeiro,
Rio de Janeiro, RJ, Brazil
Email: leandro@ime.uerj.br

## Mauricio Lima Pilla

Centre for Technological Development,
Universidade Federal de Pelotas (UFPel),
1st Gomes Cameiro Street, Pelotas,
RS 96010-610, Brazil
Email: pilla@inf.ufpel.edu.br



**Biographical notes:** Gabriele Mencagli is an Assistant Professor at the Computer Science Department of the University of Pisa, Italy. He received his PhD from the University of Pisa in 2012. He is co-author of more than 50 peer-reviewed papers in international conferences, workshops and journals, and of one book. His research interests are in the area of parallel and distributed systems and data stream processing. He is a member of the Editorial Board of *Future Generation Computer Systems* (Elsevier) and *Cluster Computing* (Springer).

Cristiana B. Bentes received her BSc degree in Mathematics from State University of Rio de Janeiro, Brazil, and her MSc and DSc degrees in Systems Engineering and Computer Science at Federal University of Rio de Janeiro. She is currently a Full Professor in the Department of Systems Engineering and Computer Science at State University of Rio de Janeiro. Her currently research interests include high-performance computing, parallel programming, novel parallel architectures and cloud computing.

Leandro A.J. Marzulo received his MSc (2007) and DSc (2011) in Systems Engineering and Computer Science from Universidade Federal do Rio de Janeiro, Brazil. He is currently an Associate Professor at Universidade do Estado do Rio de Janeiro and has been working for the past 10 years in dataflow execution models for high performance computing. His areas of interest are computer architecture, dataflow computing, parallel programming, compilers and distributed algorithms.

Mauricio Lima Pilla is an Associate Professor at the Computer Science courses in the Center for Technological Development of the Federal University of Pelotas, Brazil. He is a member of the Laboratory of Ubiquitous and Parallel Systems (LUPS) at UFPEL, with research interests in the fields of parallel and distributed systems, computer architectures, and quantum computing simulation.

Writing parallel applications is still a tedious and cumbersome task, which however is mandatory if practitioners or researchers want to fully exploit the potential of modern multi-core processors. This task becomes even harder as different computation devices, such as General Purpose Graphic Processing Units (GPGPUs), Field Programmable Gate Arrays (FPGAs), and other many-core accelerators, are employed to build heterogeneous systems. This imposes new challenges to the scientific community: the creation of *unconventional models* to ease the parallelism exploitation and reduce the programming burden by freeing the programmer to explicitly cope with the low-level details of the underlying physical devices.

Another important programming issue to consider, especially in applications that run on large systems and manipulate big datasets, is the trade-off between moving data to a remote powerful processing platform supporting very high levels of parallelism (e.g., on cloud systems) and applying the computations locally, near the data sources, in order to reduce communication and energy costs while keeping satisfactory performance levels on more limited resources. *Edge*, *fog* and *in-situ* computing paradigms (Lopez et al., 2015; Bonomi et al., 2012) intend to tackle this issue by adding computing capabilities to network devices (such as network interface cards, switches and routers), storage devices or even memory. Those 'smart' devices could perform part of the computation and would reduce data transmission over the network. Furthermore, the current trend related to the *Internet-of-Things* (IoT) predicts that an increasing number of devices, such as sensors, actuators and processing elements, will be spread around the world, forming a huge network to collect information relevant to a broad range of applications. All the aforementioned aspects make computing systems even more heterogeneous, intensifying the need for novel programming models with properly designed high-level abstractions.

This special issue on 'Advancements in High-level Parallel Programming Models for Edge/Fog/In-situ Computing' aimed at attracting researchers interested in presenting contributions about the evolution of existing models or in proposing novel ones, considering the trend of processing/accelerator devices in the context of edge/fog/in-situ computing. This Special Issue is based on extended versions of the best papers presented at two workshops, MPP 2017 and WAMCA 2017. After a rigorous evaluation process with up to three rounds of reviews, five papers have been accepted and finally included in this special issue.

Edge/fog/in-situ computing has been gaining a huge interest, with initiatives such as the OpenFog Consortium (see https://www.openfogconsortium.org/) and the Internet of Things Consortium (see https://iofthings.org/). Therefore, we hope that the included papers will be able to attract readers from both academia and industry, including independent research groups and universities.

## *Themes of this special issue*

This special issue covers a set of distinct topics about the definition of new parallel programming models merging high-level abstractions for parallel computing and features targeting edge/fog/in-Situ computing environments.

The first paper, titled 'HPSM: a programming framework to exploit multi-CPU and multi-GPU systems simultaneously', presents a high-level C++ framework to exploit multi-CPU and multi-GPU systems. The approach enables parallel loops and reduction tasks executed simultaneously on CPUs and GPUs by synergically using both the components of a heterogeneous system, that is supporting a fully hybrid processing model.

The second paper, titled 'An efficient pathfinding system in FPGA for edge/fog computing', investigates parallel pathfinding algorithms, which are widely used in GPS navigation, artificial intelligence and in IoT environments. The proposal is a parallel approach to the problem written using a FPGA, in order to effectively mitigate the network traffic problem with a suitable implementation to be placed at the edge of the network.

The goal of the third paper, titled 'A network coding protocol for wireless sensor fog computing', is to propose a new network coding technique for efficient data transmission protocols on wireless environments. The experimental analysis conducted by the authors demonstrates the suitability of the approach for Fog computing scenarios.

The paper 'A dataflow runtime environment and static scheduler for edge, fog and in-situ computing' follows the renewed interest in the dataflow programming paradigm in new highly distributed contexts. The approach extends the features of the Sucuri framework, a Python library for dataflow programming. The approach and its scheduler have been completely re-designed in order to support edge/fog/in-situ environments.

Finally, the last paper, titled 'An optimised dataflow engine for GPGPU stream processing', focuses on stream processing applications with high-demanding performance requirements that are hard to tackle with traditional programming models. The approach extends an existing Java framework for dataflow programming by adding high-level constructs to offload the processing of kernels on GPU devices, thus merging the stream processing domain with SIMD architectures and programming models.

## References

Bonomi, F., Milito, R., Zhu, J. and Addepalli, S. (2012) 'Fog computing and its role in the internet of things', *Proceedings of the 1st Edition of the MCC Workshop on Mobile Cloud Computing (MCC'12)*, ACM, New York, NY, USA, pp.13–16. DOI: http://dx.doi.org/10.1145/2342509.2342513

Lopez, P.G., Montresor, A., Epema, D., Datta, A., Higashino, T., Iamnitchi, A., Barcellos, M., Felber, P. and Riviere, E. (2015) 'Edge-centric computing: vision and challenges', *SIGCOMM Computer Communication Review*, Vol. 45, No. 5, pp.37–42. DOI: https://doi.org/10.1145/2831347.2831354