

MAURIZIO DROCCO

Department of Computer Science ◊ Corso Svizzera 185, 10149 Torino, Italy

drocco@di.unito.it ◊ +39 328 2553406

<http://www.di.unito.it/~drocco>

SUMMARY

Maurizio Drocco is a Ph.D. Candidate in Computer Science at University of Turin. He is going to discuss his thesis about distributed global pointers on October 20, 2017.

He is participating in supporting parallel and distributed programming in the C++ standard.

He has been Research Intern at IBM Thomas J. Watson Lab (NY) and at IBM Dublin Research Lab, and he is research associate at University of Torino since 2009.

He has co-authored papers in international journals and peer-reviewed conference proceedings (Google h-index 7).

His research focuses on high-level parallel programming for high-performance computing, in particular models and methods for heterogeneous platforms.

EDUCATION

University of Torino January 2014 – October 2017
Ph.D. in Computer Science, defending the thesis “Parallel Programming with Global Asynchronous Memory: Models, C++ APIs and Implementations” (Supervisor: prof. M. Aldinucci).

University of Torino July 2013
Master’s Degree in Computer Science (summa cum laude), defending the thesis “Parallel stochastic simulators in systems biology: the evolution of the species” (Supervisor: prof. M. Aldinucci). The thesis has been selected for publication [11] and got a “Best Student Award”.

University of Torino March 2010
Bachelor Degree in Computer Science (summa cum laude), defending the thesis “A Simulator for a new Calculus of Biological Processes (SCWC: Stochastic Calculus of Wrapped Compartments)” (Supervisor: prof. E. Giovannetti).

EXPERIENCE

University of Torino September 2009 – present
Research Associate Torino, IT

- member of parallel computing *alpha* group (<http://alpha.di.unito.it>)
- participant in €3.5M H2020 *RePhrase* project (<http://rephrase-ict.eu/>)
- participant in €3.5M EU-FP7 *ParaPhrase* project (<http://paraphrase-ict.eu/>)
- participant in €3.5M EU-FP7 *REPARA* project (<http://repara-project.eu/>)
- participant in Regione Piemonte *BioBITS* project (<http://www.biobits.di.unipmn.it/>)
- contributor of FastFlow parallel programming framework [6]
- co-developer of a real-time video restoration filter [5], presented at Nvidia GTC 2014 [7]
- lead developer of CWC simulator tool [1]

IBM Thomas J. Watson Research Center May 2015 – November 2015
Research Intern Yorktown Heights, NY, USA

- member of the High Performance Analytics group
- co-developer of optimized tools for parallel graph processing (<http://www.graph500.org/>)

- member of the High Performance Computing group
- lead developer of a hardware/software auto tuning framework for structured parallel applications on IBM POWER platforms

RESEARCH INTERESTS

Parallel programming architectures and tools for High Performance Computing:

- High-performance runtime support for parallel computing frameworks, lock-free synchronization, deployment onto hybrid CPU/GPU/distributed architectures
- Pattern-based parallel programming models and languages
- Semi-automatic parallelization of sequential code, semi-automatic optimization and tuning of parallel applications

RESEARCH ACTIVITY

IBM Thomas J. Watson Research Center

May 2015 – November 2015

High Performance Analytics group

- My activity during the internship, with Fabio Checconi and Fabrizio Petrini as mentors, has been twofold. First, I implemented various optimizations for high performance graph processing, culminated in a fast multicore implementation of the breadth-first search, contextually to the graph500 chart; in particular, I developed a compressed graph representation as a data structure optimized for the IBM BlueGene/Q memory system. Second, I co-developed an emulator of a distributed-shared memory system for massively parallel, non-coherent multiprocessors, targeting hybrid clusters including multicore, GPU, FPGA and flash storage devices; in particular, the emulator has been implemented as a plugin for the runtime instrumentation tool valgrind, providing minimal overhead (i.e., about 20x slowdown) with respect to traditional simulators.

REPARA project

October 2013 – present

Reengineering and Enabling Performance And poweR of Applications

- I extended the FastFlow runtime with parallel patterns for heterogeneous platforms, based on hybrid C++/OpenCL code. This work culminated in the Loop-of-Stencil-Reduce pattern, which design and implementation has been published in [3].
- I worked on the parallelization of NuChart-II, a C++ tool for graph-based Hi-C data processing. Employed techniques include working set reduction, optimized memory allocators, and cache-friendly data structures with reduced impact on memory bandwidth. This work has been published e.g. in [13].

IBM Dublin Research Lab

July 2013 – November 2013

High Performance Computing group

- I developed a framework for combined hardware-software semi-automatic tuning of FastFlow parallel patterns on POWER architectures. As enabling features, I exploited configurable hardware parameters provided by IBM POWER7/7+ architectures (e.g. SMT level, cache pre-fetching) and FastFlow structured parallel programming. The framework allows tuning for specific performance metrics, including throughput, latency, and energy efficiency.

ParaPhrase project

October 2011 – March 2015

Parallel Patterns for Adaptive Heterogeneous Multicore Systems

- I extended the FastFlow runtime with parallel patterns for heterogeneous platforms, based on hybrid C++/CUDA code.

- As a test bed, I developed a high-performance filter for high-quality, high-frequency video restoration. The filter has been presented at Nvidia GTC 2014 [7] and published e.g. in [5].

BioBITs

September 2009 – December 2012

Developing White and Green Biotechnologies by Converging Platforms from Biology and Information Technology towards Metagenomics

- I explored parallelization of systems biology simulators, culminating in a C++ parallel simulator, on top of FastFlow. I parallelized the whole simulation-mining pipeline at different levels, from pattern-matching vectorization with Intel SSE intrinsics to fine-grained scheduling of simulation quanta based on simulated time, guaranteeing pseudo-optimal load balancing. Related works have been published in [1, 8, 10, 9].

PUBLICATIONS AND CONFERENCES

I have co-authored more than 20 papers in international conference proceedings and scientific journals. I have been participating in conferences as author, speaker and member of organizing and program committees.

A complete list of my publications is available at: <http://www.di.unito.it/~drocco>.

Journals

- *Journal of Super Computing* [18, 2].
- *Parallel Processing Letters* [15].
- *International Journal of High Performance Computing Applications* [5, 17].
- *Frontiers in Genetics* [14].
- *BioMed Research International* [1].
- *Briefings in Bioinformatics* [8].
- *Theoretical Computer Science* [10].
- *Transactions on Computational Systems Biology* [9].

Conferences

- *15th IEEE International Symposium on Parallel and Distributed Processing with Applications*, 2017. Guangzhou, China. Member of **Program Committee** for the “Technologies and Tools” track .
- *16th IEEE International Conference on Scalable Computing and Communications*, 2016. Toulouse, France. Member of **Program Committee** for the “Extreme scalable systems and applications” track.
- *24th Euromicro International Conference on Parallel Distributed and network-based Processing*, 2016. Heraklion, Greece. Member of **Program Committee** for the “Advances in High-Performance Bioinformatics, Systems and Synthetic Biology” special session, **Session Chair** and **Speaker** [12].
- *23rd Euromicro International Conference on Parallel Distributed and network-based Processing*, 2015. Turku, Finland. Member of **Program Committee** for the “Advances in High-Performance Bioinformatics, Systems and Synthetic Biology” special session.
- *34th IEEE International Conference on Distributed Computing Systems Workshops*, 2014. Madrid, Spain. **Speaker** [4]
- *11th International meeting on Computational Intelligence methods for Bioinformatics and Biostatistics*, 2014. Cambridge, UK. **Speaker** [16]
- *22nd International Euromicro Parallel Distributed and network-based Processing*, 2014. Turin, Italy. Member of **Organizing Committee**.

SCHOOLS

UPMARC Multicore Computing Summer School

organized by UPMARC project

June 2016

Uppsala, Sweden

- Lectures on parallel programming models, memory systems and multicore computing
- I delivered the seminar “Sequential Consistency and other Myths” at CS department in Turin

International Summer School in Parallel Patterns

organized by EU-FP7 ParaPhrase project

June 2014

NCI, Dublin, Ireland

- Lectures on skeleton-based parallel programming, FastFlow, large-scale industrial HPC applications
- Hands-on sessions and tutorials

10th International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems

organized by HiPEAC consortium

July 2014

Fiuggi, Italy

- Lectures on foundational aspects of parallel computing, OpenCL compilation on heterogeneous platforms, low-energy aspects in HPC applications
- I Presented the poster “A Dynamic Memory Allocator for heterogeneous platforms”

ACHIEVEMENTS AND AWARDS

IBM Research

Ph.D. Scholarship Award

June 2015

Torino, Italy

- 20,000\$ Scholarship award for accomplishment reached during the internship at IBM Thomas J. Watson Research Center.

University of Torino

Best Student Award

October 2014

Torino, Italy

- I am the recipient of the “Best Student Award” in Computer Science Master programme 2012/13 of University of Torino.

GRANTS

Ph.D. three years Grant (2014 – 2017)

Italian Ministry of Education, Universities and Research

October 2013

Torino, Italy

- Ranked 1st in the competitive selection at Computer Science Department of University of Torino

10th International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems (ACACES)

organized by HiPEAC consortium

July 2014

Fiuggi, Italy

- Full registration fee for selected Ph.D. students from HiPEAC institutions

List of Publications

- [1] M. Aldinucci, C. Calcagno, M. Coppo, F. Damiani, M. Drocco, E. Sciacca, S. Spinella, M. Torquati, and A. Troina. On designing multicore-aware simulators for systems biology endowed with on-line statistics. *BioMed Research International*, 2014.
- [2] M. Aldinucci, M. Danelutto, M. Drocco, P. Kilpatrick, C. Misale, G. Peretti Pezzi, and M. Torquati. A parallel pattern for iterative stencil + reduce. *Journal of Supercomputing*, pages 1–16, 2016.

- [3] M. Aldinucci, M. Danelutto, M. Drocco, P. Kilpatrick, G. Peretti Pezzi, and M. Torquati. The loop-of-stencil-reduce paradigm. In *Proc. of Intl. Workshop on Reengineering for Parallelism in Heterogeneous Parallel Platforms (RePara)*, pages 172–177, Helsinki, Finland, Aug. 2015. IEEE.
- [4] M. Aldinucci, M. Drocco, G. Peretti Pezzi, C. Misale, F. Tordini, and M. Torquati. Exercising high-level parallel programming on streams: a systems biology use case. In *Proc. of the 2014 IEEE 34th Intl. Conference on Distributed Computing Systems Workshops (ICDCS)*, Madrid, Spain, 2014. IEEE.
- [5] M. Aldinucci, G. Peretti Pezzi, M. Drocco, C. Spampinato, and M. Torquati. Parallel visual data restoration on multi-GPGPUs using stencil-reduce pattern. *International Journal of High Performance Computing Applications*, 29(4):461–472, 2015.
- [6] M. Aldinucci and M. Torquati. *FastFlow website*, 2009. <http://mc-fastflow.sourceforge.net/>.
- [7] M. Aldinucci, M. Torquati, M. Drocco, G. Peretti Pezzi, and C. Spampinato. Fastflow: Combining pattern-level abstraction and efficiency in GPGPUs. In *GPU Technology Conference (GTC 2014)*, San Jose, CA, USA, Mar. 2014.
- [8] M. Aldinucci, M. Torquati, C. Spampinato, M. Drocco, C. Misale, C. Calcagno, and M. Coppo. Parallel stochastic systems biology in the cloud. *Briefings in Bioinformatics*, 15(5):798–813, 2014.
- [9] M. Coppo, F. Damiani, M. Drocco, E. Grassi, M. Guether, and A. Troina. Modelling ammonium transporters in arbuscular mycorrhiza symbiosis. *Transactions on Computational Systems Biology (TCS)*, 6575(13):85–109, 2011.
- [10] M. Coppo, F. Damiani, M. Drocco, E. Grassi, E. Sciacca, S. Spinella, and A. Troina. Simulation techniques for the calculus of wrapped compartments. *Theoretical Computer Science*, 431:75–95, 2012.
- [11] M. Drocco. Parallel stochastic simulators in systems biology: the evolution of the species. Master’s thesis, Computer Science Department, University of Torino, Italy, July 2013.
- [12] M. Drocco, C. Misale, and M. Aldinucci. A cluster-as-accelerator approach for SPMD-free data parallelism. In *Proc. of Intl. Euromicro PDP 2016: Parallel Distributed and network-based Processing*, pages 350–353, Crete, Greece, 2016. IEEE.
- [13] M. Drocco, C. Misale, G. Peretti Pezzi, F. Tordini, and M. Aldinucci. Memory-optimised parallel processing of Hi-C data. In *Proc. of Intl. Euromicro PDP 2015: Parallel Distributed and network-based Processing*, pages 1–8. IEEE, Mar. 2015.
- [14] I. Merelli, F. Tordini, M. Drocco, M. Aldinucci, P. Liò, and L. Milanesi. Integrating multi-omic features exploiting Chromosome Conformation Capture data. *Frontiers in Genetics*, 6(40), 2015.
- [15] C. Misale, M. Drocco, M. Aldinucci, and G. Tremblay. A comparison of big data frameworks on a layered dataflow model. *Parallel Processing Letters*, 27(01):1740003, 2017.
- [16] F. Tordini, M. Drocco, I. Merelli, L. Milanesi, P. Liò, and M. Aldinucci. NuChart-II: a graph-based approach for the analysis and interpretation of Hi-C data. In C. D. Serio, P. Liò, A. Nonis, and R. Tagliaferri, editors, *Computational Intelligence Methods for Bioinformatics and Biostatistics - 11th International Meeting, CIBB 2014, Cambridge, UK, June 26-28, 2014, Revised Selected Papers*, volume 8623 of *LNCS*, pages 298–311, Cambridge, UK, 2015. Springer.
- [17] F. Tordini, M. Drocco, C. Misale, L. Milanesi, P. Liò, I. Merelli, M. Torquati, and M. Aldinucci. NuChart-II: the road to a fast and scalable tool for Hi-C data analysis. *International Journal of High Performance Computing Applications (IJHPCA)*, pages 1–16, 2016.

- [18] M. Torquati, G. Mencagli, M. Drocco, M. Aldinucci, T. De Matteis, and M. Danelutto. On dynamic memory allocation in sliding-window parallel patterns for streaming analytics. *Journal of Supercomputing*, 2017. To appear.