

Workflow vs. Dataflow for HPC: Concepts, challenges, and simulation

Ph.D. proposal

Executive summary

The future SKA radio-telescope poses severe challenges to the computing infrastructure that will process the produced data. The constraints require to design both an innovative and highly heterogeneous architecture, and dedicated data processing algorithms. The key difficulty resides in the need to co-develop the hardware and software together. The goal of this work is to contribute to a prototyping tool for the prospective profiling of dataflow applications during their development.

Key skills required: C++ programming; Software engineering.

Research Unit: IRISA laboratory, Rennes. Team Myriads.

Co-advisors:

- Martin Quinson, ENS Rennes. <http://people.irisa.fr/Martin.Quinson/>
- Frédéric Suter, CC-IN2P3. <https://cclab.pages.in2p3.fr/frederic.suter/index.html>

PhD funding granted by the Dark-Era ANR project <https://dark-era.pages.centralesupelec.fr/>

To apply to this internship, please send your resume and cover letter to martin.quinson@ens-rennes.fr and frederic.suter@cc.in2p3.fr.

Context

The future *Square Kilometer Array (SKA)* radio telescope poses unprecedented challenges to the underlying computational system. The instrument is expected to produce Terabytes of data per second, mandating on-site pre-processing to reduce the size of data to be transferred. But the electromagnetic noise of a traditional computing center would hinder the quality of the measurements if located near to the instrument. As a result, the Science Data Processor (SDP) pipeline will only have an energy budget of only 1 MWatt for 250 Petaflops to execute in real time its complex algorithm chain over more than 7 Terabytes of data produced each second by the instrument.

Such energy and computation requirements imply the SDP to be an innovative dataflow oriented and heterogeneous architecture. On the hardware side, this supercomputer will combine standard HPC systems with dedicated accelerators such as GPU, FPGA, or components such as the Kalray Massively Parallel Processor Array (MPPA). One crucial challenge is to assess the performance both in time and energy of new complex scientific dataflow algorithms on not-yet-existing complex computing infrastructures. This will hardly be possible without efficient co-design methods and rapid prototyping tools.

The envisioned work will contribute to SimSDP, a prototyping tool for SKA-like dataflow applications. This tool shall provide early analyses in terms of memory usage, latency, throughput, and energy consumption. It will be based on two existing tools: PREESM [1], to evaluate the performance of heterogeneous single nodes, and SimGrid [2], to simulate inter-node communications. Adequately combining these tools is necessary to tackle the challenges posed to the SDP infrastructure through the evaluation of new algorithms leveraging the potential of low power accelerators.

Work plan

One first goal will be to extend the workflow application model [3] exhibited in SimGrid to allow the expression of dataflow applications [4] such as the those processed on the SDP. These models are conceptually close as they both exhibit task and data parallelism, but in a very different way. A first milestone will be to identify the common concepts and main differences between these two application models. The rest of the first year will be devoted to the adaptation, extension, and/or modification of the workflow-oriented programming interface of the SimGrid toolkit to enable the simulation of dataflow-oriented applications.

During the second year, the candidate will interface PREESM and SimGrid, to enable the simulation of dataflow-oriented applications on heterogeneous HPC platforms. PREESM will be used to predict the performance of each

task in the application graph, while SimGrid will be used to simulate the interaction between these elements. A first research question will be to find the right articulation between the practical concepts exhibited by PREESM and those of SimGrid. Then, as the algorithms envisioned for the SDP are hierarchical and multi-parametric and multiple potential configurations of the heterogeneous nodes have to be considered, the design space to explore will be extremely large. A second research question will then be about developing efficient techniques to explore this design space and minimize the number of simulations to perform.

Before writing the manuscript, the third year will be devoted to the development of a full-fledged simulator and the performance evaluation of specific numerical algorithms developed by the project partners in French Observatories [5].

Required skills

In addition to the skills that can reasonably be expected from Master-level students, the applicant should have solid skills in C++ programming (preferably under Linux), and some background in software engineering. This could be complemented or partially compensated by a very good background in concurrent and distributed algorithms. Good reading and writing skills in English are also mandated.

Benefits for the candidate

The envisioned work is expected to lead to several publications to top-level conferences and journals in computer science, such as SuperComputing, HPDC or IPDPS, constituting a strong academic file for the candidate. This scientific endeavor will naturally contribute to the development of the scientific and soft skills that any doctor should acquire. In addition, the applicant is also expected to acquire or improve several technical skills related to the development of actual HPC applications. Finally, the consortium of the Dark-Era project is highly interdisciplinary, gathering together computer scientists, specialists of applied mathematics and numerical algorithms, together with astrophysicists.

After this PhD work, the future doctor will be a natural candidate to classical research positions in academia, but they will also be a strong candidate to any R&D positions in the industry.

References

- [1] M. Pelcat, K. Desnos, J. Heulot, C. Guy, J.-F. Nezan, and S. Aridhi. PREESM: A Dataflow-based Rapid Prototyping Framework for Simplifying Multicore DSP Programming. In Proceedings of the 6th European Embedded Design in Education and Research Conference (EDERC), pages 36–40, 2014.
- [2] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter. Versatile, scalable, and accurate simulation of distributed applications and platforms. *Journal of Parallel and Distributed Computing*, 74(10):2899–2917, 2014.
- [3] E. Deelman, T. Peterka, I. Altintas, et al. The Future of Scientific Workflows. *The International Journal of High Performance Computing Applications*. 32(1):159-175, 2018.
- [4] E. A. Lee and D. G. Messerschmitt. Synchronous dataflow. In Proceedings of the IEEE 75(9):1235–1245, 1987.
- [5] C. Tasse, B. Hugo, M. Mirmont, O. Smirnov, M. Atemkeng, L. Bester, M. J. Hardcastle, R. Lakhoo, S. Perkins and T. Shimwel. Faceting for direction-dependent spectral deconvolution, *Astronomy & Astrophysics*, vol. 611, p. A87, Apr. 2018.