HPC-Big Data convergence at processing level
by bridging in-situ/in-transit processing with Big Data analytics

Advisors:

- Alexandru Costan - alexandru.costan@irisa.fr
- Gabriel Antoniu - gabriel.antoniu@inria.fr

Keywords

HPC, Big Data Analytics, in-situ/in-transit processing, stream processing, Spark

Subject

In more and more application areas, data volumes increase exponentially and are produced at increasing speed. Keeping pace with the overwhelming data volumes and velocity is critical, in order to allow meaningful extraction of insights from data and enable precise predictions and relevant decision making, despite this increasingly challenging context. To achieve this goal, the ability to perform precise analytics at extreme scales gains major importance, in a context where the scale at which data is produced and consumed is also increasing.

In the **High Performance Computing (HPC)** area, the need to get fast and relevant insights from massive amounts of data generated by extreme-scale computations led to the emergence of ***in-situ*** and ***in-transit*** **processing** approaches. They allow data to be visualized and processed in real-time, in an interactive way, as they are produced, as opposed to the traditional approach consisting of transferring data off-site after the end of the computation, for offline analysis. In the **Big Data Analytics (BDA)** area, the search for real-time, fast analysis was materialized through a different approach: **stream-based processing**. As the tools and cultures of HPC and BDA have evolved in divergent directions motivated by different optimization criteria, it becomes clear today that leveraging together the progresses achieved in the two areas can be an efficient means addresses the Big Data challenges, which are now relevant for both HPC and BDA.

Internship goal. This Master internship will be hosted by the KerData team at IRISA/Inria Rennes Bretagne Atantique(https://team.inria.fr/kerdata/). It aims to propose an approach enabling **HPC-Big Data convergence at the data processing level**, by exploring alternative solutions to architecture a unified framework for extreme-scale data processing. The architecture of such a framework will leverage the extreme scalability demonstrated by in situ/in transit data processing approaches originated in the HPC area, in conjunction with Big Data processing approaches emerged in the BDA area (batch-based, streaming-based and hybrid). The high-level goal of this framework is to enable the usage of a large spectrum of Big Data analytics techniques at extreme scales, to support precise predictions in real-time and fast decision making.

Enabling techniques. In the process of designing the unified data processing framework, we will leverage the **Damaris** framework [1,4] for scalable, asynchronous I/O and in-situ and in-transit visualization and processing, developed by KerData at Inria (https://project.inria.fr/damaris). Damaris already demonstrated its scalability up to 16,000 cores on some of the top supercomputers of Top500, including Titan, Jaguar and Kraken. For the purpose of this internship, **Damaris will have to be extended to support Big Data analytics for data processing** (e.g., based on the **Apache Flink** [2] and **Apache Spark** [3] engines and on their higher-level machine-learning libraries). The work will involve software architecture design, implementation and experimental validation.

In particular, adding programmatic support for Big Data analytics on dedicated cores in Damaris will focus on distributed, high-performance, always-available, elastic and accurate

data processing. Furthermore, we plan to design, implement and experimentally validate a mechanism for incrementally migrating running stream tasks from the in-situ processing backend to the in-transit one without stopping the query execution.

This work opens a huge opportunity: the possibility to substantially extend and enrich the knowledge already generated from **past data** (using traditional batch/stream processing with BDA) with knowledge based on **future data** (or simulated data, from HPC simulations). This can provide a richer tool for much deeper interpretation of measured real-time data, enabling more reliable, transparent and innovative decision making.

This subject will be done in close collaboration with Bogdan Nicolae and Pierre Matri from **Argonne National Laboratory (ANL, USA)**. In addition to the possibility to interact with top-level researchers and scientists from ANL, the work can involve the use of large-scale HPC experimental facilities available at ANL as well as Inria's Grid5000 distributed testbed - https://www.grid5000.fr/

Depending on the quality of the achievements, this internship could be followed by a Ph.D. thesis in the same area, in the same collaboration context.

Skills and abilities

- Programming skills (e.g. C/C++, Java, Python)
- Knowledge of computer networks and distributed systems
- Familiarity with Big Data Analytics or HPC is an advantage

Bibliography

1. Matthieu Dorier, Gabriel Antoniu, Franck Cappello, Marc Snir, Leigh Orf. Damaris: How to Efficiently Leverage Multicore Parallelism to Achieve Scalable, Jitter-free I/O. In Proc. CLUSTER – IEEE International Conference on Cluster Computing, Sep 2012, Beijing, China. URL: https://hal.inria.fr/hal-00715252
2. Carbone, Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., & Tzoumas, K. (2016). Apache Flink TM : Stream and Batch Processing in a Single Engine.
3. Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: cluster computing with working sets. In Proceedings of the 2nd USENIX conference on Hot topics in cloud computing (HotCloud'10). USENIX Association, Berkeley, CA, USA
4. The Damaris project. URL: https://project.inria.fr/damaris