Internship Proposal: I/O performance of the DDF pipeline

Collaboration between Inria - TADaaM and LAB

Funded by the NumPEx Exa-DoST project

October 2024

Introduction I - Parallel I/O for HPC

In high-performance computing (HPC) platforms, the supercomputers, applications running on the compute nodes access persistent data in a remote parallel file system (PFS), which is deployed over a set of dedicated servers. Popular examples of PFS are Lustre (https://www.lustre.org/) and BeeGFS (https://www.beegfs.io/c/) Each PFS server has one or more storage targets (OSTs), usually each associated with a different storage device. In these systems, files are broken into fixed-size stripes and distributed across the storage targets, so that different stripes can be accessed in parallel. The access to PFS is thus called parallel I/O [1].

This research field is important because there is a historical gap between processing and I/O speeds in HPC systems. Consequently, even if compute-intensive, many HPC applications spend a lot of their execution time on I/O operations, which prevents them from scaling.

Introduction II - The SKA project

The international SKA (Square Kilometre Array) project aims to build the largest radio astronomy observatory, with an effective collecting area of 1 km². SKA observatory will operate at radio frequencies between 0.3 GHz and 30 GHz and will be installed in two desert locations in Australia (low-frequency) and South Africa (mid-frequency). The observatory will continuously generate up to 10 Tb/s of raw data. This data will be analyzed on-site by the SKA “Science Data Processor” (SDP) system, an exascale supercomputer capable of performing 10^18 floating-point operations per second, which will be supplied by France. One of the final products will be calibrated high-resolution images that could be then used for astrophysical purposes.

Problem to be tackled in the internship

Data generated by the telescopes from the SKA project need to be treated by the DDF pipeline (https://github.com/mhardcastle/ddf-pipeline) before being analyzed by scientists. This pipeline needs to

be able to scale and ingest a large amount of data quickly. An important part of its operation being I/O operations, making sure they happen efficiently is paramount to achieve this goal.

Internship goals

The goal of this internship is to understand how the DDF pipeline performs I/O operations, identify what are the issues that prevent it from reaching the required scale, and to propose solutions to these issues.

The recruited person will start by running the pipeline in a cluster, under realistic workloads representative of actual observing scenarios with large radio-interferometers, with some profiling and/or tracing tool to characterize its behavior. Previous work has identified that the most commonly used I/O profiling tool, Darshan, is NOT able to work with the pipeline, which means a new solution has to be found. If no adequate tool is found, the intern will need to modify the source code (Python) to explicitly collect information such as the time spent on different parts of it.

The next step will be to create an I/O kernel of the pipeline, which is a code that mimics an application's I/O behavior, but without doing the other things that the application does such as compute, communicate, etc [2]. The goal of such an I/O kernel is twofold. First, it is an easier way of prototyping and testing ideas for improving the I/O behavior of the application. Second, it can be used by people working to improve I/O infrastructures to test their ideas.

Practical information

This internship will happen in the TADaaM team of the Inria Center at the University of Bordeaux, advised by Francieli Boito and Luan Teylo. However, the project is a collaboration with the Laboratoire d'Astrophysique de Bordeaux (LAB), so the person will spend part of the time working there to have the application point of view.

Desired profile

We are looking for someone that:

- is familiar with Unix command line, ssh, compiling and running codes;
- has some knowledge of HPC and the utilization of clusters;
- can communicate in English (it does not have to be at a great level…);
- is motivated by the subject and to work in a research team.

How to apply

Send your CV to francieli.zanon-boito@u-bordeaux.fr and luan.teylo@inria.fr

References

[1] Francieli Boito, Eduardo Inacio, Jean Bez, Philippe Navaux, Mario Dantas, Yves Denneulin, A Checkpoint of Research on Parallel I/O for High Performance Computing. ACM Computing Surveys, 2018. https://hal.univ-grenoble-alpes.fr/hal-01591755v1/

[2] Francieli Boito, Antonio Tadeu A. Gomes, Louis Peyrondet, Luan Teylo. I/O performance of multiscale finite element simulations on HPC environments. WAMCA 2022 - 13th Workshop on Applications for Multi-Core Architectures, Nov 2022. https://inria.hal.science/hal-03808833/