



Offer #2025-09049

Post-Doctoral Research Visit F/M Leveraging AI for better I/O resources management

Contract type : Fixed-term contract

Renewable contract : Yes

Level of qualifications required : PhD or equivalent

Fonction : Post-Doctoral Research Visit

About the research centre or Inria department

The Inria center at the University of Bordeaux is one of the nine Inria centers in France and has about twenty research teams.. The Inria centre is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative SMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute...

Context

With the increasing scale of applications and emerging technologies such as digital twins, artificial intelligence (AI), climate forecasting, materials science, and engineering problems, high-performance computing (HPC) provides the essential computational power necessary to analyze and interpret large datasets, accelerating discoveries and innovations. In 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 and BeeGFS. 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 input/output (I/O).

The research field of parallel I/O 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.

The performance observed by applications when accessing the I/O infrastructure is heavily impacted by the way this access is done: how many files, if they are shared by processes, how many processes and nodes are involved, how much data is moved, in requests of what size, what is the position of these requests in the files, etc. This set of characteristics is commonly called the application's *access pattern*. This happens because these parameters impact network performance, the efficacy of caching policies, storage devices' behaviors, etc.

Despite this being a well-known and documented phenomenon in the parallel I/O field, over the years studies into the behavior of scientific applications have consistently pointed to the widespread use of low-performance access patterns. There are many reasons for that, but years of such observations make it evident that it is not realistic to force users to adapt to the particularities of HPC I/O: instead, we believe the I/O infrastructure should be able to adapt to all applications' behaviors transparently.

Assignment

Because current systems cannot adapt to applications' characteristics, they are configured to some common case which is suboptimal. They use conservative algorithms (e.g. first-come-first-served scheduling) and default values for the number of storage targets and I/O nodes. Furthermore, it has been observed that, when accessing the parallel file system concurrently, applications will experience interference that is not necessarily uniform, but depends on their access patterns. In other words, some applications will harm (or be harmed by) others more because of their characteristics, and hence should not be allowed to share resources. Because current systems are like that, many applications observe poor I/O performance, high performance variability (because they may be heavily impacted by other jobs), and the expensive HPC resources are not efficiently used, as applications occupy the computing resources while waiting for I/O.

In this context, the goal of the project as a whole is to promote a smarter I/O infrastructure that can adapt itself and make good resource allocation decisions by taking applications' characteristics into account. In this context, the work of the post-doctoral researcher will consist of developing solutions to tackle various challenges, discussed below.

Main activities

1. Estimating the impact of application characteristics on application

performance: The parameters influencing I/O are so numerous that it is extremely difficult to study them in a holistic manner to predict optimal configurations. Our view is that instead of characterizing individual applications, it is best to generate classes of applications that have the same behavior in what concerns a certain system parameter – for example, applications for which the bandwidth varies

similarly as we increase the number of used OSTs. Since the number of classes is expected to be much smaller than the number of applications, this strategy greatly facilitates the task of estimating the impact of configurations on applications.

The post-doctoral researcher will work to identify classes of applications according to different types of I/O resource usage behavior (local devices, burst buffer nodes, OSTs, I/O nodes, etc) and propose techniques - using AI - to predict the class of an application at execution time. That must be done while minimizing the amount of required information and the cost of obtaining it.

2. Classification of temporal I/O behavior: In recent work [1], we studied over 440,000 traces of I/O activity from four HPC systems. We proposed a classification of application temporal I/O behaviors through clustering of time series. Our goal was to identify common patterns exhibited by applications in practice. Jobs, represented as time series, were separated into groups according to duration and peak I/O performance. Then, k-means clustering was applied inside each group, using Dynamic Time Warping (DTW) to compute the similarity between jobs. By manually investigating and labeling each cluster across all the groups, we identified a set of commonly occurring categories that cover most of the jobs. The fact that we found a relatively small number of temporal I/O behavior patterns indicate it is possible to target specific behaviors for I/O scheduling. Moreover, the fact that there are dominating clusters in most groups says temporal I/O behavior is somewhat predictable.

In this context, the work of the post-doctoral researcher will be to continue the work of [1] by including more data sets in the classification, to confirm our preliminary findings and have a more complete view of the classes of temporal I/O behavior. Then, the goal will be to propose an approach of classification that is automatic, without human intervention. This approach should then be adapted for use at execution time.

3. Profiling and simulation of I/O infrastructures: Because the parameters that impact I/O performance are so numerous, and because I/O experiments are inherently slow, simulation could be a powerful tool for I/O improvement techniques at execution time. However, existing simulators are not very accurate at job-level I/O performance due to not covering these numerous parameters. At the same time, improving these simulators is too difficult due to having to model the impact of all these parameters.

The work related to this challenge consists in improving an existing tool for profiling I/O infrastructures, called IOPS [2], to efficiently explore the large parameter space that impacts I/O performance. The result of IOPS must be a model that, given to a simulator, is able to more accurately represent the system. A first idea on how to do that is with Bayesian optimization and surrogate models.

[1] Francieli Boito, Luan Teylo, Mihail Popov, Théo Jolivel, François Tessier, et al.. A Deep Look Into the Temporal I/O Behavior of HPC Applications. 39th IEEE International Parallel & Distributed Processing Symposium (IPDPS), Jun 2025, Milan, Italy. <https://inria.hal.science/hal-04887809>

[2] <https://gitlab.inria.fr/lgouveia/iops>

Skills

The profile we are looking for for this job is someone that communicates well in English and has initiative and autonomy to conduct this research project while tackling the diverse related technical challenges. The person should also have experience in AI techniques, HPC, and in performance evaluation and statistical methods.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours) + possibility of exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

Remuneration

The gross monthly salary will be 2788€ (before sociale security contributions and withholding tax)

General Information

- **Theme/Domain** : Distributed and High Performance Computing System & Networks (BAP E)
- **Town/city** : Talence
- **Inria Center** : [Centre Inria de l'université de Bordeaux](#)
- **Starting date** : 2025-11-01
- **Duration of contract** : 12 months
- **Deadline to apply** : 2025-07-31

Contacts

- **Inria Team** : [TADAAM](#)
- **Recruiter** :
Zanon-boito Francieli / francieli.zanon-boito@inria.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers

working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

The keys to success

The profile we are looking for for this job is someone that communicates well in English and has initiative and autonomy to conduct this research project while tackling the diverse related technical challenges. The person should also have experience in AI techniques, HPC, and in performance evaluation and statistical methods.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

If you are interested, thanks to candidate by jobs.inria with the following documents :

- cv
- cover letter

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.