Offer #2025-09107

PhD Position F/M Application-aware I/O scheduling in HPC systems

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

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

High-performance computing (HPC) plays a critical role in enabling scientific advancements across domains such as artificial intelligence, climate modeling, and engineering by providing the necessary computational power to process massive datasets. In HPC systems, applications running on compute nodes access data stored in parallel file systems (PFS) like Lustre or BeeGFS. These file systems distribute data across multiple storage devices to allow simultaneous access, which is essential to achieve high throughput. However, there remains a persistent gap between the speeds of processing and input/output (I/O), causing many HPC applications to spend considerable time on I/O operations instead of computation.

The efficiency of I/O heavily depends on how applications interact with the file system—specifically, their access patterns. Factors such as the number of files accessed, request sizes, and whether data is shared across processes can drastically affect performance. Moreover, when multiple applications access the PFS concurrently, their access patterns can interfere with one another, degrading overall performance. This interference is not uniform—some applications may disrupt others more depending on their I/O behavior. Such contention can delay job completion, waste compute resources, and increase performance variability. In some cases, poorly optimized applications can consume a disproportionate share of bandwidth, slowing down others and reducing the overall efficiency of the HPC system. Because the effects of this interference are not yet fully predictable, there is a strong need for smarter, application-aware I/O scheduling strategies in HPC environments.

Assignment

The goal of this thesis is to apply I/O scheduling to mitigate contention and improve I/O performance in HPC systems. This thesis will be done in collaboration between two Inria teams (TADaaM, in Bordeaux, and KerData, in Rennes) and in the context of the NumPEx project, which aims at building the software stack of the French exascale systems.

Main activities

Achieving this goal involves tackling different issues, discussed below.

Identification of common access patterns from large data sets: Publicly available datasets of Darshan traces from HPC systems can be used for this purpose. These traces cover a large portion of the jobs executed on a system over a given period and provide aggregated counters for each file accessed by a job. The goal of this activity is to study these datasets to identify common patterns of application behavior. This work will build on existing tools such as MOSAIC. Depending on the number of patterns identified, clustering techniques may be used to group them into representative categories.

Study of the interference between the common access patterns: The next step is to evaluate how these access patterns interfere with each other. Experiments will be conducted on different systems to determine the impact of specific patterns on one another and on overall system performance. The expected outcome is to identify which patterns degrade the performance of others or of the system as a whole. Given the potentially large number of cases to analyze and the time-consuming nature of I/O experiments, a solid and efficient methodology will be required to explore this parameter space effectively.

Access pattern-aware I/O scheduling: Based on the insights from the interference study, an I/O scheduling strategy will be proposed that accounts for application characteristics and their impact on other applications. One idea is to adapt an existing scheduling approach (IO-Sets) originally designed for periodic applications with varying execution frequencies. The key concept is to avoid allowing incompatible applications to access the file system simultaneously. This idea seems promising, but further study is required to determine how applications should be grouped and how the priority of each group should be computed. Additionally, the proposed strategy should take into account temporal I/O behaviors when organizing applications into scheduling sets.

Skills

Other aspects that may be taken into consideration for the selection are:

- having obtained a Master’s in a subject related to high-performance computing;
- some experience in research (internships);
- communication abilities: can communicate well in meetings, most notably in English;

- technical abilities: low-level programming (C, C++, etc), script programming (Bash, Python, etc), use of Linux in command line and through ssh, data analysis (Pandas, etc), writing papers with LaTeX.

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 monthly salary will be 2200€ in 2025 and 2300€ in 2026 (before social security contributions and monthly witholding 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 :** 3 years
- **Deadline to apply :** 2025-08-08

Contacts

- **Inria Team :** TADAAM
- **PhD Supervisor :**
  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 most important characteristic of the ideal candidate is to be motivated to conduct this project and have the autonomy and initiative to learn whatever is required to do it.

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, please could you apply on website 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.