



Offer #2023-06907

PhD Position F/M Modeling and Simulation of Exascale Storage Systems

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Other valued qualifications : Master's degree

Fonction : PhD Position

About the research centre or Inria department

The Inria Rennes - Bretagne Atlantique Centre is one of Inria's eight centres and has more than thirty research teams. The Inria Center 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 PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc.

Context

Context

This thesis is placed in the context of NumPEX (<https://numpex.fr/>), a key national project whose goal is to **co-design the software stack for the exascale era** and prepare applications accordingly. This thesis will be **co-supervised by Inria and CEA**, respectively the Inria center at the University of Rennes and the CEA center at Bruyères-Le-Châtel, near Paris. Beyond the supervision, collaborations within NumPEX with the different partners of the consortium are to be expected.

Logo Inria



PhD Advisors

- François Tessier (Inria)
- Gabriel Antoniu (Inria)
- Philippe Deniel (CEA)
- Thomas Leibovici (CEA)

Location and Mobility

The thesis, co-supervised by Inria and CEA, will be hosted by the KerData team at the **Inria research center of Rennes** and will include regular visits at the CEA Center of Bruyères-le-Châtel. Rennes is the capital city of Brittany, in the western part of France. It is easy to reach thanks to the high-speed train line to Paris. Rennes is a dynamic, lively city and a major center for higher education and research: 25% of its population are students.

This thesis will also include **collaborations with international partners**, especially from the US.

The KerData team in a nutshell for candidates

- KerData is a human-sized team currently comprising 5 permanent researchers, 2 contract researchers, 1 engineer and 5 PhD students. You will work in a caring environment, offering a **good work-life balance**.
- KerData is leading multiple projects in top-level national and international collaborative environments such as within the **Joint-Laboratory on Extreme-Scale Computing** (<https://jlesc.github.io>). Our team has active collaboration with high-profile academic institutions all around the world (including the USA, Spain, Germany or Japan) and with industry.
- Our team strongly favors **experimental research**, validated by implementation and

experimentation of software prototypes with real-world applications on real-world platforms including some of the most powerful supercomputers worldwide.


- The KerData team is committed to personalized advising and coaching, to help PhD candidates train and grow in all directions that are critical in the process of becoming successful researchers.
- Check our **website** for more about the KerData team here <https://team.inria.fr/kerdata/>

Assignment

Introduction

Nowadays, there are many scientific fields such as radio-astronomy or weather forecast for example where the need for computing power and data processing capacity goes beyond what current machines can provide. For these workloads, the resources required are such that **supercomputers capable of reaching the exascale** become necessary. To date, only a few machines such as Frontier at Oak Ridge National Laboratory (USA) have this capability, but in the coming months, new systems will be deployed. However, the efficient use of these systems raises new challenges, especially regarding data management.

Indeed, even though High-Performance Computing (HPC) systems are increasingly powerful, there has been a **relative decline in I/O bandwidth**. Over the past ten years, the ratio of I/O bandwidth to computing power of the top three supercomputers has been divided by 10 (see Figure below) while in some scientific computing centers the volume of data stored has been multiplied by 41 [1]. This tends to accentuate congestion and performance variability on often centralized storage systems [2,3]. To mitigate that, **new levels of storage** have been added to recently deployed supercomputers, increasing their complexity. Harnessing this additional storage capacity is an active research topic but little has been done about how to efficiently provision it [4,5].



Thesis proposal

Dealing with this high degree of storage heterogeneity is a real challenge for scientific workflows and applications. **This PhD thesis proposes to model and simulate heterogeneous storage systems in order to study their behavior, predict their performance and propose innovative algorithmic approaches for better resource utilization.**

Main activities

One of the aims of this thesis is to **make better use of storage resources** for scientific applications and workflows that are destined to run on exascale supercomputers. Initially, **storage systems such as Lustre and DAOS will be studied, modeled and simulated in an existing WRENCH-based [6] simulator, called StorAlloc [5]**, developed in the team. This study will shed light on the criteria influencing the performance of these systems. Secondly, **advanced resource allocation algorithms will be proposed, implemented and evaluated in the simulator** to overcome the limitations of existing methods (e.g. Lustre uses the disks of its storage system in a simple round-robin manner). Multiple criteria can be taken into account in those algorithms such as contention or energy. Tools developed by the CEA, including the Robinhood policy engine [7] and the outcomes from the IO-SEA European Project [8] will also be used to validate these contributions on real systems. For this work, a strong emphasis will be put on **international collaborations**, especially with the University of Manoa (HI, USA), and **national partnership** such as with the French **SKA** team providing a relevant use-case for this work. The candidate will also have the opportunity to be hosted for 3-6 month internships abroad to strengthen the international visibility of his/her work and benefit from the expertise of other researchers in the field.

References

[1] GK. Lockwood, D. Hazen, Q. Kozioł, RS. Canon, K. Antypas, and J. Balewski. "Storage 2020: A Vision for the Future of HPC Storage". In: Report: LBNL-2001072. Lawrence Berkeley National Laboratory, 2017.

[2] O. Yildiz, M. Dorier, S. Ibrahim, R. Ross, and G. Antoniu. "On the Root Causes of Cross-Application I/O Interference in HPC Storage Systems". In: 2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS). 2016, pp. 750–759

[3] F. Tessier, V. Vishwanath. "Reproducibility and Variability of I/O Performance on BG/Q: Lessons Learned from a Data Aggregation Algorithm". United States: N. p., 2017. Web. doi:10.2172/1414287

[4] F. Tessier, M. Martinasso, M. Chesi, M. Klein, M. Gila. "Dynamic Provisioning of Storage Resources: A Case Study with Burst Buffers". In: IPDPSW 2020 - IEEE International Parallel and Distributed Processing Symposium Workshops, May 2020, New Orleans, United States.

[5] J. Monniot, F. Tessier, M. Robert, G. Antoniu. "StorAlloc: A Simulator for Job Scheduling on Heterogeneous Storage Resources". In: HeteroPar 2022, Aug 2022, Glasgow, United Kingdom.

[6] H. Casanova, R. Ferreira da Silva, R. Tanaka, S. Pandey, G. Jethwani, W. Koch, S. Albrecht, J. Oeth, and F. Suter. "Developing Accurate and Scalable Simulators of Production Workflow Management Systems with WRENCH". In: Future Generation Computer Systems, vol. 112, p. 162-175, 2020.

[7] <https://github.com/cea-hpc/robinhood>

[8] <https://iosea-project.eu/>

Skills

- An excellent Master degree in computer science or equivalent
- Completion of a teaching unit in high-performance computing or distributed computing is an advantage
- Programming skills in C/C++ and Python
- Good communication skills in oral and written English.
- Open-mindedness, strong integration skills and team spirit

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Possibility of teleworking (90 days per year) and flexible organization of working hours
- Partial payment of insurance costs

Remuneration

monthly gross salary amounting to 2051 euros for the first and second years and 2158 euros for the third year

General Information

- **Theme/Domain** : Distributed and High Performance Computing Scientific computing (BAP E)
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2024-09-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2024-12-31

Contacts

- **Inria Team** : [KERDATA](#)
- **PhD Supervisor** :
Tessier François / francois.tessier@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 candidate will have to show motivation, autonomy and an ability to initiate links between the research activities carried out at INRIA and at the CEA center.

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

Please submit online : your resume, cover letter and letters of recommendation eventually

For more information, please contact francois.tessier@inria.fr

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.