



Offer #2024-08227

PhD Position F/M Design and Implementation of a Scalable Naming Service in Shared Memory

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 Saclay-Île-de-France Research Centre was established in 2008. It has developed as part of the Saclay site in partnership with **Paris-Saclay University** and with the **Institut Polytechnique de Paris**.

The centre has [40 project teams](#), 32 of which operate jointly with Paris-Saclay University and the Institut Polytechnique de Paris; Its activities occupy over 600 people, scientists and research and innovation support staff, including 44 different nationalities.

Assignment

Context

The CXL [1] standard will profoundly impact resource management in data centers. CXL defines a cache coherency domain that not only includes system memory and CPUs, but also PCIe devices. It opens the way to fully disaggregated data centers, as the PCIe buses of a cluster of machines can be connected through a *CXL fabric* [2, 3] which allows the loads and stores emitted by a processor to be transparently routed to the memory of the receiver through a cluster-scale cache-coherency protocol. At the software level, *far* memory located in another machine can be accessed as transparently as local memory: a simple statement such as $a = 42$ can be routed seamlessly to any memory of any machine connected to the CXL fabric.

In this context, the traditional design in which independent storage nodes are accessed by compute nodes over a network becomes inadequate. This approach has been dominant in the past due to its ability to independently scale computation and storage. However, it incurs significant costs from data exchange and transformation which can be avoided by taking advantage of the efficient cluster-scale cache-coherency protocol provided by CXL.

PhD Topic

In this project, we aim to reimagine the architecture of cloud applications in the CXL era. Our approach decouples memory from processes, enabling global memory sharing across processes, similar to how threads share memory in multi-threaded applications. However, unlike the multi-threaded model, memory objects in this design can persist beyond the lifespan of individual processes, acting as long-term storage for ephemeral processes that are launched on demand to serve clients or handle large-scale data analytics. Since any process can directly access the global memory, the architecture avoids the high cost of transforming data when it is exchanged between the processes.

Central to this architecture is a naming service, which, at a high level, is reminiscent of a classical file system. This naming service must make it possible for each process to retrieve objects produced by other processes. It must reside in shared memory and scale to thousands of processes. As a PhD student, you will study how such a naming service can be designed, leveraging low-level hardware features such as virtualization extensions to alleviate bottlenecks.

References

[1] Debendra Das Sharma. Compute Express Link (CXL): enabling heterogeneous data-centric computing with heterogeneous memory hierarchy. *IEEE Micro*, 2022.

[2] Donghyun Gouk, Sangwon Lee, Miryeong Kwon, and Myoungsoo Jung. Direct access, high-performance memory disaggregation with DirectCXL. In *Proceedings of the USENIX Annual Technical Conference, USENIX ATC'22*, 2022.

[3] Huaicheng Li, Daniel S. Berger, Lisa Hsu, Daniel Ernst, Pantea Zardoshti, Stanko Novakovic, Monish Shah, Samir Rajadnya, Scott Lee, Ishwar Agarwal, Mark D. Hill, Marcus Fontoura, and Ricardo Bianchini. Pond: CXL-based memory pooling systems for cloud platforms. In *Proceedings of the conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS'23*, 2023.

Main activities

Main activities:

- Design and implement a system
- Evaluate the performance of the proposed system
- Write reports and papers

Additional activities:

- Present the work as well as related work
- Attend seminars, workshops, and/or conferences

Skills

The candidate must have a good background in system programming, concurrent programming, distributed systems, and C.

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

Remuneration

- 1st and 2nd year : 2100€ gross/month
- 3rd year : 2190€ gross/month

General Information

- **Theme/Domain** : Distributed Systems and middleware System & Networks (BAP E)
- **Town/city** : Palaiseau
- **Inria Center** : [Centre Inria de Saclay](#)
- **Starting date** : 2024-12-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2024-11-30

Contacts

- **Inria Team** : [BENAGIL](#)
- **PhD Supervisor** :
Lozi Jean-pierre / jean-pierre.lozi@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.

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

Inria is committed to the employment of people with disabilities. All of the Institute's positions are open to disabled workers. Inria implements personalized follow-up during recruitment, when the person starts to work and throughout their career.

In addition, every year Inria opens up positions to people who are recognized as disabled workers. These recruitments are offered in a variety of professional fields ranging from research to management and the operation of the institute.

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.