



Offer #2025-09329

PhD Position F/M Cost and Performance- Efficient Caching for Massively Distributed 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 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

Financial and working environment.

This PhD will be in the context of IPCEI-CIS (Important Project of Common European Interest – Next Generation Cloud Infrastructure and Services) DXP (Data Exchange Platform) project involving Amadeus and three Inria research teams (COAST, CEDAR and MAGELLAN). This project aims to design and develop an open-source management solution for a federated and distributed data exchange platform (DXP), operating in an open, scalable, and massively distributed environment (cloud-edge continuum).

The PhD student will be recruited and hosted at the Inria Center at Rennes University; and the work will be carried out within the MAGELLAN team in collaboration with other partners.

The PhD student will be supervised by:

- Shadi Ibrahim, MAGELLAN team in Rennes
- Cedric Tedeschi, MAGELLAN team in Rennes

Assignment

Context

The ever-growing number of services and Internet of Things (IoT) devices has resulted in data being distributed across different locations (regions and countries). Additionally, data exhibits different usage patterns, including cold data (written once and never read), stream data (produced once and consumed by many), and hot data (written once and consumed by many). Furthermore, these data types have different performance and dependability requirements (e.g., low latency for data streams).

Data caching is a widely used technique that improves application performance by storing data on high-speed devices close to end users. Most research on data caching has focused on the benefits of different data placement strategies (i.e., which data to place in the cache), data movement, cache partitioning, cache eviction [1, 2, 3, 4, 5, 6, 7, 8], and on realizing cost-efficient data redundancy techniques in caching systems [9]. However, few efforts have studied data management when caches are distributed across different platforms (Edge-to-Cloud), utilize heterogeneous storage devices (in terms of performance and cost), and serve multiple, diverse applications, including traditional data services, serverless workflows and data streaming.

References:

- [1] Asit Dan and Don Towsley. 1990. An Approximate Analysis of the LRU and FIFO Buffer Replacement Schemes. *SIGMETRICS Perform. Eval. Rev.* 18, 1 (apr 1990), 143–152. <https://doi.org/10.1145/98460.98525>
- [2] Marek Chrobak and John Noga. 1999. LRU is better than FIFO. *Algorithmica* 23 (02 1999), 180–185. <https://doi.org/10.1007/PL00009255>
- [3] Blankstein, Aaron, Siddhartha Sen, and Michael J. Freedman. “Hyperbolic caching: Flexible caching for web applications.” 2017 USENIX Annual Technical Conference (USENIX ATC 17). 2017.
- [4] Cristian Ungureanu, Biplob Debnath, Stephen Rago, and Akshat Aranya. 2013. TBF: A memory-efficient replacement policy for flash- based caches. In 2013 IEEE 29th International Conference on Data Engineering (ICDE). 1117–1128. <https://doi.org/10.1109/ICDE.2013.6544902>
- [5] Orcun Yildiz, Amelie Chi Zhou, Shadi Ibrahim. 2018. Improving the Effectiveness of Burst Buffers for Big Data Processing in HPC Systems with Eley. *Future Generation Computer Systems*, Volume 86, 2018, Pages 308-318, ISSN 0167-739X, <https://doi.org/10.1016/j.future.2018.03.029>.
- [6] G. Aupy, O. Beaumont and L. Eyraud-Dubois, "Sizing and Partitioning Strategies for Burst-Buffers to Reduce IO Contention," *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Rio de Janeiro, Brazil, 2019,
- [7] ZHANG, Yazhuo, YANG, Juncheng, YUE, Yao, *et al.* {SIEVE} is simpler than {LRU}: an efficient {Turn-Key} eviction algorithm for web caches. In : *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)* . 2024. p. 1229-1246.

[8] Juncheng Yang, Ziming Mao, Yao Yue, and K. V. Rashmi. GL-Cache: Group-level learning for efficient and high-performance caching. FAST'23, pages 115–134, 2023.

[9] RASHMI, K. V., CHOWDHURY, Mosharaf, KOSAIAN, Jack, *et al.* {EC-Cache} : {Load-Balanced}, {Low-Latency} cluster caching with online erasure coding. In : *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 2016. p. 401-417.

Main activities

The goal is to design cost- and performance-efficient distributed smart caching middleware that facilitates data exchange within and across data providers (and producers) and consumers (users) while considering the temperature of the data, its frequency, and the heterogeneity and dynamics of the infrastructure. Specifically, we aim to address these research questions:

- How can we seamlessly aggregate the caches of many distributed, heterogeneous machines?
- Where to place data across different sites/organizations (across the IoT to cloud continuum).
- Which data should be cached and for how long, how to resize the caches between different applications/users, when to empty caches, etc
- How to exploit data caches for data streams and how to efficiently share caches with hot data.
- In addition, it is important to study the right number of replicas to meet the users demands.

Research Methodology: This research addresses the challenges of data management in distributed, heterogeneous caches by designing and implementing novel middleware, models, algorithms, and a framework to answer the above questions. All solutions will be validated through simulations or on a real distributed infrastructure, such as Grid'5000 and Amazon Web Services.

Skills

- An excellent Master degree in computer science or equivalent
- Strong knowledge of distributed systems
- Knowledge of storage and distributed file systems
- Ability to conduct experimental systems research
- Strong programming skills (C/C++, Python)
- Working experience in the areas of Big Data management, Cloud Computing, and Data Analytics are advantageous
- Very good communication skills in oral and written English

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 (after 6 months of employment) 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

monthly gross salary 2200 euros

General Information

- **Theme/Domain** : Distributed Systems and middleware System & Networks (BAP E)
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2025-11-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2026-04-30

Contacts

- **Inria Team** : [MAGELLAN](#)
- **PhD Supervisor** :
Ibrahim Shadi / Shadi.Ibrahim@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

-

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

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.