



Offre n°2025-09182

Post-Doctoral Research Visit F/M Scheduling Data-Intensive Applications in P2P Environments

Type de contrat : Fixed-term contract

Niveau de diplôme exigé : PhD or equivalent

Fonction : Post-Doctoral Research Visit

A propos du centre ou de la direction fonctionnelle

The Inria Centre at Rennes University is one of Inria's nine centres and has more than thirty 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, etc.

Contexte et atouts du poste

Financial and working environment.

This postdoc position **will be funded by the Cupseli Inria Challenge (Défi Inria)**. The Cupseli Inria Challenge brings together 11 Inria teams distributed over 6 Inria centers and the Hive startup company based in Cannes. The position 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 position is for one year.

About Hive:

Hive intends to play the role of a next generation cloud provider in the context of Web 3.0. Hive aims to exploit the unused capacity of computers to offer the general public a greener and more sovereign alternative to the existing clouds where the true power lies in the hands of the users. It relies both on distributed peer-to-peer networks, on the encryption of end-to-end data and on blockchain technology.

Mission confiée

Context:

Large-scale P2P environments are characterized by a high number of node failures and churns [1]. This can lead to unwanted delays in the completion time of running applications and makes both scalability and reliability critical when running data-intensive applications (e.g., MapReduce applications [2]) in a peer-to-peer compute environment. We are interested in investigating how to optimize the execution of data-intensive applications in the presence of failures and churns by leveraging P2P storage services (e.g., hive-Disk platform [3]), using checkpoints and making job scheduling failure-aware.

Objectives:

General purpose fault tolerant strategies lead to excessive execution of recovery tasks (re-execution of tasks on failed machines). Therefore, we will investigate how to adapt fault-tolerance techniques to P2P systems by making job scheduling failure-aware (leveraging our previous experience and work with Hadoop clusters [4, 5]) and by enabling checkpoint/restart so that we can roll back execution from the last checkpoint instead of restarting the execution after a failure [6]. We will present a performance model for checkpoint/restart in P2P systems and introduce a scheduling framework that decides when and where to trigger checkpoints and where to restart, and when and where to execute recovery tasks, taking into account failure distribution, data location, and resource heterogeneity. We will also explore how to use P2P storage services (e.g., hive-Disk platform) to store checkpoints and temporary data (e.g., map outputs in MapReduce).

[1] Apostolos Malatras. "State-of-the-art survey on P2P overlay networks in pervasive computing environments". In: *Journal of Network and Computer Applications* 55 (2015), pp. 1–23.

[2] Jeffrey Dean and Sanjay Ghemawat. "MapReduce: simplified data processing on large clusters." *Communications of the ACM* 51.1 (2008): 107-113.

[3] <https://www.hivenet.com/store-with-hivenet-cloud-storage>

[4] Orcun Yildiz, Shadi Ibrahim, Tran Anh Phuong, and Gabriel Antoniu. Chronos: Failure-aware scheduling in shared hadoop clusters. In 2015 IEEE International Conference on Big Data (Big Data), pages 313–318. IEEE, 2015.

[5] Orcun Yildiz, Shadi Ibrahim, and Gabriel Antoniu. Enabling fast failure recovery in shared hadoop clusters: towards failure-aware scheduling. *Future Generation Computer Systems*, 74:208–219, 2017.

[6] Ifeanyi P Ekwutuoha, David Levy, Bran Selic, and Shiping Chen. A survey of fault tolerance mechanisms and checkpoint/restart

implementations

Principales activités

- Read and synthesize literature work.
- Design new scheduling policies for Data-Intensive Applications in P2P Environments.
- Implementation and large-scale validation.
- Participate in project meetings and discussions with other partners.
- Write research papers and disseminate results through presentations at project meetings, conferences, and workshops.

Compétences

- A Ph.D. in computer science
- A solid background in the area of distributed systems
- Ability to conduct experimental systems research
- Experience with building systems and tools
- Working experience in the areas of Big Data management, Cloud Computing, Data Analytics are advantageous
- Very good communication skills in oral and written English

Avantages

- 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

Rémunération

Monthly gross salary amounting to 2788 euros

Informations générales

- **Thème/Domaine** : Distributed Systems and middleware System & Networks (BAP E)
- **Ville** : Rennes
- **Centre Inria** : [Centre Inria de l'Université de Rennes](#)

- **Date de prise de fonction souhaitée** : 2026-01-01
- **Durée de contrat** : 12 months
- **Date limite pour postuler** : 2025-09-18

Contacts

- **Équipe Inria** : [MAGELLAN](#)
- **Recruteur** :
Ibrahim Shadi / Shadi.Ibrahim@inria.fr

A propos d'Inria

Inria est l'institut national de recherche dédié aux sciences et technologies du numérique. Il emploie 2600 personnes. Ses 215 équipes-projets agiles, en général communes avec des partenaires académiques, impliquent plus de 3900 scientifiques pour relever les défis du numérique, souvent à l'interface d'autres disciplines. L'institut fait appel à de nombreux talents dans plus d'une quarantaine de métiers différents. 900 personnels d'appui à la recherche et à l'innovation contribuent à faire émerger et grandir des projets scientifiques ou entrepreneuriaux qui impactent le monde. Inria travaille avec de nombreuses entreprises et a accompagné la création de plus de 200 start-up. L'institut s'efforce ainsi de répondre aux enjeux de la transformation numérique de la science, de la société et de l'économie.

Attention: Les candidatures doivent être déposées en ligne sur le site Inria. Le traitement des candidatures adressées par d'autres canaux n'est pas garanti.

Consignes pour postuler

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

Sécurité défense :

Ce poste est susceptible d'être affecté dans une zone à régime restrictif (ZRR), telle que définie dans le décret n°2011-1425 relatif à la protection du potentiel scientifique et technique de la nation (PPST). L'autorisation d'accès à une zone est délivrée par le chef d'établissement, après avis ministériel favorable, tel que défini dans l'arrêté du 03 juillet 2012, relatif à la PPST. Un avis ministériel défavorable pour un poste affecté dans une ZRR aurait pour conséquence l'annulation du recrutement.

Politique de recrutement :

Dans le cadre de sa politique diversité, tous les postes Inria sont accessibles aux personnes en situation de handicap.