



Offre n°2025-09040

PhD Position F/M Resilient network communications, tolerant to failures and volatility, and suited for artificial intelligence applications.

Type de contrat : Fixed-term contract

Niveau de diplôme exigé : Graduate degree or equivalent

Fonction : PhD Position

Niveau d'expérience souhaité : Recently graduated

A propos du centre ou de la direction fonctionnelle

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...

Contexte et atouts du poste

The position is part of a collaboration between the Hivenet company and the TOPAL team. The successful candidate will be part of the TOPAL team based at Inria Bordeaux.

About TOPAL: In TOPAL, we're tackling the evolving challenges at the intersection of high-performance computing (HPC), numerical simulation, and machine learning. As computing platforms grow in scale and complexity—featuring millions of cores and diverse hardware like GPUs—we're leveraging our long-standing expertise in dynamic runtime systems to make the most of these resources. Our work helps adapt to unpredictable workloads and optimize task scheduling without relying on rigid, pre-planned execution. We're now extending this know-how to emerging applications like deep neural network training, which places unique demands on both computation and memory. At the same time, we're addressing the urgent need to reduce energy consumption and carbon footprint in HPC. That means rethinking algorithms, data movement, and hardware use to build more sustainable systems. A key focus for us is efficient data management, since moving and storing data is increasingly more costly than computing itself. By combining our background in linear algebra, resource

scheduling, and algorithmic optimization with these new demands, we're shaping the future of scientific computing on next-generation platforms.

About the Hivenet company: Hivenet is shaping the future of cloud computing by leveraging unused computing capacity to provide a decentralized, environmentally friendly, and user-empowered alternative to traditional cloud services.

Mission confiée

Context

Hivenet is a company that offers individuals and businesses the ability to make their unused computing resources available. Hivenet offers a data storage service called **HiveDisk**, which utilizes the storage space contributed by HiveDisk users. This enables HiveDisk users to benefit from geo-distributed and replicated storage. Similarly, Hivenet aims to share, via **HiveCompute**, unused computing resources (primarily GPUs) to perform tasks related mainly to the training and inference of artificial intelligence applications. Through a web interface, users can request the allocation of a certain number of GPUs distributed across different machines and then access them to run their computations.

Initially, the allocated GPUs will be located on machines within the same local network (for example, within a company's site network or a [PoliCloud](#) container), but the long-term goal is to leverage GPUs located across different enterprise or community networks in various locations (e.g., all the sites of a company across a country) [1].

This project presents numerous challenges, mainly because the targeted environment differs from traditional HPC environments. From a hardware perspective, the machines are less powerful, heterogeneous, and interconnected via standard networks that are less efficient and reliable than HPC networks. It is also important to consider that computing resources are not available at all times (for example, machines may be less available during the day when employees are using them) and are more likely to become unavailable unexpectedly. Furthermore, using machines spread across different geographical locations results in a network with heterogeneous performance: latency between two sites is much higher than within a single site.

Objectives

The objective of this PhD thesis is to explore issues related to network communications in such a context. This will require an analysis of existing communication libraries (such as [PCCL](#) [3], MPI [2], or [libp2p](#)) in order to determine which one(s) can be adapted to the targeted environment. Once a suitable communication model is established, the work will focus—given a set of machines and their topology—on adapting communication patterns in learning applications to minimize communication overhead. This may involve using routing algorithms and better distribution of computation and data, tailored to the characteristics of the network interconnecting the machines.

The system will also need to detect the addition or loss of machines and respond accordingly—for example, by ignoring the contributions of lost machines in a

data-parallel setup, or by redistributing the data and computation.

In a second phase, the project will consider how to manage network usage when HiveDisk and HiveCompute are active simultaneously on the same networks and machines. The goal will be to maintain acceptable performance levels for both services by dynamically adjusting quality-of-service parameters based on network conditions and user requirements.

References

- [1] N. T. Karonis, B. de Supinski, I. Foster, W. Gropp and E. Lusk, "A Multilevel Approach to Topology-Aware Collective Operations in Computational Grids." arXiv preprint cs/0206038, 2002
- [2] L. Shalev, H. Ayoub, N. Bshara, and E. Sabbag, "A Cloud-Optimized Transport Protocol for Elastic and Scalable HPC", *IEEE micro*, 40(6), 67-73.
- [3] M. Keiblinger, M. Sieg, J. Min Ong, S. Jaghouar and J. Hagemann, "Prime Collective Communications Library -- Technical Report", *arXiv preprint arXiv:2505.14065*, 2025

Principales activités

The PhD student will conduct original researches on the topic described above, and will collaborate with colleagues in the TOPAL team and Hivenet partners.

Activities includes, but are not limited to: bibliographical synthesis, research, software implementation, presentation of results at conferences, attending research schools, etc

Compétences

Technical skills and level required :

- Solid understanding of network communication (socket, TCP/IP);
- Proficiency in systems programming (C/C++) and high level language (Python).
- Experience with AI frameworks (PyTorch, Tensorflow, ...) and software performance measurement is a plus
- Interest in distributed systems and HPC are also a plus

Languages :

Good communication skills in English (French is a plus)

Relational skills :

Ability to work collaboratively in an academic–industry setting

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 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

The monthly salary will be 2200€, and 2300€ in 2026 (before social security contributions and withholding tax)

Informations générales

- **Thème/Domaine** : Distributed and High Performance Computing
Scientific computing (BAP E)
- **Ville** : Talence
- **Centre Inria** : [Centre Inria de l'université de Bordeaux](#)
- **Date de prise de fonction souhaitée** : 2025-09-01
- **Durée de contrat** : 3 years
- **Date limite pour postuler** : 2025-07-31

Contacts

- **Équipe Inria** : [TOPAL](#)
- **Directeur de thèse** :
Herault Thomas / thomas.herault@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'engage ainsi de répondre aux enjeux de la transformation numérique de la science, de la société et de l'économie.

L'essentiel pour réussir

Feeling comfortable in a dynamic scientific environment, enjoying learning and tackling real-world problems, and having an experimental mindset are essential qualities for succeeding in this role.

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

If you are interested, thanks to candidate on [jobs.inria](https://jobs.inria.fr) with the following documents :

- cv
- cover letter

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.