



Offre n°2020-02408

Doctorant F/H Supporting In-Situ Machine Learning for Fast Data Analytics

Type de contrat : CDD

Niveau de diplôme exigé : Bac + 5 ou équivalent

Autre diplôme apprécié : Master of Science or Engineering

Fonction : Doctorant

A propos du centre ou de la direction fonctionnelle

Le centre Inria Rennes - Bretagne Atlantique est un des huit centres d'Inria et compte plus d'une trentaine d'équipes de recherche. Le centre Inria est un acteur majeur et reconnu dans le domaine des sciences numériques. Il est au cœur d'un riche écosystème de R&D et d'innovation : PME fortement innovantes, grands groupes industriels, pôles de compétitivité, acteurs de la recherche et de l'enseignement supérieur, laboratoires d'excellence, institut de recherche technologique.

Contexte et atouts du poste

- **Advisors:** [Alexandru Costan](#), [Gabriel Antoniu](#) (KerData team)
- **Main contacts:** [gabriel.antoniu \(at\) inria.fr](mailto:gabriel.antoniu@inria.fr), [gabriel.antoniu \(at\) inria.fr](mailto:gabriel.antoniu@inria.fr)
- **Collaboration context:** [JLESC International Laboratory](#) on Extreme-Scale Computing
- **Expected start date:** October 1st, 2019

Location and Mobility

The thesis will be mainly hosted by the KerData team at Inria Rennes Bretagne Atlantique. It will include collaborations with **Argonne National Lab, USA** (which provides one of the target applications, where the student is expected to be hosted for a 3-month internship). 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.

The KerData team in a nutshell for candidates

- As a PhD student mainly hosted in the KerData team, you will join a dynamic and enthusiastic group, committed to top-level research in the areas of High-Performance Computing and Big Data Analytics. Check the team's web site: <https://team.inria.fr/kerdata/>.
- The team is leading multiple projects in top-level national and international collaborative environments, e.g., the JLESC international Laboratory on Extreme-Scale Computing: <https://jlesc.github.io>. It has active collaborations with top-level academic institutions all around the world (including the USA, Mexico, Spain, Germany, Japan, Romania, etc.). The team has close connections with the industry (e.g., Microsoft, Huawei, Total).
- The KerData team's publication policy targets the best-level international journals and conferences of its scientific area. The team also strongly favors experimental research, validated by implementation and experimentation of software prototypes with real-world applications on real-world platforms, e.g., clouds such as Microsoft Azure and some of the most powerful supercomputers in the world.

Why joining the KerData team is an opportunity for you

- The team's top-level collaborations strongly favor successful PhD theses dedicated to solving challenging problems at the edge of knowledge, in close interaction with top-level experts from both academia and industry.
 - To follow the career of our former PhD students, have a look here: <https://team.inria.fr/kerdata/team-members/>.

- **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, top-level researchers.
- You will have the opportunity to present your work in top level venues where you will meet the best experts in the field.
- **What you will learn.** Beyond learning how to perform meaningful and impactful research, you will acquire useful skills for communication both in written form (how to write a good paper, how to design a convincing poster) and in oral form (how to present their work in a clear, well-structured and convincing way).
 - This is how some of our PhD students received awards in recognition to the quality of their research. Have a look here: <https://team.inria.fr/kerdata/awards/>.
- **Additional complementary training** will be available, with the goal of preparing the PhD candidates for their postdoctoral career, should it be envisioned in academia, industry or in an entrepreneurial context, to create a startup company.

Mission confiée

- **Main contacts:** gabriel.antoniou (at) inria.fr, alexandru.costan (at) inria.fr
- **Expected start date:** October 1st, 2019
- **Application deadline:** as early as possible, no later than May 30, 2019

Description

Fast Data Analytics refers to the process of examining and extracting relevant knowledge from sets of data which are so huge, exhibit such a high format variety and are generated at such a high speed that traditional systems for data storage and processing cannot be used any longer in an efficient way to extract knowledge in an acceptable time. Potentially coming from a very large variety of sources (e.g., sensors from the Internet of Things, social networks, business applications), such data arrive with a high rate and need to be reacted to in real-time, hence they are often referred as Fast Data. They are curated, (sometimes partially) stored, processed and fed into analysis engines that build representations through data-driven models that further enable descriptive, predictive and prescriptive analytics to get valuable insights for decision making.

Why in situ and in transit analytics? A traditional approach for analyzing such data generated by simulations consists of transferring them from the supercomputer to persistent storage, from where they can be visualized and analyzed after the end of the computation (offline analysis). However, as the exascale is getting near, with the challenging volumes of data generated at increasing rates, storing the data for later, “offline” analysis, becomes infeasible. Hence, many computational scientists have moved to “in situ” analysis strategies, in which analysis tasks are executed on the same nodes (where the simulation is running). This brings several advantages : it is not necessarily useful any longer to store the whole data before it is analyzed; data transfers could actually even be fully avoided if the whole analytics can be performed “in situ.” In addition, it becomes possible to get early insights from the simulation while it runs, which makes it possible to detect ways to act on the simulation (e.g., change its parameters, increase the simulation resolution, etc.). In the case where the analytics is heavier, data generated by the simulations can be transferred from the simulation resources to dedicated resources for what is called “in transit analytics” (still performed while the simulation is running, but more loosely coupled with it), before they are finally transferred on persistent storage. These techniques (in situ/in transit analytics) usually rely on libraries or middlewares that interfaces the simulation with analysis codes.

Why Artificial Intelligence is a catalyst for Fast Data Analytics. Fast Data Analytics increasingly relies on Machine Learning (ML), a subfield of Artificial Intelligence typically used for data classification and feature extraction. While traditional ML deals with tractable feature extraction, Deep Learning recently attracted a very high interest as a particularly efficient approach when classical machine learning is intractable. DL relies on neural network representations with a high number of layers, able to learn very complex representations and subsequently use them for predictions (inference). It uses

dense linear algebra kernels and allows for lower-precision representation and arithmetic, for which general-purpose GPU (GPGPU) accelerators (increasingly available on HPC systems) are a relevant infrastructure. Consequently, DL-based Fast Data Analytics generates workloads that naturally fit HPC systems, thereby acting as a catalyst for HPC-Fast Data convergence.

Challenge: support machine learning based in situ/in transit analytics. To enable in situ/in transit analysis, various software technologies such as Damaris [1,2], Decaf [3], ADIOS [4] and FlowVR [5] have been developed. In all of the mentioned technologies, a subset of available resources (e.g. cores/nodes) are allocated to data analysis tasks. But, in current solutions, the focus is mainly on improving I/O and the support for data analysis is very limited. To this date, no existing in situ / in transit analysis middleware offers programmatic support for Big Data analytics.

Focus of the thesis: the data processing level. In the high-performance computing area (HPC), the need to get fast and relevant insights from massive amounts of data generated by extreme-scale computations led to the emergence of in situ and in transit processing approaches. They allow data to be visualized and processed in real-time, in an interactive way, as they are produced, as opposed to traditional approach consisting of transferring data off-site after the end of the computation, for offline analysis. In the Fast and Big Data area, the search for real-time, fast analysis was materialized through a different approach: stream-based processing, in support to intelligent, ML-based data analytics.

Thesis goal. This PhD thesis aims to propose a paradigm shift at the in-situ processing level: for the first time, we plan to introduce programmatic support for Big Data analytics on platforms that were traditionally used only for simulations. In particular, we will target dedicated Machine Learning support. To this end, we plan to leverage an existing building block - the Damaris middleware, developed by the KerData team at Inria. The goal is to extend Damaris towards a more elaborate system, providing the possibility to use dedicated cores or dedicated nodes for in-situ processing. In particular, we will add Big Data analytics support by means of dedicated plug-ins for distributed, high-performance, always-available, elastic and accurate data processing. Furthermore, we plan to design a mechanism for incrementally migrating running stream tasks from the in-situ processing backend to the in-transit one without stopping the query execution.

Principales activités

Target use case. This PhD will analyze and address the requirements of a use case on machine learning coherent diffraction data made available by the group of Tom Peterka at Argonne National Lab, with which the KerData team is collaborating. X-ray coherent diffractive imaging (CDI) is a remarkable tool for seeing nanoscale materials. Classifying features such as defects in those materials, however, is difficult because state-of-the-art iterative algorithms for CDI phase retrieval are compute-intensive, and classification is done manually. The goal is to replace current compute- and human- driven methods with data-driven ones. The proposal studies whether atomic-scale defects such as dislocations as well as lattice positions in real space can be learned from diffraction patterns in reciprocal space without being defined analytically and without running expensive computations to reconstruct the real-space image. To understand bottlenecks and inefficiencies, we will setup and conduct experiments for the overall workflow of this use case using the in-situ Machine Learning based data analytics, analyze the logs and use the whole workflow as a baseline to compare and validate the outcome of the other work packages (algorithms, frameworks, resource management techniques).

Enabling techniques. In the process of designing the unified data processing framework, we will leverage in particular techniques for data processing already investigated by the participating teams as proof-of-concept software, validated in real-life environments:

- The **Damaris [1]** framework for scalable, asynchronous I/O and in situ and in transit visualization and processing (developed at Inria, <https://project.inria.fr/damaris/>). Damaris already demonstrated its scalability up to 16,000 cores on some of the top supercomputers of Top500, including Titan, Jaguar and Kraken). Developments are currently in progress in a contractual framework between Total and Inria to use Damaris for in situ visualization for extreme-scales simulations at Total. For the purpose of this work, Damaris will have to be extended to support Big Data analytics

plugins for data processing (e.g., based on the Flink and Spark engines and on their higher-level machine-learning libraries).

- The **KerA [6]** approach for low-latency storage for stream processing (currently under development at Inria, in collaboration with UPM, in the framework of a contractual partnership between Inria and Huawei Munich). By eliminating storage redundancies between data ingestion and storage, preliminary experiments with KerA successfully demonstrated its capability to increase throughput for stream processing. Kera is now subject of interest for exploitation plans by Huawei.

The resulting framework will be integrated in a state-of-the-art data processing ecosystem (Spark or Flink) and allow to apply in situ/in transit advanced tools for Big Data analytics (e.g. ML-based) using stream-based techniques, to combine the result with historical data and thereby derive insights from data in real time. These insights can further be used to steer the simulation.

Partnership and valorisation

Leveraging the ongoing collaboration with TU Berlin (the main contributor to the Apache Flink reference framework), the prototype designed during this PhD will be **integrated in the Apache Flink state-of-the-art data processing ecosystem**. It will allow to perform both Edge and Cloud analytics and thereby derive insights from data in real-time.

References

- [1] M. Dorian, G. Antoniu, F. Cappello, M. Snir, L. Orf. "Damaris: How to Efficiently Leverage Multicore Parallelism to Achieve Scalable, Jitter-free I/O", In Proc. CLUSTER – IEEE International Conference on Cluster Computing, Sep 2012, Beijing, China. URL: <https://hal.inria.fr/hal-00715252>
- [2] M. Dorian, R. Sisneros, T. Peterka, G. Antoniu, D. Semeraro, "Damaris/Viz: a Nonintrusive, Adaptable and User-Friendly In Situ Visualization Framework", Proc. Lдав – IEEE Symposium on Large-Scale Data Analysis and Visualization, Oct 2013, Atlanta, USA. URL: <https://hal.inria.fr/hal-00859603>
- [3] M. Dreher, T. Peterka, "Decaf: Decoupled Dataflows for In Situ High-Performance Workflows", Technical Report, United States, doi:10.2172/1372113, <https://www.osti.gov/servlets/purl/1372113/>, 2017.
- [4] Q. Liu, J. Logan, Y. Tian, H. Abbasi, N. Podhorszki, J. Y. Choi, S. Klasky, R. Tchoua, J. Lofstead, R. Oldfield, M. Parashar, N. Samatova, K. Schwan, A. Shoshani, M. Wolf, K. Wu, W. Yu, "Hello ADIOS: The Challenges and Lessons of Developing Leadership Class I/O Frameworks", Concurrency & Computation: Practice and Experience, v.26, n.7, pp. 1453-1473, 2013.
- [5] M. Dreher, Bruno Raffin, "A Flexible Framework for Asynchronous In Situ and In Transit Analytics for Scientific Simulations", ACM/IEEE International Symposium on Cluster, Cloud and Grid Computing (CCGRID), Chicago, IL, 2014.
- [6] O.C. Marcu, A. Costan, G. Antoniu, M. Pérez-Hernández, B. Nicolae, et al.. "KerA: Scalable Data Ingestion for Stream Processing". ICDCS 2018 – 38th IEEE International Conference on Distributed Computing Systems, Vienna, Austria, pp.1480-1485, 2018,

Compétences

- An excellent Master degree in computer science or equivalent
- Strong knowledge of computer networks and distributed systems
- Knowledge on storage and (distributed) file systems
- Ability and motivation to conduct high-quality research, including publishing the results in relevant venues
- Strong programming skills (e.g. C/C++, Java, Python).
- Very good communication skills in oral and written English.
- Open-mindedness, strong integration skills and team spirit.
- Working experience in the areas of Big Data management, Cloud computing, HPC, is an advantage.

Avantages

- Restauration subventionnée

- Transports publics remboursés partiellement
- Congés: 7 semaines de congés annuels + 10 jours de RTT (base temps plein)
- Équipements professionnels à disposition (visioconférence, prêts de matériels informatiques, etc.)
- Prestations sociales, culturelles et sportives (Association de gestion des œuvres sociales d'Inria)
- Accès à la formation professionnelle

Rémunération

Rémunération mensuelle brute de 1982 euros les deux premières années et 2085 euros la troisième année

Informations générales

- **Thème/Domaine** : Calcul distribué et à haute performance
Système & réseaux (BAP E)
- **Ville** : Rennes
- **Centre Inria** : [Centre Inria de l'Université de Rennes](#)
- **Date de prise de fonction souhaitée** : 2020-10-01
- **Durée de contrat** : 3 ans
- **Date limite pour postuler** : 2020-06-30

Contacts

- **Équipe Inria** : [KERDATA](#)
- **Directeur de thèse** :
Antoniou Gabriel / gabriel.antoniou@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.

L'essentiel pour réussir

- An excellent Master degree in computer science or equivalent
- Strong knowledge of computer networks and distributed systems
- Knowledge on storage and (distributed) file systems
- Ability and motivation to conduct high-quality research, including publishing the results in relevant venues
- Strong programming skills (e.g. C/C++, Java, Python).
- Very good communication skills in oral and written English.
- Open-mindedness, strong integration skills and team spirit.
- Working experience in the areas of Big Data management, Cloud computing, HPC, is an advantage.

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

Merci de déposer en ligne CV, lettre de motivation et éventuelles recommandations

Pour plus d'information, contactez gabriel.antoniou@inria.fr

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.