



Offre n°2020-02409

Doctorant F/H Uniform Cloud and Edge Stream Processing for Fast and Intelligent Big 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)
- **Expected start date:** October 1st, 2020
- **Application deadline:** as early as possible, no later than March 10, 2020

Location and Mobility

The thesis will be mainly hosted by the KerData team at Inria Rennes Bretagne Atlantique. It will include collaborations with **Politehnica University of Bucharest** and the **University of Düsseldorf**. 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, 2020
- **Application deadline:** as early as possible, no later than March 10, 2020

Description

Context

The recent spectacular rise of the **Internet of Things (IoT)** and the associated augmentation of the data deluge motivated the emergence of **Edge computing** [1] as a means to distribute processing from centralized Clouds towards decentralized processing units close to the data sources. The key idea is to leverage computing and storage resources at the “edge” of the network, i.e., near the places where data is produced (e.g., sensors, routers, etc.). They can be used to filter and to pre-process data or to perform (simple) local computations (for instance, a home assistant may perform a first lexical analysis before requesting a translation to the Cloud).

This scheme led to new challenges regarding the ways to distribute processing across Cloud-based, Edge-based or hybrid Cloud/Edge-based infrastructures. State-of-the-art approaches advocate either “100% Cloud” or “100% Edge” solutions. In the former case, a plethora of Big Data stream processing engines like Apache Spark [2] and Apache Flink [3] emerged for data analytics and persistence, using the Edge devices just as proxies only to forward data to the Cloud. In the latter case, Edge devices are used to take local decisions and enable the real-time promise of the analytics, improving the reactivity and “freshness” of the results. Several Edge analytics engines emerged lately (e.g. Apache Edgent [4], Apache Minifi [5]) enabling basic local stream processing on low performance IoT devices. The relative efficiency of a method over the other may vary. Intuitively, it depends on many parameters, including network technology, hardware characteristics, volume of data or computing power, processing framework configuration and application requirements, to cite a few.

Problem statement

Hybrid approaches, combining both Edge and Cloud processing require deploying two such different processing engines, one for each platform. This involves two different programming models, synchronization overheads between the two engines and empirical splits of the processing workflow schedule across the two infrastructures, which lead eventually to sub-optimal performance.

Thesis goal

This PhD thesis aims to propose a *unified stream processing model for both Edge and Cloud platforms*. In particular, the thesis will seek potential answers to the following research question: how much can one improve (or degrade) the performance of an application by performing computation closer to the data sources rather than keeping it in the Cloud? To this end, the PhD thesis will first devise a methodology to understand the performance trade-offs of Edge-Cloud executions, and then design a unified processing model capable of exploiting the semantics of both platforms. The model will be implemented and experimentally evaluated with representative real-life stream processing use-cases executed on hybrid Edge-Cloud testbeds. The high-level goal of this thesis is to enable the usage of a large spectrum of Big Data analytics techniques at extreme scales, to support fast decision making in real-time.

Principales activités

Target use-case

The unified processing model will be evaluated using the requirements of a real-life production application, provided by our partners from **University Politehnica of Bucharest**. The unified model will be used as the processing engine of the **MonALISA** [6] monitoring system of the **ALICE** experiment at **CERN** [7]. ALICE (A Large Ion Collider Experiment) is one of the four LHC (Large Hadron Collider) experiments run at CERN

(European Organization for Nuclear Research). ALICE collects data at a rate of up to 4 Petabytes per run and produces more than 109 data files per year. Tens of thousands of CPUs are required to process and analyze them. More than 350 MonALISA services are running at sites around the world, collecting information about ALICE computing facilities, local- and wide-area network traffic, and the state and progress of the many thousands of concurrently running jobs. Currently, all the monitoring and alerting logic is implemented in the Cloud, with high latency. With the proof-of-concept envisioned by this thesis, the goal is to enable Edge/Cloud monitoring data processing and to enable faster alerts and decision making, as soon as the data is collected.

Enabling technologies

In the process of designing the unified Edge-Cloud 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 **KerA** [8] approach for Cloud-based low-latency storage for stream processing (currently under development at Inria, in collaboration with **Universidad Politécnica de Madrid**, 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.
- The **Planner** [9] middleware for cost-efficient execution plans placement for uniform stream analytics on Edge and Cloud. Planner automatically selects which parts of the execution graph will be executed at the Edge in order to minimize the network cost. Real-world micro-benchmarks show that Planner reduces the network usage by 40% and the makespan (end-to-end processing time) by 15% compared to state-of-the-art.

References

- [1] M. Satyanarayanan, "The emergence of edge computing," *Computer*, 2017.
- [2] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. "Spark: cluster computing with working sets". In *Proceedings of the 2nd USENIX conference on Hot topics in cloud computing (HotCloud'10)*. USENIX Association, Berkeley, CA, USA, 2010.
- [3] P. Carbone, S. Ewen, G. Fóra, S. Haridi, S. Richter, and K. Tzoumas. "State management in Apache Flink: consistent stateful distributed stream processing". *Proc. VLDB Endow.* 10, 12 2017, 1718-1729.
- [4] Apache Edgent, <http://edgent.apache.org>, accessed on January 2019
- [5] Apache Minifi, <https://nifi.apache.org/minifi/>, accessed on January 2019
- [6] I. Legrand, R. Voicu, C. Cirstoiu, C. Grigoras, L. Betev, and A. Costan. "Monitoring and control of large systems with MonALISA". *Communications of the ACM* 52, 9, September, 2009, 49-55.
- [7] The ALICE LHC Experiment, <https://home.cern/science/experiments/alice>, accessed on January 2019.
- [8] 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,
- [9] L.Proserpi, A.Costan, P.Silva and G.Antoniu, "Planner:Cost-efficient Execution Plans Placement for Uniform Stream Analytics on Edge and Cloud," in *WORKS 2018: 13th Workflows in Support of Large-Scale Science Workshop*, held in conjunction with the IEEE/ACM SC18 conference, 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** :
Antoni 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.