



Offer #2024-08203

Post-Doctoral Research Visit F/M Resource allocation and scheduling for data stream processing in shared Fog environments

Contract type : Fixed-term contract

Level of qualifications required : PhD or equivalent

Fonction : Post-Doctoral Research Visit

About the research centre or Inria department

The Inria Centre at Rennes University is one of Inria's eight 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 PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc

Context

Financial and working environment

This post-doctoral position is part of the PEPR Cloud - Taranis project funded by the French government (France 2030). 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 close collaboration with the DiverSE team and other partners in the Taranis project.

Assignment

Context:

The mutual low-latency objective for both Data Stream Processing (DSP) and Fog environments has resulted in a continuous growth of DSP deployments on Fogs [1]. However, the success of the deployment and running of stream data applications in the Fog relies on how to efficiently allocate resources and schedule tasks (operators) to achieve the desired performance. Previous efforts on deploying stream data applications in the Fog have focused on reducing the volume of communication overhead between nodes (inter-node communication) and dividing the computation between Fog servers and Clouds [2, 3]. Unfortunately, they are oblivious to (1) the dynamic nature of data streams (i.e., data volatility and bursts) and to (2) the bandwidth and resource heterogeneity in the Fog, which negatively affects the performance of stream data applications [4][5].

Objectives:

The goal of this postdoc project is to investigate how to optimize resource and task allocation when deploying data streaming processing applications in the Fog. In particular, we want to investigate new optimization metrics and objectives when deploying streaming processing applications in the Fog, including latency, throughput, and maximum sustainable throughput. Accordingly, we will develop a new scheduling framework that relies, among others, on Machine Learning/Deep Learning models to decide on resource allocation and operator placement at runtime (based on the collected data and given the cost model of redeployment and process migration). The proposed framework will be integrated in one of state-of-the-art data stream engines such as Flick [6], Storm [7] or Spark [8] and evaluated at large-scale using syntactic applications and real-world stream data application.

References :

[1] Noghabi, Shadi A., Landon Cox, Sharad Agarwal, and Ganesh Ananthanarayanan. "The emerging landscape of edge computing." *GetMobile: Mobile Computing and Communications* 23, no. 4 (2020): 11-20.

[2] Nardelli, Matteo, Valeria Cardellini, Vincenzo Grassi, and Francesco Lo Presti. "Efficient operator placement for distributed data stream processing applications." *IEEE Transactions on Parallel and Distributed Systems* 30, no. 8 (2019): 1753-1767.

[3] Renart, Eduard Gibert, Alexandre Da Silva Veith, Daniel Balouek-Thomert, Marcos Dias De Assunção, Laurent Lefevre, and Manish Parashar. "Distributed operator placement for IoT data analytics across edge and cloud resources." In *2019 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, pp. 459-468. IEEE, 2019.

[4] Lambert, Thomas, David Guyon, and Shadi Ibrahim. "Rethinking operators placement of stream data application in the edge." In Proceedings of the 29th ACM International Conference on Information & Knowledge Management, pp. 2101-2104. 2020.

[5] Arsalane, Khaled, Guillaume Pierre, and Shadi Ibrahim. "Toward Stream Processing Elasticity in Realistic Geo-Distributed Environments." In IC2E 2024-12th IEEE International Conference on Cloud Engineering, pp. 1-9. 2024.

[6] "Apache flink," <https://flink.apache.org>.

[7] Apache Storm. 2020. <https://storm.apache.org/>

[8] Zaharia, Matei, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica. "Discretized streams: Fault-tolerant streaming computation at scale." In *Proceedings of the twenty-fourth ACM symposium on operating systems principles*, pp. 423-438. 2013.

Main activities

- Read and synthesize literature work.
- Design new resource allocation and scheduling policies for data stream processing in the Fog.
- 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.

Skills

- 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

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

he postdoctoral researcher will receive a gross monthly salary of 2,788 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** : 2024-12-01
- **Duration of contract** : 1 year, 6 months
- **Deadline to apply** : 2024-12-25

Contacts

- **Inria Team** : [MAGELLAN](#)
- **Recruiter** : 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.

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.