



Offer #2025-09183

Post-Doctoral Research Visit F/M Data management and job scheduling for Geo-distributed Workflows

Contract type : Fixed-term contract

Renewable contract : Yes

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 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 PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc.

Context

Financial and working environment.

This postdoc position will be in the context of IPCEI-CIS (Important Project of Common European Interest – Next Generation Cloud Infrastructure and Services) DXP (Data Exchange Platform) project involving Amadeus and three Inria research teams (COAST, CEDAR and MAGELLAN). This project aims to design and develop an open-source management solution for a federated and distributed data exchange platform (DXP), operating in an open, scalable, and massively distributed environment (cloud-edge continuum). 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, with the possibility of an extension to 24 months.

Assignment

Context:

Data are usually hosted in multiple geo-distributed locations (private and public clouds, distributed caches and across the edge-to-cloud continuum). Collectively processing these data is a must (e.g. distributed data analysis and queries), but this presents several challenges to existing data-intensive distributed workflow frameworks (e.g. MapReduce [1], Spark [2], TensorFlow [3] and Dataflow [4]). This is due to the low capacity of wide area network (WAN) links, as well as the heterogeneity of networks, computation power and monetary cost in geo-distributed environments [5].

Much effort has devoted on optimizing the performance of geo-distributed workflows [5, 6, 7, 8, 9]. These efforts mainly focus on reducing cross-data center data transfer and optimal task placement according to performance heterogeneity of the data centers [5, 6, 7]. However, these efforts do not consider the storage services heterogeneity [10] (input and intermediate data are stored on different devices), monetary cost heterogeneity in terms of computing, storage and network, or multi-tenancy when multiple workflows run concurrently.

Objectives:

We will study the interplay and correlation between the different factors that contribute to the performance of geo-distributed workflows (i.e., input data, intermediate data, number of iterations, capacity of site, etc). Accordingly, we will design scheduling policies and associated data movement to improve the overall performance, monetary cost, and resource utilization by considering the input location, the network cost and status, intermediate data, and the capacity of the different sites, when scheduling multiple workflows across massively distributed environments.

[1] J. Dean and S. Ghemawat, “MapReduce: Simplified Data Processing on Large Clusters,” *Commun. ACM*, pp. 107–113, Jan. 2008.

[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, 2010, pp. 10–10.

[3] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: a system for large-scale machine learning. In Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation (OSDI'16).

[4] Tyler Akidau, Robert Bradshaw, Craig Chambers, Slava Chernyak, Rafael J. Fernández-Moctezuma, Reuven Lax, Sam McVeety, Daniel Mills, Frances Perry, Eric Schmidt, and Sam Whittle. 2015. The dataflow model: a practical approach to balancing correctness, latency, and cost in massive-scale, unbounded, out-of-order data processing. Proc. VLDB Endow.

[5] Qifan Pu, Ganesh Ananthanarayanan, Peter Bodik, Srikanth Kandula, Aditya Akella, Paramvir Bahl, and Ion Stoica. Low latency geo-distributed data analytics. In Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM'15, pages 421–434, New York, NY, USA, 2015. ACM.

[6] Sharad Agarwal, John Dunagan, Navendu Jain, Stefan Saroiu, Alec Wolman, and Harbinder Bhogan. Volley: Automated data placement for geo-distributed cloud services. In Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation, NSDI'10, pages 2–2, Berkeley, CA, USA, 2010. USENIX Association.

[7] WANG, Qingyuan, GAO, Bin, ZHOU, Zhi, *et al.* Dag-aware optimization for geo-distributed data analytics. In: *Proceedings of the 52nd International Conference on Parallel Processing*. 2023. p. 472-481.

[8] Ashish Vulimiri, Carlo Curino, Philip Brighten Godfrey, Thomas Jungblut, Konstantinos Karanasos, Jitendra Padhye, and George Varghese. Wanalytics: Geo-distributed analytics for a data intensive world. In Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, SIGMOD'15, pages 1087–1092, New York, NY, USA, 2015. ACM.

[9] Chien-Chun Hung, Leana Golubchik, and Minlan Yu. Scheduling jobs across geo-distributed datacenters. In Proceedings of the Sixth ACM Symposium on Cloud Computing, SoCC'15, pages 111–124, New York, NY, USA, 2015. ACM.

[10] Hao Wu, Junxiao Deng, Hao Fan, Shadi Ibrahim, Song Wu, Hai Jin, "QoS-Aware and Cost-Efficient Dynamic Resource Allocation for Serverless ML Workflows," *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, St. Petersburg, FL, USA, 2023, pp. 886-896.

Main activities

- Read and synthesize literature work.
- Design new scheduling policies and data management for Geo-distributed Workflows
- 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

Monthly gross salary amounting to 2788 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** : 2025-11-01
- **Duration of contract** : 12 months
- **Deadline to apply** : 2025-09-18

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.