



Offer #2025-08998

PhD Position F/M Modelization of HPC Jobs and Resources to Minimize Energy Waste

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

About the research centre or Inria department

The Centre Inria de l'Université de Grenoble groups together almost 600 people in 24 research teams and 9 research support departments.

Staff is present on three campuses in Grenoble, in close collaboration with other research and higher education institutions (Université Grenoble Alpes, CNRS, CEA, INRAE, ...), but also with key economic players in the area.

The Centre Inria de l'Université Grenoble Alpes is active in the fields of high-performance computing, verification and embedded systems, modeling of the environment at multiple levels, and data science and artificial intelligence. The center is a top-level scientific institute with an extensive network of international collaborations in Europe and the rest of the world.

Context

co-advised by Raphaël Bleuse and Eric Rutten (Ctrl-A), LIG/INRIA, and

Franck Corset (LJK, ASAR)

Within the framework of the Taranis project in the PEPR Cloud.

Assignment

Soberness—in terms of electrical power—of Data Centers and High-Performance Computing (HPC) systems is becoming an important design issue, as the global energy consumption of Information Technologies is rising at considerable levels. Large-scale computing infrastructures are processing vaster amount of data or solving problems requiring vaster amount of computing power. The behavior of large scale infrastructures has become more variable and difficult to model, especially with respect to power consumption and application performance. Therefore, dealing with time variations and unpredictable disturbances demands to automate the management (i.e., configuration) of the infrastructures. This automatic management can be done by periodically monitoring the state of the system, and updating the configuration to activate relevant mechanisms.

This work takes root in the field of autonomic computing [6], and aims at designing efficient feedback loops to automatically manage the resources of a HPC (high performance computing) infrastructure. The use of feedback loops is widespread in various fields of engineering, but recent in the computer science field.

Main activities

Resource harvesting

The Resource and Job Management System (*RJMS*) is a key component to operate a HPC system [5, 7]. Users submit *jobs*: a description of their computation, data, and resource requirements. With respect to the resources status reported by the resource manager, the scheduler decides which resources to allocate to a job and assigns a time slot for the job's execution.

The RJMS is however unable to fully exploit the resources in a HPC cluster: the presence of unused resources results from the limits of the scheduling of jobs. The allocation of resources to jobs, while respecting all constraints, leaves resources idle. Such inefficiency is sometimes referred as *fragmentation*. The loss of computing power resulting from the fragmentation represents an exploitable pool of resource.

CiGri [1] is a lightweight, scalable and fault-tolerant grid system that plugs into the RJMS. CiGri aims at minimizing the waste of computing resources to leverage the pool of unused resources. Yet, some computations (best-effort computations) can still lead to wasted resources if they are stopped before completion. In particular, by integrating information from the RJMS one could avoid to start best-effort computations if there is not enough time to execute them.

Modeling jobs and resources

The authors in [4] show that the jobs execution times are not necessary deterministic and can be modeled by an Exponential, Weibull, log-Normal or Normal distributions. Moreover, the high variance of execution times and the diversity of the jobs, whether in terms of the nature of the data, the application domain or the size of the problem have to be taking into account, by considering for instance a mixture of distributions (see [3]). Thus, we propose to apply some Machine Learning

techniques, e.g., a clustering of all jobs submitted during the last decade, in order to take into account this heterogeneity. In a second step, we propose to take this information as a prior distribution in a Bayesian setting in order to improve the accuracy of the execution times estimations (see [8]). Furthermore, prior work mostly focuses on models that are job-centric and based on post-execution data [2]: they neglect to model the resources used for the computation.

In this work, we want to model the availability of the platform in order to insert the best-effort jobs in a frugal way. The design of models will have to consider that the allocation decisions are taken online, with partial information unveiling during the platform life-cycle.

References

- [1] Bruno Bzeznik and Ghislain Charrier, *CiGri*. lic: GPL-3.0-or-later. url: <http://cigri.imag.fr/>, vcs: <https://github.com/oar-team/cigri>.
- [2] Dror G. Feitelson, Dan Tsafir, and David Krakov. “Experience with using the Parallel Workloads Archive.” In: *J. Parallel Distributed Comput.* 74.10 (Oct. 2014), pp. 2967–2982. doi: [10.1016/J.JPDC.2014.06.013](https://doi.org/10.1016/J.JPDC.2014.06.013).
- [3] Sylvia Fruhwirth-Schnatter, Gilles Celeux, and Christian P Robert. *Handbook of mixture analysis*. CRC press, 2019.
- [4] Ana Gainaru, Hongyang Sun, Guillaume Aupy, Yuankai Huo, Bennett A. Landman, and Padma Raghavan. “On-the-fly scheduling versus reservation-based scheduling for unpredictable workflows.” In: *Int. J. High Perform. Comput. Appl.* 33.6 (2019). doi: [10.1177/1094342019841681](https://doi.org/10.1177/1094342019841681).
- [5] Yiannis Georgiou. “Contributions for Resource and Job Management in High Performance Computing.” PhD thesis. LIG, Univ. Grenoble Alpes, France, Nov. 2010. url: <https://tel.archives-ouvertes.fr/tel-01499598> (visited on 2023-10-11).
- [6] Jeffrey O. Kephart and David M. Chess. “The Vision of Autonomic Computing.” In: *IEEE Computer* 36.1 (Jan. 2003), pp. 41–50. doi: [10.1109/MC.2003.1160055](https://doi.org/10.1109/MC.2003.1160055).
- [7] Albert Reuther et al. “Scalable system scheduling for HPC and big data.” In: *Journal of Parallel and Distributed Computing* 111 (Jan. 2018), pp. 76–92. doi: [10.1016/j.jpdc.2017.06.009](https://doi.org/10.1016/j.jpdc.2017.06.009).
- [8] Christian P Robert et al. *The Bayesian choice: from decision-theoretic foundations to computational implementation*. Vol. 2. Springer.

Skills

The PhD candidate must have:

- a MSc degree in Computer science or Statistics.
- skills in programming languages, software engineering
- Knowledge in the domains of HPC

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

Base of 2200 euros gross / month

General Information

- **Theme/Domain** : Distributed Systems and middleware System & Networks (BAP E)
- **Town/city** : Grenoble
- **Inria Center** : Centre Inria de l'Université Grenoble Alpes
- **Starting date** : 2025-10-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2025-07-16

Contacts

- **Inria Team** : CTRL-A
- **PhD Supervisor** :
Rutten Eric / eric.rutten@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

Applications must be submitted online via the Inria website. Processing of applications submitted via other channels is not guaranteed.

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.