



Offer #2024-07244

R&D Engineer in Exascale High-Performance Computing - Damaris

Contract type : Fixed-term contract

Renewable contract : Yes

Level of qualifications required : PhD or equivalent

Fonction : Temporary scientific engineer

About the research centre or Inria department

Inria, the French national research institute for the digital sciences, promotes scientific excellence and technology transfer to maximise its impact.

It employs 2,400 people. Its 200 agile project teams, generally with academic partners, involve more than 3,000 scientists in meeting the challenges of computer science and mathematics, often at the interface of other disciplines.

Inria works with many companies and has assisted in the creation of over 160 startups.

It strives to meet the challenges of the digital transformation of science, society and the economy.

Context

About Inria, the team and the position

Inria is the only French public research body fully dedicated to computational sciences. Inria's missions are to produce outstanding research in the computing and mathematical fields of digital sciences and to ensure its impact on the economy and society through technology transfer and innovation. Throughout its research centers and its approximately 200 project teams, Inria has a workforce of 3 400 scientists with an annual budget of 265 million euros, 29% of which coming from its own resources. The Inria Research Center at Rennes University is one of the ten sites of Inria. This publicly funded research center has a workforce of about 620 people, including full-time research scientists, faculty staff, engineers and support staff, distributed in 33 teams and support services.

The hired engineer will be a member of the KerData Inria team (<https://team.inria.fr/kerdata/>) led by Gabriel Antoniu. KerData is a joint research team of Inria's Research Centre at Rennes University and INSA Rennes, and also a team of the IRISA lab. KerData's main research activities address the area of distributed data management at challenging scales, with a current focus on pre-Exascale/Exascale HPC supercomputers, clouds and edge-based systems and on hybrid combinations of those. In particular, we address the needs of data-intensive high-performance applications. For this position, the engineer will be working on projects associated with Damaris, a library developed by the KerData research team as a result of a collaboration within JLESC (Joint INRIA-ANL-UIUC-BSC Lab for Extreme Scale Computing: <https://jlesc.github.io/>). Damaris is a middleware for managing I/O and in situ processing of Big Data on HPC infrastructures.

The position is made available in the context of the Exa-DoST project (Data-oriented Software and Tools for the Exascale) of the NumPEX National Programme: <https://numpex.org/exadost-data-oriented-software-and-tools-for-the-exascale/>. NumPEX aims to build the software infrastructure for the first Exascale supercomputer expected to be installed in France in 2025 (Jules Verne project).

Assignment

Mission overview

By joining our team you will participate in a dynamic work environment with exceptionally talented and friendly coworkers who are committed to high-quality research and development practices. You will collaborate with esteemed researchers from around the world by taking the technical responsibility for the development of the Damaris software, with the following global missions:

1. Maintain Damaris as a distributable, professional-quality software (continuous integration, technical support, documentation, management of the web site);
2. Contribute to the design of new and improved features for Damaris. This may include enabling new and improved data processing server-side plugins, designing a Python client-side library, designing and testing improved process placement capabilities, support for dynamically triggered in situ analysis, support for GPU-based analytics, etc.
3. Contribute to project work and dissemination actions by interacting with potential users, in particular in the context of the NumPEX National programme (<https://numpex.org/>) and the EUPEX

- EuroHPC JU project (<https://eupex.eu/>) commitments.
4. Perform large-scale experiments with Damaris (while making the necessary extensions) to support its efficient execution on emerging pre-Exascale/Exascale platforms (such as EUPEX or Jules Verne, the first Exascale supercomputer expected to be installed in France).

Main activities

Detailed missions

1. Improve and extend the Damaris code, perform robustness and performance tests, maintain a continuous code integration process;
2. Create a unified data processing framework combining Damaris (as in situ/in transit data processing framework) with Big Data analytics plugins to support batch-based or stream-based data processing (e.g., currently based on components from the Python Dask ecosystem, however adding capabilities with other big-data ecosystems is of interest);
3. Adapt Damaris to use recent advanced communication and multithreading technologies (examples: MPI_Sessions, ANL Mercury, Argobots);
4. Develop and enhance software connectors allowing Damaris to use state-of-the-art visualization tools such as VisIt, ParaView and Ascent, and interface with other in situ libraries, such as PDI and Melissa.
5. Extend Damaris to support other programming languages (e.g. creation of Python and/or Julia bindings).
6. Facilitate the dissemination of Damaris and of the unified data processing framework through the following means:
 - Design and implementation of software demonstrators in collaboration with users and project work packages;
 - Extend existing example codes to facilitate learning of the interface by new users;
 - Maintain a professional-quality website facilitating the distribution of the code and of its documentation (reference manual, user manual);
 - Make demos at forums such as the Supercomputing conference, the main international forum of the HPC community;
 - Create and animate a user community around the library (maintain a mailing list, a mechanism for bug report and solving, user support, etc.).

Skills

Required qualifications

- Excellent, demonstrated programming skills in C, C++, Python;
- Knowledge of hardware and software technologies in the area of HPC including MPI (which is a must), resource managers such as SLURM/PBS/OAR/etc., module system, multi-core and GPU based libraries and programming interfaces, use of parallel debuggers;
- Experience with Linux operating system, software code repository and build systems (git, GitLab/GitHub/Bitbucket, CMake, Spack);
- Some familiarity with big-data tools (Dask, Spark, Flink)
- Knowledge of methodologies for managing software projects;
- Ability to analyze and synthesize user requirements;
- Ability to develop benchmarking suites, analyze and present results.
- Ability to communicate and work in collaboration with experts in the same area and in other areas, in English;
- Taste for transmitting and sharing knowledge, results, progress with the facility to present results in written and oral form.
- Autonomy in leading and performing the tasks;
- Sense of partnership and team spirit.

Benefits package

- Subsidised catering service
- Partially-reimbursed public transport

Remuneration

monthly gross salary from 2695 euros according to diploma and experience

General Information

- **Theme/Domain** : Distributed and High Performance Computing
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2024-05-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2024-06-30

Contacts

- Inria Team : [KERDATA](#)
- Recruiter :
Antoni Gabriel / gabriel.antoniou@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

For more information, please contact gabriel.antoniou@inria.fr

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.