



Offer #2024-07213

PhD Position F/M Dynamic in situ and in transit data analysis for Exascale Computing using Damaris

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Other valued qualifications : Master's degree

Fonction : PhD Position

About the research centre or Inria department

The Inria Rennes - Bretagne Atlantique Centre is one of Inria's eight centres and has more than thirty research teams. The Inria Center 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 SMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc.

Context

Context

The field of high-performance computing has reached a new milestone, with the world's most powerful supercomputers exceeding the exaflop threshold. These machines will make it possible to process data on an unprecedented scale, enabling simulations of complex phenomena to be carried out with superior precision in a wide range of application fields: astrophysics, particle physics, healthcare, genomics, and more. By way of example, it is estimated that the SKA project will process **one exabyte of raw data per day**. In France, the **installation of the first Exascale supercomputer** is scheduled for 2025. Major players in the French scientific community in the field of high-performance computing (HPC) have joined forces within the PEPR NumPEX program (<https://numpex.fr/>) to carry out research aimed at contributing to the design and implementation of this machine's software infrastructure. This thesis is part of the Exa-DoST project of NumPEX, focusing on **Exascale data management challenges**.

Logo Inria



PhD Advisors

- Gabriel Antoniu (Inria)
- Laurent Colombet (CEA)
- Julien Bigot (Maison de la Simulation, CEA)

Location and Mobility

The thesis, co-supervised by Inria and CEA, will be hosted by the KerData team at the **Inria Research Center at Rennes University** and will include regular visits at CEA at Bruyères-le-Châtel. 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.

This thesis will also include **collaborations with international partners**, especially from the USA.

The KerData team in a nutshell for candidates

- KerData is a human-sized team currently comprising 5 permanent researchers, 2 contract researchers, 1 engineer and 5 PhD students. You will work in a caring environment, offering a **good work-life balance**.
- KerData is leading multiple projects in top-level national and international collaborative environments such as within the **Joint-Laboratory on Extreme-Scale Computing** <https://jlesc.github.io>. Our team has active collaboration with high-profile academic institutions all around the world (including the USA, Spain, Germany or Japan) and with industry.
- Our team strongly favors **experimental research**, validated by implementation and experimentation of software prototypes with real-world applications on real-world platforms including some of the most powerful supercomputers worldwide.
- 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 researchers.
- Check our **website** for more about the KerData team here <https://team.inria.fr/kerdata/>

Assignment

Introduction

Without a major change in practices, the increased computing capacity of the next generation of computers will lead to an explosion in the volume of data produced by numerical simulations. Managing this data, from production to analysis, is a major challenge.

The use of simulation results is based on a well-established calculation-storage-calculation protocol. The difference in capacity between computers and file systems makes it inevitable that the latter will be clogged. While it is not conceivable to do without a storage system, many experiments are aimed at reducing its use. In this field, the emergence of a varied vocabulary reflects the diversity of approaches: in situ processing, in transit processing, staging nodes, helper cores. All of these approaches underline the desire to replace the usual write-read process for linking applications with in-flight intervention. Analysis carried out at the same time as simulation is a capability of particular interest to CEA physicists. This need has led to the first implementations of in situ or in transit analysis systems in simulation codes, and to the creation of specific middleware such as Damaris [1,2,3,4].

Technological and application context

Damaris (<https://project.inria.fr/damaris/>) is a middleware for I/O management and real-time processing of data from large-scale MPI-based HPC simulations. It initially proposed to dedicate cores to asynchronous I/O in the multicore nodes of recent HPC platforms, focusing on ease of integration into existing simulations, efficient use of resources (thanks to the use of shared memory) and simplicity of extension via plug-ins. Over the years, Damaris has evolved into a more elaborate system, offering the possibility of using dedicated cores or nodes to perform in situ data processing and visualization. It offers a seamless connection to VisIt or ParaView visualization software to enable in situ visualization with minimal impact on simulation runtime. Damaris provides an extremely simple user interface (API) and can be easily integrated into existing large-scale simulations. It has been validated up to 14,000 cores on supercomputers such as Titan (1st in the Top500 at the time of the experiments), Jaguar, Kraken, etc., with numerous simulation codes. Damaris is one of the software building blocks to be used on the first exaflop-scale supercomputer to be installed in France.

PDI is an API that enables weak coupling between simulation codes and data management libraries for intensive computing. The approach consists in instrumenting the codes to identify where and when data becomes present in memory, and where and when the memory will be reused to store new values. This instrumentation is entirely independent of the data management libraries used. A separate file in Yaml format is used to specify what to do with the data in the code and with which library. This approach corresponds to aspect-oriented programming. The various aspects that can be taken into account in this way are numerous and include, for example, reading code parameters, data initialization, in-situ post-processing, visualization or storage of results on disk, fault tolerance, inclusion in a code coupling or in an overall simulation. All these aspects can be managed, each thanks to a plugin giving access to a different dedicated library. PDI merely arbitrates between these different plug-ins, but offers no functionality of its own, such as the provision of dedicated cores in Damaris. PDI is used in production codes such as Gysela, which uses the most powerful Petaflop machines available today, such as Fugaku, Adastra and Exa1-HF. PDI's intra-process architecture ensures that the scalability offered is exactly that of the libraries used on the back-end.

Challenges

Initial feedback from application users, using an in situ system, clearly shows the need to develop a system that can dynamically manage the addition or deletion of analyses during the execution of a simulation. For example, during a simulation study of a material's behavior under stress, different observations and analyses are frequently requested as the simulation progresses. Indeed, the elastoplastic properties of a material may change over time, triggering new analyses to understand underlying physical phenomena such as dislocation propagation and possible phase changes (solid vs. liquid or solid vs. solid). In this context, to save time and computational resources, it is important to trigger the activation of new analyses at the right moment during the simulation run. Note that the event can be detected either by the simulation code or by an analysis. Of course, in order to maintain high performance results, it is essential to manage the placement of new analyses on GPU nodes. These dynamic analysis management capabilities are not yet effectively available in Damaris.

Coddex is a simulation code that solves the equations of continuum mechanics in dynamic hyperelasticity (shocks or rapid loading). It also incorporates the description of behavioral discontinuities (cf. Figure 1) of change or maillage. Coddex stands for Code de Dynamique des Discontinuités pour l'Étude des cristaux.




Figure 1: Deformation map of a TATB polycrystal, an ultra-anisotropic energetic material, using the Coddex code

Example scenarios for implementing the in-situ system in Coddex via Damaris

- Programmed" analyses: the physicist user defines a complete physical simulation and a list of in-situ analyses (Coddex-Damaris-Paraview links) with variable execution frequencies. In this scenario, outputs are produced without operator assistance (programmed outputs), in the form of statistics files or 2D (images) or 3D visuals (e.g. 3D iso-surface outputs by Paraview in the form of .obj files).
- On-the-fly" analyses: the user can also (and independently) launch new analyses via the pause system (waiting for requests from the orchestrator). An analysis triggers the creation and implementation of a new analysis with a specific frequency. New data outputs (initially unplanned) are produced. The list of analyses can be modified on the fly in this way. This "on-the-fly" scenario can then be incorporated into a "scheduled" analysis.
- Triggered" analyses: the code detects a discontinuity (creation of a new phase, for example) and triggers an appropriate analysis, selected from a bank of typical analyses.

Objectives

The research work proposed in this thesis consists in designing an innovative model for the dynamic management of in situ and in transit analyses, proposing its implementation in the Damaris middleware and validating it with simulations performed using the Coddex code.

Main activities

After studying the state of the art and getting to grips with the Damaris architecture and Coddex code, the candidate will study, propose and develop innovative solutions, which he or she will publish in the best journals and conferences in the field. The candidate will work in a multidisciplinary environment (computer science and physics) thanks to the INRIA-CEA collaboration within the Exa-DoST project of the NumPEX PEPR, and will have privileged access to very large-scale computers for experimentation.

References

- [1] M. Dreher, B. Raffin, "A Flexible Framework for Asynchronous In Situ and In Transit Analytics for Scientific Simulations", 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (2014).
- [2] M. Dorier, G. Antoniu, F. Cappello, M. Snir, and L. Orf, "Damaris : How to Efficiently Leverage Multicore Parallelism to Achieve Scalable, Jitter-free I/O", in CLUSTER – IEEE. International Conference on Cluster Computing. IEEE, Sep. 2012.
- [3] M. Dorier, M. Dreher, T. Peterka, J. Wozniak, G. Antoniu and B. Raffin, "Lessons Learned from Building In Situ Coupling Frameworks", in Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization, Austin, 2015.
- [4] E. Dirand, L. Colombet, B. Raffin, "TINS: A Task-Based Dynamic Helper Core Strategy for In Situ Analytics", in Proceedings of Asian Conference on Supercomputing Frontiers, Singapore 2018.

Skills

- An excellent Master degree in computer science or equivalent
- Strong knowledge of 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 (Python, C/C++)
- Working experience in the areas of HPC and Big Data management is an advantage
- Very good communication skills in oral and written English
- Open-mindedness, strong integration skills and team spirit

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Possibility of teleworking (90 days per year) and flexible organization of working hours
- Partial payment of insurance costs

Remuneration

monthly gross salary amounting to 2100 euros for the first and second years and 2200 euros for the third year

General Information

- **Theme/Domain** : Distributed and High Performance Computing
Scientific computing (BAP E)
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2024-09-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2024-07-20

Contacts

- **Inria Team** : [KERDATA](#)
- **PhD Supervisor** :
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.

The keys to success

The candidate will have to show motivation, autonomy and an ability to initiate links between the research activities carried out at Inria and at the CEA.

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.