



Offre n°2024-08263

Internship: Python Data Processing on Supercomputers for Large Parallel Numerical Simulations.

Le descriptif de l'offre ci-dessous est en Anglais

Type de contrat : Convention de stage

Niveau de diplôme exigé : Bac + 4 ou équivalent

Fonction : Stagiaire de la recherche

A propos du centre ou de la direction fonctionnelle

The Centre Inria de l'Université de Grenoble groups together almost 600 people in 23 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 Alpe 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.

Contexte et atouts du poste

The internship will take place at the DataMove team located in the IMAG building on the campus of Saint Martin d'Herès (Univ. Grenoble Alpes) near Grenoble, under the supervision of Bruno Raffin (bruno.raffin@inria.fr), Andres Bermeo (andres.bermeo-marinelli@inria.fr) and Yushan Wang (yushan.wang@cea.fr)

The length of the internship is *4 months minimum* and the start date is flexible, but need a 2 month delay before starting the interhsip due to administrative constraints. The DataMove team is a friendly and stimulating environment that gathers Professors, Researchers, PhD and Master students all leading research on High-Performance Computing. The city of Grenoble is a student-friendly city surrounded by the Alps mountains, offering a high quality of life and where you can experience all kinds of mountain-related outdoor activities.

Mission confiée

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 unprecedented quantities of data, which can be used to simulate complex phenomena with superior precision in a wide range of application fields: astrophysics, particle physics, healthcare, genomics, etc.

Without a significant 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. For instance, the Gysela code in production mode can produce up to 5TB of data per iteration. It is obvious that storing 5TB of data is not feasible at high frequency. What's more, loading this quantity of data for later analysis and visualization is also a difficult task. To bypass this difficulty, we choose to rely on the in situ data analysis approach.

In situ consists in coupling the parallel simulation code, Gysela for instance, with a data analytics code that processes the data online, as soon as produced. In situ enables reducing the amount of data to write to disk, limiting the pressure on the file system. This is a mandatory approach to run massive simulations like Gysela on the latest Exascale supercomputers.

We developed an in situ data processing approach, called Deisa, relying on Dask, a Python environment for distributed tasks. Dask defines tasks that are executed asynchronously on workers once their input data are available. The user defines a graph of tasks to be executed. This graph is then forwarded to the

Dask scheduler. The scheduler is in charge of (1) optimizing the task graph and (2) distributing the tasks for execution on the different workers according to a scheduling algorithm aiming at minimizing the graph execution time.

Deisa extends Dask so it becomes possible to couple a MPI-based parallel simulation code with Dask. Deisa enables the simulation code to directly send newly produced data into the worker memories, notify the Dask scheduler that these data are available for analysis and that associated tasks can then be scheduled for execution.

Compared to previous in situ approaches that are mainly MPI-based, our approach relying on Python tasks makes for a good tradeoff between programming ease and runtime performance.

The goal of this internship is to investigate solutions to improve task placement and thus performance enabling tasks to be scheduled in process (into the simulation processes), in situ (running on external processes but on the same compute nodes that also run the simulation code), in transit (on dedicated nodes different from the simulation nodes). Running closer to the simulation reduces the need for data movements, but can potentially steal resources (CPU, GPU, network, memory, cache) from the simulation and slow it down. Dask task graph optimization is a good starting point to develop such approaches.

References

1. Dask - <https://www.dask.org/>
2. Deisa Paper: Dask-enabled in situ analytics. Amal Gueroudji, Julien Bigot, Bruno Raffin. Hipc 2021. <https://hal.inria.fr/hal-03509198v1>
3. Deisa Paper: [Dask-Extended External Tasks for HPC/ML In Transit Workflows](#) Amal Gueroudji, Julien Bigot, Bruno Raffin, Robert Ross. Work workshop at Supercomputing 23. <https://hal.science/hal-04409157v1>
4. Deisa Code: <https://github.com/pdidev/deisa>
5. Ray - <https://github.com/ray-project/ray>
6. Damaris: How to Efficiently Leverage Multicore Parallelism to Achieve Scalable, Jitter-free I/O. Matthieu Dorier, Gabriel Antoniu, Franck Cappello, Marc Snir, Leigh Orf. IEEE Cluster 2012. <https://inria.hal.science/hal-00715252>

Principales activités

After getting familiar with the base concepts, existing work, deploying and running some simple in situ data processing with Deisa on a supercomputer, the candidate will investigate solutions to enabling a better task placement (in process, in situ, in transit), develop a prototype and experiments to measure the performance gain.

Compétences

Expected skills include

- Knowledge on distributed, parallel computing and numerical simulations.
- Python, Numpy, Parallel programming (MPI)
- English (working language)

Avantages

- 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 (90 days / year) and flexible organization of working hours
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage under conditions

Rémunération

€4.35 per hour of actual presence at 1 January 2024.

About 590€ gross per month (internship allowance)

Informations générales

- **Thème/Domaine** : Calcul distribué et à haute performance
Calcul Scientifique (BAP E)
- **Ville** : Saint Martin d'Hères
- **Centre Inria** : [Centre Inria de l'Université Grenoble Alpes](#)
- **Date de prise de fonction souhaitée** : 2025-02-01

- **Durée de contrat :** 6 mois
- **Date limite pour postuler :** 2024-11-21

Contacts

- **Équipe Inria :** [DATAMOVE](#)
- **Recruteur :**
Raffin Bruno / bruno.raffin@inria.fr

A propos d'Inria

Inria est l'institut national de recherche dédié aux sciences et technologies du numérique. Il emploie 2600 personnes. Ses 215 équipes-projets agiles, en général communes avec des partenaires académiques, impliquent plus de 3900 scientifiques pour relever les défis du numérique, souvent à l'interface d'autres disciplines. L'institut fait appel à de nombreux talents dans plus d'une quarantaine de métiers différents. 900 personnels d'appui à la recherche et à l'innovation contribuent à faire émerger et grandir des projets scientifiques ou entrepreneuriaux qui impactent le monde. Inria travaille avec de nombreuses entreprises et a accompagné la création de plus de 200 start-up. L'institut s'efforce ainsi de répondre aux enjeux de la transformation numérique de la science, de la société et de l'économie.

L'essentiel pour réussir

Possibility to pursue with a PhD degree after this internship.

Attention: Les candidatures doivent être déposées en ligne sur le site Inria. Le traitement des candidatures adressées par d'autres canaux n'est pas garanti.

Consignes pour postuler

CV + cover letter

Les candidatures doivent être déposées en ligne sur le site Inria.
Le traitement des candidatures adressées par d'autres canaux n'est pas garanti.

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

Sécurité défense :

Ce poste est susceptible d'être affecté dans une zone à régime restrictif (ZRR), telle que définie dans le décret n°2011-1425 relatif à la protection du potentiel scientifique et technique de la nation (PPST). L'autorisation d'accès à une zone est délivrée par le chef d'établissement, après avis ministériel favorable, tel que défini dans l'arrêté du 03 juillet 2012, relatif à la PPST. Un avis ministériel défavorable pour un poste affecté dans une ZRR aurait pour conséquence l'annulation du recrutement.

Politique de recrutement :

Dans le cadre de sa politique diversité, tous les postes Inria sont accessibles aux personnes en situation de handicap.