



Job vacancy #2023-06619

PhD Position F/M Topology-aware load balancing for ocean simulation on heterogeneous platforms.

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Fonction : PhD Position

About the research centre or Inria department

The Inria center at the University of Bordeaux is one of the nine Inria centers in France and has about twenty 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 SMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute..

Context

CROCO (Coastal and Regional Ocean Community) is an oceanic modeling system (<https://www.croco-ocean.org>). An important objective for CROCO is to resolve very fine scales (especially in the coastal area), and their interactions with larger scales. It includes new capabilities such as a non-hydrostatic solver, ocean-wave-atmosphere coupling, evolving sediment dynamics and marine biogeochemistry, and new high-order numerical schemes for advection and mixing.

Various HPC improvements of the CROCO model itself are currently carried out with respect to a sustainable support of GPUs and different parallel programming models. Indeed, the current trend in high-performance computing architectures is going even more towards increasing heterogeneity. This is omnipresent on the intra-node computation with accelerator cards as well as on the inter-node level with different hardware and communication behaviors.

However, on the application and scheduling side, this trend is often ignored: scheduling of applications, in particular CROCO, still assumes homogeneity across the hardware stack. This leads to a mismatch between applications and the underlying HPC system, resulting in a poor performance in particular in the strong scaling case.

The AIRSEA team in Grenoble is one of the main developers of the CROCO model and the Tadaam team in Bordeaux has the expertise in load-balancing and topology-aware algorithms. Therefore, **this PhD will be carried out mainly in Bordeaux but with strong collaboration with Grenoble**: visits and exchanges will be organized regularly between the two locations.

Assignment

The CROCO ocean model has a very complex workload model including non-homogeneous workload, adaptive mesh refinement with nested grids as well as existing support for hybrid CPUs and GPUs. Optimization attempts without application-driven information are therefore prone to fail. The goal of this PhD is to work on optimizing the execution of the CROCO model on supercomputers by developing and investigating new load-balancing algorithms.

Even if CROCO relies on structured meshes, load imbalance appears between the different computing units due to varying runtime of solvers. Moreover, as the topology of a heterogeneous machine can be extremely complex, the cost of communication can be very high depending on the location of the sender and the receiver. Hence, it is necessary to carefully optimize the mapping of the compute process and the load balance between them to optimize the computation and communication costs of the CROCO model.

Main activities

The Phd Candidate will work on the following workplan:

1. Understanding the CROCO model and the computation/communication graph of the application
2. Work on the state-of-the-art of load-balancing and topology-aware algorithms.
3. Collaborate in the development of a microbenchmark that mimics the behavior of the CORCO model in terms of imbalance and communication on a fixed adaptive mesh.
4. Develop a performance model of the application/microbenchmark that will be used by the algorithmic engine
5. Propose a static load-balancing algorithm for the heterogeneous case (CPU)
6. Evaluate this algorithm on real testcases and real supercomputers.
7. Enhance the solution toward heterogeneous resources (first GPUs, then hybrid) and at runtime.

Skills

- Mandatory:
 - High-performance computing
 - Parallel programming models (MPI, OpenMP)
 - Parallel programming models for heterogeneous computing (GPU/CPU)
 - Performance modeling
 - Strong programming skills
 - Graph:
 - Graph theory
 - Optimization and algorithms
- Optional:
 - Numerics
 - Usage of large-scale super computers
 - Able to cope with operational forecasting codes / Fortran 90

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Possibility of partial teleworking 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

Gross monthly salary :

- 2051€ / month (before taxes) during the first 2 years,
- 2158€ / month (before taxes) during the third year.

General Information

- **Theme/Domain** : Distributed and High Performance Computing Scientific computing (BAP E)
- **Town/city** : Talence
- **Inria Center** : [Centre Inria de l'université de Bordeaux](#)
- **Starting date** : 2023-10-01
- **Duration of contract** : 3 years
- **Deadline to apply** : 2023-11-01

Contacts

- **Inria Team** : [TADAAM](#)
- **PhD Supervisor** :
Jeannot Emmanuel / Emmanuel.Jeannot@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

To succeed the Phd Candidate needs to have the following personal skills:

- Able to autonomously work on topics

- Collaborative skills including online collaborations with remote collaborators
- Excellent communication skills in English
- Strong ability to work on long-term plan

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

Your application must include the following documents:

- CV
- Cover letter
- Master marks and ranking
- Support letter(s)

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.