



2-years Post-Doc position: Data-driven code-coupling at Exascale

Maison de la Simulation offers a 2-year Post-Doc position in the framework of the EoCoE-II European project entitled “Data-driven code-coupling at Exascale”.

- Contact:** Julien Bigot, CEA Researcher, julien.bigot@cea.fr, Tel: 01 69 08 01 75
Location: Maison de la Simulation, bât. 565, CEA Saclay, 91191 Gif-sur-Yvette CEDEX
Employer: CEA, 91191 Gif-sur-Yvette CEDEX
Salary: Based on CEA salary grid, typically in the 35k€-40k€ gross yearly range
Duration: 2 years starting as early as possible from November 1st 2019

Description

With the advent of more and more complex software architecture for high-performance computing, code complexity is becoming an urgent issue. Two trends emerge in this context. On the one hand, simulations implemented as a workflow involving multiple codes would benefit from improved data exchange to overcome the file-system communications performance bottleneck. On the other hand, complex monolithic simulation codes would benefit from improved modularization for maintainability and to make it possible to choose the best suited programming model for each part of the code.

One interesting direction to answer these needs is to design application as efficient loose coupling of independent modules. Traditionally, performance-oriented coupling solutions such as scientific workflows and software components have put an emphasis on control-driven coupling where connections between interacting modules have to be defined by hand. The complexity of these approach means that many code couplings still rely on the file-system for interactions.

Distributed programming models in the field of data analytics have adopted a different approach where most interactions are defined by the data themselves. Typically in the Map-Reduce model, the reducer handling a piece of data is not explicitly specified but implicitly by a key included in the data. The PDI library developed at MdS offers a comparable approach where coupling is achieved by sharing access to a logical data store and reacting to changes in this store. This library is however limited to process-local interactions.

The goal of the proposed Post-Doc position is to evaluate and implement a solution to support data-driven coupling at the scale of a complete Exascale machine. In order to achieve this goal, the candidate will have to take into account problems that arise at this scale, including but not limited to the following. Data distributions and re-distributions between modules will have to be handled. Achieving good performance while taking into account data transfer times will require asynchronous solutions. Unique values identification in the presence of asynchronism will require additional information, for example to handle values from different time-steps.

Skills

The successful candidate will master the following skills and knowledge:

- strong experience in software engineering,
- knowledge of parallel computing (including the MPI library)
- proficiency in C++11.

In addition, the following skills and knowledge will be considered a plus:

- knowledge of the C, Python & FORTRAN languages,
- knowledge of HPDA programming models including the Map-Reduce approach,
- understanding of the PGAS programming model.

Context

Maison de la Simulation (MdlS – <http://www.maisondelasimulation.fr/>) is a joint laboratory located on the Saclay plateau between: the French Alternative Energies and Atomic Energy Commission (CEA), the French National Center for Scientific Research (CNRS), the University of Versailles Saint-Quentin (UVSQ) and the University of Paris-Sud. The laboratory groups activities around high performance computing (HPC): research in computer science and applied mathematics, engineering and development for high-performance simulation applications. Of specific interest at MdlS are software engineering aspects of HPC, especially regarding separation of concern.

The **Energy oriented Center of Excellence** (EoCoE – <https://eocoe2.eu/>) is a European project lead by MdlS at the cross-road of the energy and digital revolutions. EoCoE develops and applies cutting-edge computational methods in its mission to accelerate the transition to the production, storage and management of clean, decarbonized energy. EoCoE is anchored in the HPC community and targets research institutes, key commercial players and SMEs who develop and enable energy-relevant numerical models to be run on exascale supercomputers, demonstrating their benefits for low-carbon energy technology. EoCoE drives its efforts into 5 scientific Exascale challenges in the low-carbon sectors of energy: Meteorology, Materials, Water, Wind and Fusion.

The **Portable Data Interface** (PDI – <https://pdi.julien-bigot.fr/>) is a library designed and developed at Maison de la Simulation for process-local loose coupling in high-performance simulation codes. It offers a reference system similar to Python or C++ `shared_ptr` with locking so as to ensure coherent access by coupled modules. It provides a global namespace (the data store) to share references and implements the *Observer* pattern to enable modules to react to data availability and modifications. It implements a metadata system that can be used to specify a dynamic type for references based on the value of other data (eg. array size based on the value of a shared integer). Codes using PDI declarative API expose the buffers in which they store data and notify it of significant steps in the simulation. Third-party libraries such as HDF5, SIONlib or FTI are wrapped in PDI plugins. A YAML file is used to to interleave application code with plugins use and additional code without having to modify the original application.

Work-plan and environment

Along the 2 years, the candidate will be expected to propose, implement and validate a new coupling strategy. Developments should take place both in dedicated experimental prototypes and in the existing PDI library where it makes sense. Validation will specifically focus on the example of the extreme-scale plasma simulation code Gysela that currently embeds in a monolithic Fortran application both data pre-processing, the simulation per-se, in-line and in-situ post-processing and a checkpoint-restart mechanism and additionally relies on Python post-processing scripts for data analysis. The candidate is expected to publish its results in international conferences and journals like Supercomputing, IPDPS or JPDC for example.

In order to achieve these goals, the candidate will join the PDI team including engineers working on the library and its integration in existing codes as well as a PhD student working on smart workflow for HPC. She or he will be given access to computing resources including the MdlS dedicated cluster as well as French national and European supercomputers. It will also be possible to use some of the multiple large-scale production codes ported to PDI in the EoCoE project for experiments, including Gysela, but also Parflow, Eurad-IM, Warf or Esias for example.