Title
Implementation of the parallel sparse linear algebra solver **PaStiX** developed by the **HiePACS** team in large application code and scalability study.

Hosting team
HiePACS – Inria Bordeaux Sud-Ouest

Research and Development framework
The work will be conducted within the **HiePACS** team which main expertise is the design and development of parallel numerical scientific libraries targeting large scale computing. This activity will be developed in the framework of a European project in close collaboration with the other partners. The scalability studies will be performed on large computing platforms both in France and in Europe.

Mission
For many years, the team has developed expertise in the field of high performance computing and has designed many parallel software libraries, such as **PaStiX** (http://pastix.gforge.inria.fr), for the solution of large sparse linear systems of equations. Those solvers have already been integrated in some large simulation codes and enable the solution of large 3D problems with a few million unknowns.

The main tasks in that project are applications driven:

- At every step of the project, the developments done in the **PaStiX** library will be made available to the community through the public Git repository of the library. All the work developed in the context of the project will be integrated in the unitary tests and the continuous integration system that has been put in place on the Inria gitlab repository to improve the quality of the software produced and ease its integration in applications. Additionally, a public API will be developed to make available the direct interaction with the low-rank kernels recently introduced in the **PaStiX** solver. This development will be critical to have a good interaction with applications that may use the Schur complement directly in its low-rank form.
- **PaStiX** is already used with the full rank version in applications from CEA Cesta and Cadarache. For instance, **Jorek** and **Tokam3x** are hybrid (Fortran90, MPI + OpenMP) parallel code developed for fusion simulations in the ITER project. In this task, we will work on the integration of the new **PaStiX** interface that has been developed for the BLR (block low-rank) version of the solver, and that will slightly evolve to follow the needs of the hierarchical matrix format. Furthermore, the **Tokam3x** code has recently started investigating optimizations for many-core architectures. In order to produce challenging simulations of edge plasma turbulence in X-point geometry, the linear solvers need to handle very large sparse matrices within the limited amount of memory available on such architectures. Compression techniques will help to reduce memory consumption and time to solution to study larger test cases per node.
- A parallel distributed memory implementation has already been developed in previous releases of **PaStiX** (5.*). Most of the users of our solver reclaim an MPI implementation of the new version of **PaStiX** (6.*). Since the solver can handle two runtime systems (**StarPU** and **PaRSEC**) a first attempt could consist in managing the communications through those runtime systems. But for a reference implementation we also intent to redevelop the native MPI/thread version of the solver.

Duration
12 months with possible renewal.

Required skill
Background in parallel computing (algorithms and platforms), parallel programming (MPI, OpenMP, CUDA), Knowledge in linear algebra.

Person to contact

Pierre RAMET and Mathieu FAVERGE – pierre.ramet@inria.fr / mathieu.faverge@inria.fr