

# **Proposal for the Use of High Level Support Team resources**

## **Abstract**

*(150 – 300 words)*

Edge turbulent transport is one of the key issues of magnetized fusion research. It determines both the performances of the tokamak, via its role in the building of the H-mode transport barrier, and the life expectancy of plasma facing components, via the particle and heat fluxes reaching the solid materials of the reactor wall. Benefiting from the steady increase of the available computing power, the current international effort is progressively shifting towards developing 3D fluid turbulence codes able to describe self-consistently turbulent transport in realistic geometry.

The TOKAM3X is currently being developed and used as part of this effort. Hybrid MPI+OpenMP parallelization is used to take advantage of the characteristics of large nodes machines currently available. The code has recently reached a stage of its development where a first stable version exists and has been applied to simple physical cases. More complex cases in more realistic geometry will require the use of finer meshes at the cost of a strong increase in computing time. Optimizing the performance of the code will greatly improve the accessibility of these complex cases.

The aim of the current project is to perform a complete profiling of the current version of the code. Both sequential and parallel (MPI and OpenMP) profiling are needed. The expected output is the identification of bottlenecks and most promising sources of optimization. This will be a key guideline in developing the code.

|  |   |
|--|---|
| <b>Project Title</b>                           | <i>Sequential and parallel profiling of the TOKAM3X 3D edge fluid turbulence code</i> |
| <b>Project Acronym</b><br>(up to 8 characters) | <i>TOKAPROF</i>   |

**Project coordinator:**

|                             |                             |
|-----------------------------|-----------------------------|
| <i>Name of Coordinator:</i> | Patrick Tamain              |
| <i>Institution:</i>         | CEA, France                 |
| <i>Street Address:</i>      | CEA Cadarache, IRFM bat 513 |
| <i>City:</i>                | Saint Paul lez Durance      |
| <i>Country:</i>             | FRANCE                      |
| <i>Email:</i>               | patrick.tamain@cea.fr       |
| <i>Phone:</i>               | 0033 44225 2616             |

**Principal Investigator(s) [other than coordinator]:**

|                             |                             |
|-----------------------------|-----------------------------|
| <i>Name of Coordinator:</i> | Guillaume Latu              |
| <i>Institution:</i>         | CEA, France                 |
| <i>Street Address:</i>      | CEA Cadarache, IRFM bat 513 |
| <i>City:</i>                | Saint Paul lez Durance      |
| <i>Country:</i>             | FRANCE                      |
| <i>Email:</i>               | guillaume.latu@cea.fr       |
| <i>Phone:</i>               | 0033 44225 6357             |

|                                |   |
|--------------------------------|---|
| <i>Institution:</i>            | M2P2 laboratory, Aix-Marseille University |
| <i>Names of Investigators:</i> | Eric Serre                                |
| <i>Street Address:</i>         | 38, Rue Frédéric Joliot-Curie             |
| <i>City:</i>                   | Marseille                                 |
| <i>Country:</i>                | FRANCE                                    |
| <i>Email:</i>                  | eric.serre@L3m.univ-mrs.fr                |
| <i>Phone:</i>                  | 0033 49111 8535                           |

*Please duplicate the table above if Principal Investigators from more than one Research Unit*

**Requirements for the present largest run of the code**

|  |  |
|--|--|
| <i>Total amount of CPU hours</i>                         | 50 000 core.h                              |
| <i>Architecture(s) where application is already used</i> | Intel Sandy Bridge, Intel X5675 (Westmere) |
| <i>Number of CPUs</i>                                    | 128-256 cores                              |

|   |  |
|---|--|
| <i>Memory requirements</i>                            | 500MB - 1GB per core   |
| <i>Storage requirements</i>                           | 25GB   |
| <i>Pure MPI or mixed communication (OpenMP+MPI)</i>   | Mixed OpenMP + MPI   |
| <i>Own code / 3rd party code</i>                      | Own code   |
| <i>Code publicly available (yes/no)?</i>              | No   |
| <i>Library requirements</i>                           | MPI, OpenMP, HDF5, PETSC, MUMPS, PASTIX                        |
| <i>Special requirements</i>                           |  |
| <i>Site name(s) where application is already used</i> | Mesocentre Aix-Marseille University, CINES/France, CCRT/France |
| <i>Expected usage of the IFERC computer (yes/no)?</i> | Next year  |

### Technical Improvement or adaptation work done so far

- Do you apply in parallel for similar support from other institutions?  
No.
- Has your code/project already received support (especially as part of a previous HLST call) related to improvement of its computational capabilities?  
No.

### Request for work

- Indicate nature (type) of HLST support being requested
  - Sequential and parallel profiling (MPI alone, OpenMP alone, hybrid OpenMP+MPI) of the TOKAM3X code
  - Identification of main bottlenecks and main potential sources of computing time optimization within memory limits of large nodes parallel machines, typically 24GB / node
  - Implementation in the code of some of these improvements
- Indicative level of support (in ppm<sup>1</sup>)  
12ppm

### Involvement of the project proponents

<sup>1</sup> Note that 1ppy=12ppm

Indicate the effort (in ppm<sup>1</sup>) of the projects proponents to be given (in parallel to the HLST work) to the execution of the project

1ppm

|                                |
|--------------------------------|
| <b><i>Potential Impact</i></b> |
|--------------------------------|

Indicate the estimated benefits that the HLST support activity will have on the software and physics modelling capabilities

TOKAM3X is a relatively new code. The effort has until now mainly focused on developing a first stable version and applying it to relatively simple physical studies. Although some effort was put in the parallelization from the beginning, no optimization work has been undertaken yet so that a large margin of improvement probably exists. Now that a stable version exists, this profiling and optimization effort becomes a mandatory step towards addressing more complex cases requiring larger meshes and computing time.

## Detailed Project Description (max. 1-2 pages)

The TOKAM3X code has now reached a production stage and has been used to study turbulent particle transport in simple limiter plasmas. Although the model treated by the code is expected to progressively evolve as more physics is added in the model, its basics structure can be considered as converged. The spatial discretization is based on conservative finite differences. A first order operator splitting is used to treat separately the different operators of the system. Three main types of operators can be distinguished:

1. Parallel and perpendicular advection terms: they describe a dynamics occurring on the time scale of interest and can hence be treated explicitly with a forward Euler scheme. Shock-capturing algorithms (WENO and Roe-scheme) are used to insure both high order discretization and numerical stability
2. Perpendicular and parallel diffusion terms: due to the large values of the diffusion coefficients expected to be used in some situations, these are treated with an implicit backward Euler scheme. The corresponding sparse 2D linear systems are LU-decomposed and solved using the MUMPS or PASTIX solvers.
3. The so-called “vorticity” operator, related to the current balance equation: it is an extremely anisotropic 3D laplacian operator with fast dynamics in the parallel direction. It is solved implicitly using the MUMPS or PASTIX sparse solvers. PASTIX is usually preferred due to the fact that it is multi-threaded (MPI+OpenMP versus only MPI for MUMPS). This operator is usually the most time-consuming of the code.

The HLS team will be provided with the latest stable version of the TOKAM3X code as well as a set of test cases (typically 3) to be used for the analysis of the performances of the code in different template geometries (leading to different domain decompositions and load distributions between processes). In each case, a thorough profiling of the code and its subroutines will be performed in order to:

1. Determine the current scaling of performances with mesh size and number of CPUs (strong and weak scaling) for each sub-part of the code
2. Identify bottlenecks and main potential sources of improvements

The profiling will be carried out both for the sequential performances (optimization of the algorithm) and the parallel performances (OpenMP loops and directives as well as MPI communication scheme).

Concerning point 2, optimization will be sought through both amendments of the code itself, but also the determination of good compilation flags. Special

focus will be given to putting in place a good vectorization (AVX) strategy. The coupling with the PASTIX linear sparse solver should be looked at as it is one of the most time consuming parts of the algorithm.

Some of the recommended modifications will be implemented in the code and tested. They will also be used to define a set of guide lines for future developments of the code.