

# Designing and running turbulence transport simulations using a distributed multiscale computing approach

O. Hoenen<sup>1</sup>, L. Fazendeiro<sup>2</sup>, B D. Scott<sup>1</sup>, J. Borgdorff<sup>3</sup>, A G. Hoekstra<sup>3</sup>, P. Strand<sup>2</sup>, D P. Coster<sup>1</sup>

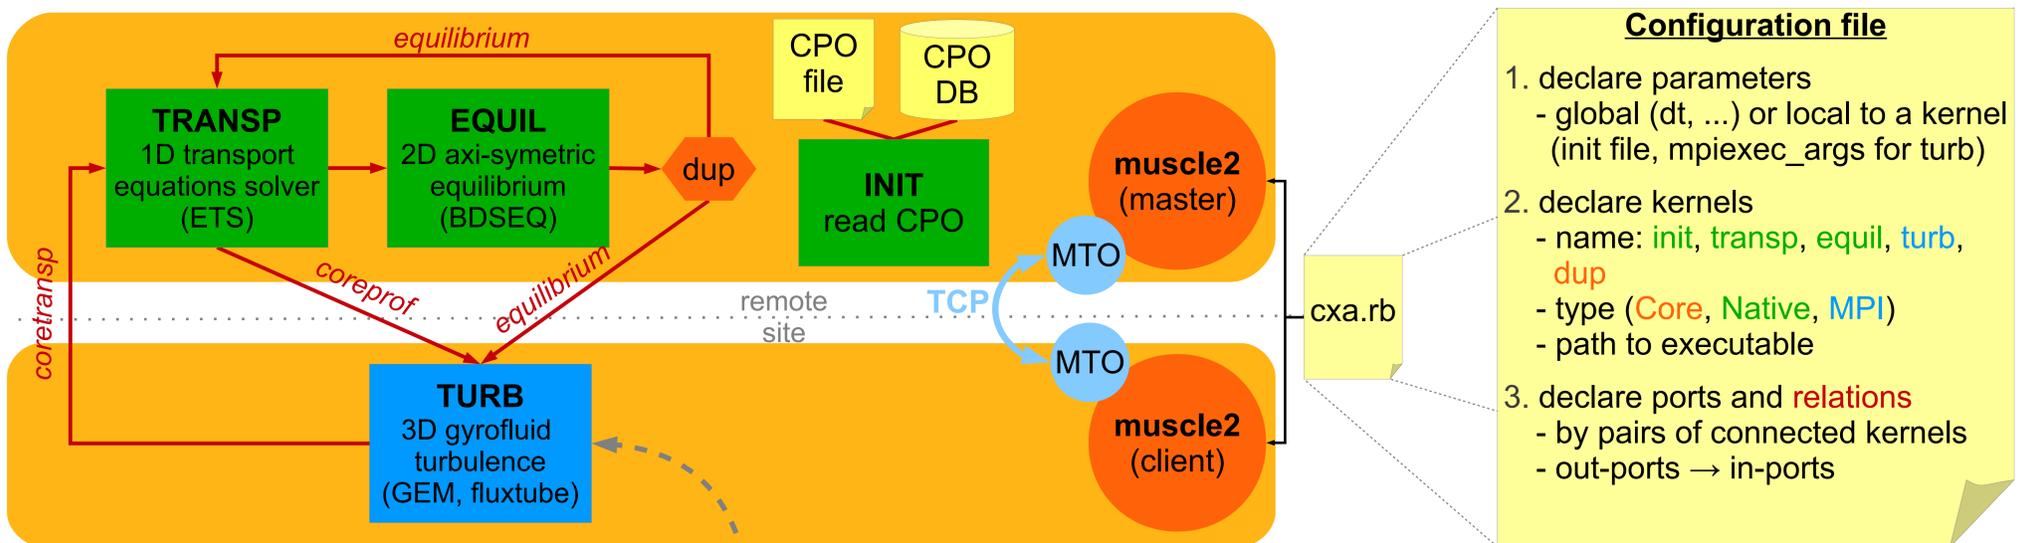
<sup>1</sup> Max-Planck-Institut für Plasmaphysik, Boltzmannstr. 2, D-85748, Garching, Germany <sup>2</sup> Chalmers University of Technology, SE-412 96 Göteborg, Sweden <sup>3</sup> Institute of Informatics, University of Amsterdam, 1090 GH, The Netherlands

## Introduction

- Multiscale simulation involving slow transport and fast turbulent scales is a challenging computational problem.
- Modelling a multiscale problem as a **set of coupled single scale submodels** prevents complexity of monolithic codes.
- Scales can be spatio-temporal domains or multi-physics.
- Such approach requires using some **coupling framework**.
- Interface agreement, data exchange:
  - submodels using **ITM's generic datastructure**, CPO [1]
- Need for **distributed simulation** capabilities:
  - execute a submodel on a specific hardware (accelerators or bigger HPC systems)
  - access to a local database (simulation or experiment)
- Target workflow: Transport-Equilibrium-Turbulence

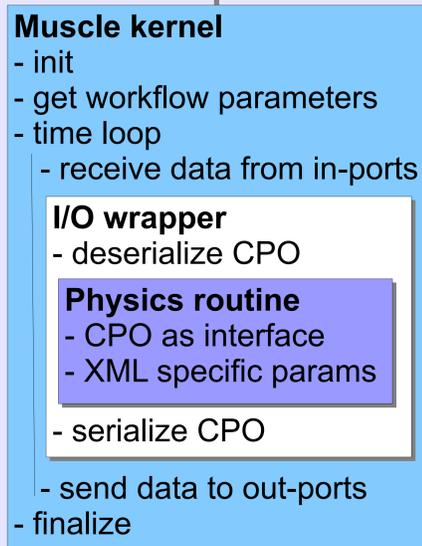
## MAPPER

- Multiscale **APPL**ications on European e-InfRastructure <http://www.mapper-project.eu>
- Provides a formal framework, tools, software and services for building and running **distributed multiscale** applications [2]
- Assists scientists developing multiscale applications from scratch or from **legacy codes**
- Two types of application:
  - loosely-coupled: acyclic, data exchange with file
  - **tightly-coupled**: cyclic, data exchange with coupling library
- Software stack on three levels:
  - high-level tools: web-based GUI, design and execution
  - middleware: distributed resource manager (grid and HPC)
  - **coupling library (MUSCLE2)**: schedules work, transfers data



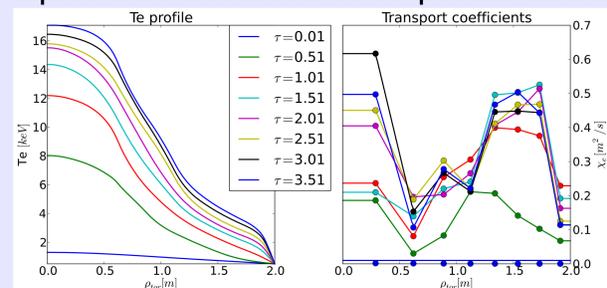
## MUSCLE2

- Java-based library
- **Simple API** in C/C++ and F90:
  - MUSCLE\_init/finalize
  - MUSCLE\_get\_property
  - MUSCLE\_send/recv
- Communication: **blocking**, interoperable types, **byte streams**
- Coupled simulation requires:
  - muscle2 bootstrap program
  - submodels implementation (kernels)
  - configuration file
- **Transparent remote transfer** (clients just need master IP:port)
- **Firewall traversal** through Muscle Transfer Overlay (MTO)

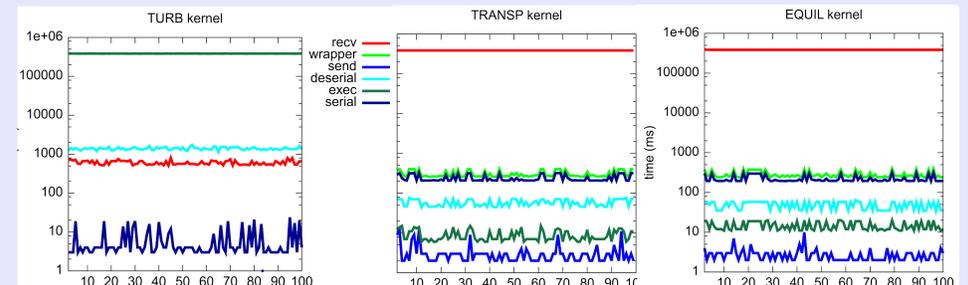


## Results

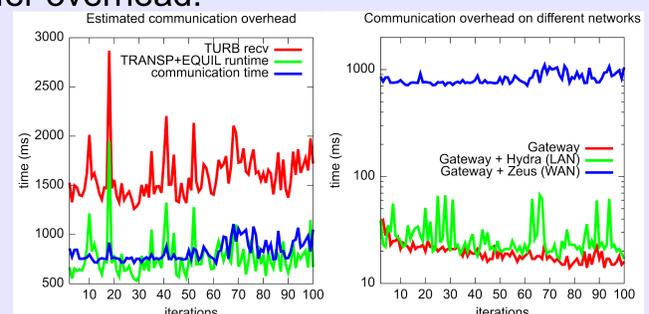
Proof of principle: time evolution of Te profile and transport coefficients



Muscle's kernel serialization/deserialization overhead:



Data transfer overhead:



## Conclusion

- Simple to use, non-intrusive, adaptable through a config file
- MPI: only main process interacts with muscle ⇒ broadcast
- Serialization overhead is not negligible for fast kernels

### References:

- [1] F. Imbeaux, Computer Physics Communications **181**, 987 (2010)  
 [2] J. Borgdorff, et al., Procedia Computer Science **9**, 596 (ICCS 2012)