Fusion Engineering and Design 87 (2012) 2057–2062

Contents lists available at SciVerse ScienceDirect

Fusion Engineering and Design

journal homepage: www.elsevier.com/locate/fusengdes

Easy use of high performance computers for fusion simulations

Y. Frauel[a], D. Coster[b], B. Guillerminet[a,*], F. Imbeaux[a], A. Jackson[c], C. Konz[b], M. Owsiak[d], M. Plociennik[d], B. Scott[b], P. Strand[e], European Task Force on Integrated Tokamak Modelling Activity[2]

[a] *CEA, IRFM, F-13108 Saint-Paul-lez-Durance, France*
[b] *Max Planck Institut Fur Plasma Physik, Garching, Germany*
[c] *University of Edinburgh (EPCC), Edingburgh, United Kingdom*
[d] *Poznan Supercomputing and Networking Centre (PSNC), Poznan, Poland*
[e] *Chalmers University of Technology, Sweden*

ARTICLE INFO

Article history:
Available online 24 April 2012

Keywords:
Simulation
Fusion
ITER
GRID
HPC
ITM
EUFORIA

ABSTRACT

Fusion Modelling and Simulation are very challenging and the high performance computing issues are addressed here. Based on the framework developed by the European Integrated Tokamak Modelling project and on the EUFORIA infrastructure, a tool solving nicely these difficulties has been developed for the end users and applied to several fusion simulation cases. The first part recalls the issues with GRID and high performance computing, while the second part presents the solutions and the tool for developing easily a GRID/HPC actor. The last part reports the use of this tool in MHD equilibrium and plasma edge simulations.

1. Introduction

Simulating a whole fusion machine is a huge undertaking. The problem has traditionally been broken up into isolated simulations of elementary physical phenomena, the simulation of each of these from first principles requiring extremely demanding resources. Today, thanks to the availability of more powerful computing resources, the integration of ab initio plasma physics calculations to other physics and technology simulation modules becomes possible, in view of getting a more complete picture of the reality. Such an integration effort is conducted inside Europe's Integrated Tokamak Modelling Task Force (ITM-TF), whose aim is to gather existing simulations codes, as well as to produce new codes and a complete software infrastructure to integrate these codes together.

Several of the tools developed within the recently completed European project EUFORIA for the GRID and high performance computing (HPC) infrastructures have been adapted and extended by the ITM-TF, with the aim of easing/automating the execution of simulation codes on supercomputing resources. In particular, a tool has been developed that allows the automatic conversion of a FORTRAN or C++ code into a component that can be integrated in a simulation workflow and actually executed on an HPC or the GRID. In this paper, we describe this tool as well as the first simulations constructed with it.

The first part of this paper recalls the ITM-TF and EUFORIA projects focusing on the tools developed for the simulations on super-computers. Issues and the tool developed for easy use of the GRID and HPC infrastructures are described in the main part. We then demonstrate the first physic codes which have been implemented with the help of the latest HPC tools. We present the first results obtained within the ITM-TF framework running on mixed infrastructures, clusters and HPC.

2. Description of the context

2.1. The ITM-TF and EUFORIA projects

The European Fusion Development Agreement (EFDA) organization has recognized the need to coordinate efforts at the international level in order to foster the development of a complete infrastructure for fusion modelling. It has therefore funded

* Corresponding author.
E-mail addresses: yann.frauel@cea.fr (Y. Frauel), david.coster@ipp.mpg.de (D. Coster), bernard.guillerminet@cea.fr (B. Guillerminet), frederic.imbeaux@cea.fr (F. Imbeaux), adrianj@epcc.ed.ac.uk (A. Jackson), cjk@ipp.mpg.de (C. Konz), marcinp@man.poznan.pl (M. Owsiak), michal.owsiak@man.poznan.pl (M. Plociennik), bds@ipp.mpg.de (B. Scott), par.strand@chalmers.se (P. Strand).

[1] http://www.euforia-project.eu/.
[2] http://portal.efda-itm.eu/portal/.

0920-3796/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.fusengdes.2012.04.015

since 2005 the Integrated Tokamak Modelling Task Force (ITM-TF) [1] gathering more than 240 individuals from 25 different European institutions. In charge of "coordinating the development of a validated suite of simulation tools for ITER and DEMO plasmas", it has developed a framework for data description and management and for workflow orchestration. A common development platform ("the Gateway") has been made available to all participants [2]. To allow cross validation between different fusion devices, a standardized descriptions of the data and a device independent approach to data transport (resp. Consistent Physical Objects or CPO and Universal Access Layer or UAL [3]) have been developed within the ITM framework. The Open Source Kepler [4] tool has been chosen for the workflow orchestration. A workflow component is called an actor in Kepler and it could be itself composed of actors and workflows. Supplemented with a number of automated tools to integrate physics component into it, Kepler provides a dynamic, flexible and extensible modelling environment for fusion modelling. An advantage of this approach is that the complexity of mixed language programming and even detailed data management within the integrated modelling application is hidden from the end user.

Modelling activities will be an essential part of ITER plasma analysis and operation, and developing a predictive capability is a very challenging task requiring easy access to high performance computing infrastructures. Hence data access and code coupling technologies are required to be available for a heterogeneous, possibly distributed, environment. The developments in this area were pursued in a separate project – EUFORIA (EU Fusion for ITER Applications) [5] which started in 2008 and ended in 2010 from 14 different European institutes. EUFORIA enhanced the modelling capabilities for ITER-sized plasmas through the extension of the ITM-TF toolset to access GRID and HPC resources and by adaptation, optimization and integration of a set of critical applications for edge and core transport modelling. EUFORIA integrated codes and applications running on a set of heterogeneous platforms into a single coupled framework through the Kepler workflow engine. The extension of the workflow orchestration system and middleware to transparently incorporate GRID and HPC is based on a Roaming Access Server (RAS), in charge of accessing the various middleware (i2g, glite, globus, Unicore). The workflow engine is only connected to this stable layer and the outcome is a significant improvement of the architecture which simplifies the connection between the workflow engine based on Kepler and the various GRID/HPC infrastructures.

The tool presented in this paper, is based on the ITM framework and EUFORIA infrastructure [6,7].

2.2. Issues with GRID and high performance computing

Running jobs on GRID or HPC infrastructures is usual but the issues in the fusion simulations come from the simultaneous uses of clusters, GRID and HPC in a single workflow. Tasks and jobs are distributed among these various infrastructures sharing large datasets. Dynamic coupling and integration of codes running on a set of heterogeneous platforms into a single coupled framework through a workflow engine is very challenging. The main issues are shown below.

2.2.1. Authorization/security

The various infrastructures use different authentication mechanisms (login, certificates, registry). The integrated modelling must be able to schedule jobs on the various infrastructures using the same mechanism. EUFORIA solved this issue by using the GRID certificate and implementing a registry service on HPC infrastructures.

2.2.2. Execution

Several middleware initiatives have contributed to distributed computing infrastructures like Globus-based frameworks (gLite, i2g, etc.) and UNICORE but there is no standard able to handle all the GRID & HPC infrastructures. Even the languages for job submission are different (RSL, JDL, JSDL, etc.). The developed tool/infrastructure must be able to handle:

- Parallel jobs mainly based on MPI. As explained below, MPI usage has been restricted to HPC infrastructures.
- Stateful jobs. In a scientific workflow, jobs are data driven and it is time consuming to restart from scratch at each time step. The batch queue mechanism is not appropriate. Remote suspend/resume commands are also suitable in the debugging phase. Solving this issue on all infrastructures is explained in the next section.
- Infrastructure errors. Due to software/hardware maintenances and various distributed architecture, we have to handle execution errors in a transparent way for the user. EUFORIA has implemented a set of Kepler actors (Serpens suite) in charge of this error handling [8].
- New middleware: implementation of new middleware or upgrading the existing one is handled by the RAS server.
- Long duration simulations. HPC systems commonly have time limits on the length of computational jobs.

2.2.3. Data access

For security reasons, external data access is generally very restrictive and a very limited number of tools are allowed to move data from outside to inside, usually based on GridFTP file transfer or other secure data transfer mechanisms (such as copy over SSH). The ITM data access is based on a standardized interface (Universal Access Layer, UAL) and it has been implemented on GRID and HPC infrastructures using the MDSplus transport layer [9].

3. A new tool to automate the use of supercomputing resources

In order to make it easier to run codes on either a GRID or an HPC, we have developed a software tool called HPC2K (which stands for "HPC to Kepler"). The purpose of this tools is to convert a physics code written either in Fortran or in C++ to a Kepler actor that automatically launches a job on a GRID or HPC. The HPC2K tool is written in Java and uses C, C++ or Fortran template for the code generation. HPC2K generates an actor (running locally) and a job running on GRID or HPC. The communication between the actor and the job is based on several Java libraries (Unicore, glite, i2g, etc.) for the control (execution, status, etc.) and on the UAL for the data access [3,7].

The code must initially be compiled as a library with a single main function that corresponds to the execution of the code. Then HPC2K is run. The tool presents a simple graphical user interface (Fig. 1), in which the user just has to fill in some fields, such as the number and types of the code arguments, the location of the code library, the compiler to use, the target system (GRID or HPC), etc. Once all the parameters are set, the creation of the Kepler actor is launched through a simple button. The tool then generates a series of files:

- A wrapper program – written in either Fortran or C++, depending on the language of the code – that deals with the remote data transfers in input and output between the database located on the ITM's Gateway computer and the code running on a remote computer.