

# XXL Simulation for XXI<sup>st</sup> Century Power Systems Operation

**IESP Workshop, 6-8 of April 2009, Santa Fe, NM-USA**

**JY Berthou and JF Hamelin and Etienne de Rocquigny**

EDF R&D, 1, Av du Général de Gaulle, BP 408, 92141 Clamart Cedex, France

E-mail: jy.berthou@edf.fr, jean-francois.hamelin@edf.fr, etienne.derocquigny@edf.fr

## **Abstract**

As an industrial user with very high stakes in the operation and maintenance of complex systems like nuclear power plants, EDF has been engaged into simulation for many years. We feel that Exaflops software should not only be thought as a way of tackling daunting research problems but should also take into account the sometimes equally daunting requirements that stem from an industrial usage perspective. We do feel that, whatever the hard changes that will probably have to be made on various software aspects, we should not lose sight that continuity paths have also to be found in order to make those big changes acceptable and profitable to many. We identified in this paper what we think are priority research themes that could benefit of an international collaboration

**Keywords :** HPC simulation, parallel programming, uncertainty and risk quantification, simulation platform, numerical solvers

## **1 Introduction**

As an industrial user with very high stakes in the operation and maintenance of complex systems like nuclear power plants, EDF has been engaged into simulation for many years. We have decided to design our own codes in order to capitalize precious knowledge on our fleet of nuclear reactors, and shorten the time to put this knowledge at work for the many engineering challenges that we have to meet. Software in the millions of lines have been written and explain why we feel very much concerned by the future requirements for Exaflops machines. We have already established the value of running our codes on 100 Tflops / 30 000 cores computers which yield a much better understanding of operating margins and in turn allow for a better optimisation of our power plants, increased safety and performance, lower environmental impact and costs and extended lifetime of assets. We have also recognized that some of our key industrial processes like waterflow within our nuclear cores or production optimisation under uncertain future are still out of reach of Petaflop grade technology and will require major changes in the way we write, validate, run and use simulation codes.

We therefore feel that Exaflops software should not only be thought as a way of tackling daunting research problems but should also take into account the sometimes equally daunting requirements that stem from an industrial usage perspective: this includes both the capacity to model very complex, possibly coupled phenomena over extended spatial and time scales, mixed with capacities like uncertainty quantification or data assimilation that are key to industrial acceptance. Our contribution to this IESP workshop is not that of software specialists but of fairly active users already engaged in the evolution of existing software for Petaflop/100 k cores machines. We will contribute the issues and problems that we are already facing at this first level, and that must find solutions for the future. We do feel that, whatever the hard changes that will probably have to be made on various software aspects, the group should not lose sight that continuity paths have also to be found in order to make those big changes acceptable and profitable to many. The context of simulation at EDF is detailed in [Hame].

## **2 Major software barriers as seen by an industrial user of HPC and propositions for an international collaboration**

One of the major difficulty will be to manage massively parallel systems, composed of approximately millions of heterogeneous cores that will appear at the end of this decade. The challenge is particularly severe for multi-physics, multi-scale simulation platforms that will have to combine massively parallel software components developed independently from each others. Another difficult issue is to deal with legacy codes, which are constantly evolving and have to stay in the forefront of their disciplines. This will require new compilers, libraries, middleware, programming environments, languages, as well as new numerical methods, code architectures, mesh generation tool, visualization tool:

We identified below what we think are priority research themes that could benefit of an international collaboration.

## 2.1 Programming massively parallel computers

### Possible joint efforts:

- **Languages/compilers/performance** analysis tools for achieving **mono-processor** high performance, specially with accelerators (Larrabe, GPU, Cell, ...)  
Goal : achieve more than 30% of the peak performance
- Efficient, “easy to use”, portable and fault tolerant implementation of **Libraries/Languages/compilers** for mixed parallelism : MPI/OpenMP/”cuda like” language  
Goal: one million cores (heterogeneous, hierarchical and massively parallel)
- **Algorithm/solvers** and **data structures** adapted to heterogeneous/hybrid, multilevel and hierarchical massively parallel machines.  
Example: dealing with non-structured irregular meshes for CFD computation on GPU  
Goals:
  - No global communication involving the complete system(avoiding MPI\_ALL-REDUCE, MPI\_BARRIER,... on 1 million of threads)
  - exhibiting different type of parallelism (MPP, SIMD, ...)
  - enabling fault tolerance techniques implementation
  - enabling efficient IO (data restructuring?)

## 2.2 A single generic interface for High Performance Solvers

**Possible joint efforts.** Defining and developing a single generic interface for High Performance Solvers

Computational scientists have developed over the past 20 years numerous[Dong] scientific libraries and solvers (direct, iterative and eigenvalue), ScaLAPACK, PETSc, HyPre, TRILINOS to cite some of them, which all have their own interface. This multiplicity of interfaces makes difficult and costly their integration and maintenance in end-user Scientific Application. It also makes tricky for a given community to test them and find the most appropriate for a given purpose. Both solver and code developers would greatly benefit of a **single generic interface for High Performance Solvers**. Moreover, coming with interfaces to freely available libraries, the sources of the codes are available. This is of great importance for industrial software stability in time. In order to be compatible with the external libraries, the necessary periodic efforts are only done once by the Interface’s development team and not many times by each client software using, for example, PETSc or HyPre separately.

A similar project called Numerical Platon[NP] is developed by the French Atomic Energy Commission (CEA) . It provides an interface to a set of parallel linear equation solvers for high-performance computers that may be used in industrial software written in various programming languages (C, C++, FORTRAN, Python...). This tool was developed as part of considerable efforts by the CEA Nuclear Energy Division in the past years to promote massively parallel software and on-shelf parallel tools (public and in-house solvers, essentially PETSc, SuperLU and HyPre) to help develop new generation simulation codes.

Moreover, at EDF R&D, collaborations are currently underway to improve the direct solvers MUMPS[Mump] and PaStiX[Past] (Out-of-core, parallelization of the analyse step, null space basis computing) and their hybrid overlays (A2S2 and HIPS). These sparse parallel solvers are natural candidates to join such a product.

### **2.3 Stochastic HPC computing for uncertainty and risk quantification**

Numerical modeling of increasing complexity are developing in order to better characterize the underlying factors : multi-physics, multi-scale or complex portfolios all imply increasing computing power. Probabilistic quantification of the associated risks and uncertainties amounts to an additional technological challenge as one needs to multiply at a large scale these already-costly unit simulations in a framework that becomes stochastic. This also alters the way the computer power is invested in the sense that massive distribution becomes necessary; to best value decision-support computing power, one needs to re-work the compromise between the sophistication of best-estimate models and meshes and the stochastic exploration. On this rapidly evolving domain, two kinds of challenges may be highlighted: those related to the development of stochastic methods, and those related the associated computer science implications.

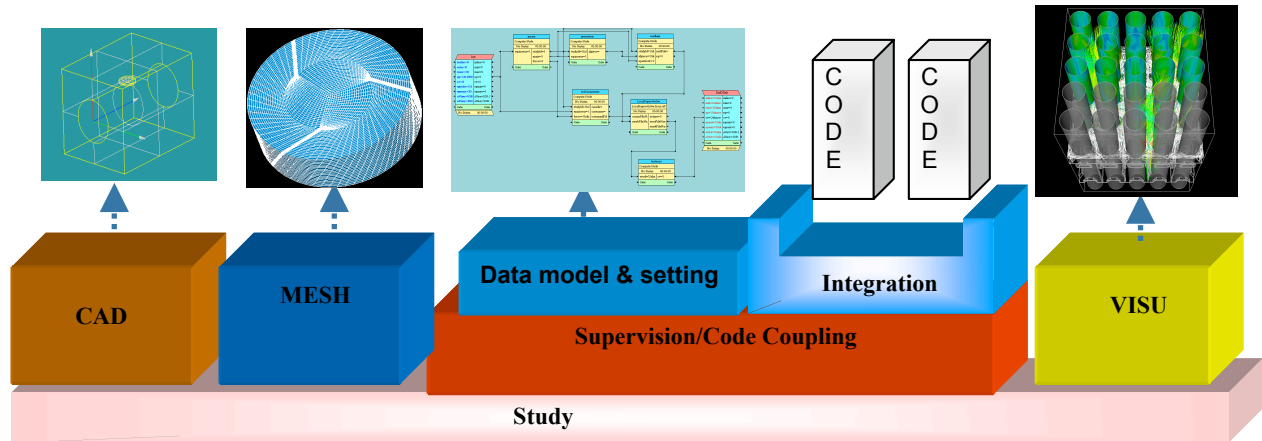
Probabilistic quantification of the risks and uncertainties affecting a best-estimate model has generated a whole domain of applied science, linking probabilistic, numerical analysis as well as physics and decision-theory [Rocq]. Beyond the traditional Monte-Carlo sampling whose history is closely linked to that of computing itself with Von Neumann's ENIAC pioneering applications, a number of uncertainty propagation and probabilistic simulation algorithms have been developed, such as accelerated sampling (importance sampling, particulate methods etc.), reliability techniques (FORM-SORM etc.), stochastic developments (e.g. chaos polynomials) and response surface techniques, yet still wanting for further development particularly regarding the challenges of low probability estimates for irregular response or high input dimension for sensitivity analysis/importance ranking or high-volatility time series.

Beyond uncertainty propagation or risk computation, even tougher challenges come with the need for inverse probabilistic techniques as the observable data to calibrate model variability generally comes on parameters different the model inputs, so that the identification of the extent of uncertainty affecting its input parameters requires the use of inverse techniques coupled with stochastic simulation. Closely related is the need for a general coupling between stochastic optimization and simulation in order to strike robust design or operational management strategies, with challenging mathematical implications that are only partially solved under existing Expectation-Maximization or stochastic dynamic programming algorithms (typically limited to close to Gaussian/linear behavior). Bayesian settings are also bound to develop to better incorporate expert knowledge in a solid decision-theory foundation.

Beyond the development of the methods itself there are key implications on the way HPC is structured and used: challenges involves striking an advanced compromise between parallel & distributed stochastic computing. While standard Monte-Carlo sampling leads to straightforward massive distribution as the runs are all fully independent, the other kinds of stochastic computing algorithms do need back-and-forth links between the various runs involved in exploring the stochastic space. For instance, past developments of uncertainty propagation such as adaptive importance sampling schemes have been designed with very limited link to the issue of computer implementation, ending up with purely-sequential formulations that fail to fully optimize the avenues offered by distributed computing to minimise the overall computing time. Adding the fact that parallel computing may be necessary to run a single simulation of complex underlying best-estimate models, optimizing the overall stochastic program becomes an insufficiently-researched domain.

## 2.4 Unified Simulation Framework and associated services

Advancing individual solvers performance is not enough to bring high performance simulation to the end-user. Each community needs a much broader set of tools in order to conduct industrial studies: CAD, mesh generation, data setting tools, computational scheme editing aids, visualization, etc.



**Figure 10.** The Salome platform, [www.platform-salome.org](http://www.platform-salome.org)

In the early 2000 EDF, together with CEA and other industrial and academic partners, started the development of an integrated toolbox Salome [www.platform-salome.org](http://www.platform-salome.org) [Ribe,Berg], with the following aims:

- reduce the cost of complex simulation platforms by mutualizing a set of common tools: pre and post-processing, calculation distribution and supervision etc.
- boost performance through easy integration of multiple solvers for multi-physics studies (via a common data model).

If Salome has been proved to be well adapted for sequential and moderately parallel simulations it has to evolve in order to support **massively parallel computing**.

**Possible joint efforts.** Building a Unified Simulation Framework and associated services adapted to massively parallel simulation:

- Common data model : designing a common data model and associated libraries for mesh and field exchange adapted to massively parallel computing would enable interoperability and the coupling of independent parallel scientific softwares. High level operations on simulation data, such as mesh projection, data interpolation, could be implemented on top of this model.
- Meshing. In 2007 it took to the EDF CFD team several months to produce the  $10^8$  cells mesh for the simulation of part of a fuel assembly with the CFD code Saturne, compare to “only” 1 month of calculation needed on 8000 BG/L processors. Generating  $\times 10^{10}$  cells mesh as targeted in 2015, requires future meshing tools to provide parallel meshing, automatic hexahedral meshing, mesh healing, CAD healing for meshing and dynamic mesh refinement.

**As an example are future works identified by a CFD Saturne code:**

- Re-evaluate if tetrahedra are really that bad
- Our extended neighborhood gradient reconstruction scheme should reduce impact of non-orthogonality
- Having mesh refinement algorithms would help, even if we don't do AMR right away
- Some octree-based techniques lead to fully hexahedral meshes:
  - · conforming using stencils and smoothing

- non-conforming with hanging nodes, using building-cube type method (also used by several codes, such as the Gerris Flow solver), combined with cut cells or immersed boundary
  - · At first, re-meshing on a low-quality, easily generated background would avoid issues with CAD interpretation and allow to easily define the local cell target size
- Using hierarchical techniques would also make multi-resolution visualization possible
  - · We have been luckier with visualization than with meshing, but tools and formats have their limits
- Parallel visualisation tools. Considering the volume of data that will be produced by Petaflop and Exaflop computers, end users are needed adapted parallel visualisation tools and specific clusters to post-treat their simulation results. The international scientific community would benefit in focusing their research efforts in few software. VISIT and Paraview seem two good candidates.
- Remote and collaborative post-treatment: the sheer volume of data produced by Petaflop/Exaflop calculations, storage and network limitations, and multi-sites teams make it necessary to further advance R&D on remote and collaborative multi-user visualisation, parallel and distributed file systems.
- Supervising and code coupling tool, coupling schemes: EDF and CEA have engaged in 2006 the development of YACS, a new generation of supervisor, intended to handle parallel multi-physics coupling scheme through a portable parallel extension to CORBA named PACO++[13] developed by INRIA. Similar works are handle in US, based on different middlewares. Implementing tightly coupled scheme, involving scientific applications developed by separated teams with such generic tools is a particularly difficult challenge. A joint collaboration on code coupling tool architecture principle, middleware for massively parallel coupled simulations seems indispensable.

The coupling using an external tool such as YACS is as less intrusive in the legacy codes as possible. On the other hand, we share the advanced coupling algorithms for all multi-physics simulation in a dedicated algorithmic box in the SALOME platform. From an algorithmic point of view, the existing couplings are mainly explicit and semi-implicit (fixed point algorithm). Current work are performed to implement Newton-like algorithms.

## Bibliography

- [Berg] Bergeaud V and Tajchman M 2007 Application of the SALOME software architecture to nuclear reactor research *SCS Spring Simulation Multiconference on High Performance Computing Symposium*, Norfolk, USA
- [Dong] J.Dongarra. *Freely available software for linear algebra on the web (sept 2006)*. <http://www.netlib.org/utk/people/JackDongarra/la-sw.html>.
- [Hame] “Jean-François Hamelin and Jean-Yves Berthou, **Getting ready for petaflop capacities and beyond: a utility perspective**, 2008 *J. Phys.: Conf. Ser.* 125 012001, July 2008”
- [Ribe] Ribes A and Caremoli C 2007 SALOME platform component model for numerical simulation *COMPSAC* july, Beijing, China
- Mumps] Mumps’ web page. <http://www.mumps.enseeiht.fr>.
- [NP] B.Secher, M.Belliard & C.Calvin. *Numerical Platon: a unified linear equation solver interface by CEA for solving open foe scientific applications*. Nuclear Engineering and design, vol. 239-1, pp87-95 (2009).
- [Past] PaStiX’s web page. <http://pastix.gforge.inria.fr/files/README-txt.html>.
- [Rocq] de Rocquigny E., Devictor N., Tarantola ed. (2008), *Uncertainty in industrial practice – A guide to Quantitative Uncertainty Management*, John Wiley & Sons