



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team ScAlApplix

*High Performance Schemes and Algorithms
for Complex Scientific Applications*

Bordeaux - Sud-Ouest

THEME NUM

Activity
R *eport*

2008

Table of contents

1. Team	1
2. Overall Objectives	2
2.1. Introduction	2
2.2. Highlights of the year	2
3. Scientific Foundations	3
3.1. Numerical schemes for fluid mechanics	3
3.2. Multiresolution and multiscale schemes and algorithms	4
3.2.1. Multiresolution algorithms	4
3.2.2. Multiscale algorithms	4
3.3. Algorithms and high-performance solvers	5
3.3.1. High-performance direct solvers for distributed clusters	5
3.3.2. High-performance iterative solvers	5
3.3.3. Fast multipole methods	6
3.4. Computational steering for parallel and distributed numerical simulations	6
4. Application Domains	8
4.1. Introduction	8
4.2. Fluid mechanics	8
4.3. Material physics	8
4.4. High performance simulation dedicated to ITER project	9
5. Software	9
5.1. Introduction	9
5.2. EPSN	10
5.3. FluidBox	10
5.4. MUMPS	11
5.5. PaStiX	11
5.6. HIPS	12
5.7. LibMultiScale	12
5.8. Scotch	13
6. New Results	14
6.1. Numerical schemes and algorithms for fluid mechanics.	14
6.1.1. Residual distribution schemes	14
6.1.2. Uncertainty quantification	14
6.1.3. Finite volume schemes for complex applications	15
6.1.4. Interface, multifluid problems and phase transition	15
6.1.5. Parallelization and fluid mechanics	15
6.1.6. Mesh adaptation	15
6.2. Multiscale schemes and algorithms	16
6.3. High performance simulation for plasma and beam physics	16
6.3.1. MHD instabilities	16
6.3.2. Very scalable semi-Lagrangian Vlasov solver	17
6.3.3. Adaptive Vlasov solver with low memory consumption	17
6.3.4. Dynamic load balancing for PIC solvers	17
6.4. Algorithms and high-performance solvers	17
6.4.1. Parallel domain decomposition and sparse matrix reordering	18
6.4.2. High-performance direct solvers on multi-platforms	19
6.4.3. High-performance hybrid direct-iterative solvers for large sparse systems	20
6.4.4. Parallel fast multipole method	21
6.4.5. Domain decomposition to solve neutron transport equations	22
6.4.6. High performance algorithms for depth imaging	22

6.5.	Computational steering environment for distributed numerical simulations	22
7.	Contracts and Grants with Industry	23
8.	Other Grants and Activities	24
8.1.	National initiatives	24
8.1.1.	AEROCV: Experimental and numerical study of the aeroacoustic in cylindric cavities	24
8.1.2.	ASTER: Adaptive MHD Simulation of Tokamak Elms for iteR	24
8.1.3.	HALOBAR: Dark matter and the dynamical evolution of barred galaxies	25
8.1.4.	MASSIM: development of a software environment for interactively processing and visualizing complex large scale data coming from numerical simulations	25
8.1.5.	NOSSI: New platform for parallel, hybrid quantum/classical simulations	25
8.1.6.	NUMASIS: adaptation and optimization of the applicative performance on NUMA architectures	26
8.1.7.	PETAL	26
8.1.8.	SOLSTICE: High performance solvers and simulations	26
8.1.9.	SYNERGHY: Simulate the beam, gravity and hydrodynamic phenomenon	27
8.2.	International initiatives	27
8.2.1.	Associated team PHyLeaS	27
8.2.2.	ADIGMA: a european project on the development of adaptive higher order variational methods for aerospace applications	28
8.2.3.	France-Berkeley Fund Award	28
8.2.4.	Collaboration project with JAEA (Japan)	29
9.	Dissemination	29
9.1.	Participation to the Scientific Community	29
9.2.	Teaching	29
10.	Bibliography	29

SCALAPPLIX is an INRIA Project-Team joint with University of Bordeaux (UBI and ENSEIRB) and CNRS (IMB, UMR 5251 and LaBRI, UMR 5800)

1. Team

Research Scientist

Olivier Coulaud [INRIA, Research Director (DR)]
Pascal Hénon [INRIA, Research Associate (CR)]
Mario Ricchiuto [INRIA, Research Associate (CR)]

Faculty Member

Jean Roman [INRIA, Team Leader, Research Director (on secondment), Professor at ENSEIRB, HdR]
Rémi Abgrall [University Bordeaux 1, Professor at MatMéca, IMB, INRIA Research scientist (on secondment since Oct. 2008), HdR]
Cécile Dobrzynski [MatMéca, Assistant Professor (MdC)]
Aurélien Esnard [University Bordeaux 1, Assistant Professor (MdC)]
Abdou Guermouche [University Bordeaux 1, Assistant Professor (MdC)]
Boniface Nkonga [MatMéca, Senior assistant Professor, Professor in Nice Sophia Antipolis since Nov. 2008, HdR]
Pierre Ramet [University Bordeaux 1, Assistant Professor (MdC)]
Guillaume Latu [University of Strasbourg, Assistant Professor (MdC), INRIA research scientist (partial secondment ended in Jul. 2008)]
François Pellegrini [ENSEIRB, Associate Professor (MdC)]

External Collaborator

Patrick Amestoy [ENSEEIHT (École Nationale Supérieure d'Électrotechnique, d'Électronique, d'Informatique, et d'Hydraulique de Toulouse), IRIT (Institut de Recherche en Informatique de Toulouse, UMR 5505), Professor (Pr), HdR]
Guillaume Anciaux [EPFL, PostDoc]
Pierre Fortin [LIP6, Assistant Professor (MdC)]
Yousef Saad [University of Minnesota, USA, Professor (Pr), HdR]

Technical Staff

Guilhem Caramel [INRIA, Expert Engineer, funding from ANR ARA Masses de données MASSIM, till Dec. 2008]
Damien Genet [INRIA, funding from ANR ARA CIS NOSSI, since Sep. 2008]
Pascal Jacq [INRIA, Expert Engineer, funding from world competitiveness cluster AESE, MACAO Project]
Mathieu Souchaud [INRIA, Associate Engineer, until Nov. 2008]
Xavier Lacoste [INRIA, Associate Engineer]

PhD Student

Guillaume Baurin [SNECMA, CIFRE, since Oct. 2008]
Rached Abdelkhalek [TOTAL, CIFRE, since Jan. 2008]
Aubin Bellue [University Bordeaux 1, MESR grant, since Sep. 2006]
Marie Billaud [University Bordeaux 1, funding from CEA, since Sep. 2006]
Mathieu Chanaud [INRIA, funding from INRIA and CEA, since Dec. 2007]
Fabrice Dupros [funding from BRGM and ANR CIGC NUMASIS, since Feb. 2007]
Mathieu Faverge [INRIA, funding from ANR CIGC NUMASIS, since Nov. 2006]
Sébastien Fourestier [CNRS, funding from CNRS and Conseil Régional d'Aquitaine, since Sep. 2008]
Jérémy Gaidamour [CNRS, funding from CNRS and Conseil Régional d'Aquitaine, since Oct. 2006]
Robin Huart [INRIA, funding from ANR CIS ASTER, since Oct. 2007]
Adam Larat [INRIA, funding from ADIGMA, since Oct. 2006]
Bruno Lathuilière [EDF, CIFRE, since Nov. 2006]

Nicolas Richart [INRIA, Conseil Régional d'Aquitaine grant, till Sep. 2008, University Bordeaux 1, ATER, since Oct. 2008]

Jiri Treflík [INRIA, funding from AEROCV and CTU, Prag, since Oct. 2008]

Christelle Wervaecke [INRIA, Conseil Régional d'Aquitaine grant, since Sep. 2007]

Post-Doctoral Fellow

Charles-Edmond Bichot [INRIA, until Jul. 2008]

Jun-Ho Her [INRIA, since Sep. 2007]

Administrative Assistant

Josy Baron [INRIA]

2. Overall Objectives

2.1. Introduction

ScAlApplix is a joint project of INRIA Futurs, LaBRI (Laboratoire Bordelais de Recherche en Informatique – CNRS UMR 5800, University of Bordeaux 1 and ENSEIRB) and IMB (Institut Mathématique de Bordeaux – CNRS UMR 5251, Universities of Bordeaux 1 and Bordeaux 2). ScAlApplix has been created on the first of November, 2002 (<http://www.labri.fr/scalapplix>).

The purpose of the ScAlApplix project is to analyze and solve scientific computation problems arising from complex research and industrial applications and involving scaling. These applications are characterized by the fact that they require enormous computing power, on the order of tens or hundreds of teraflops, and that they handle huge amounts of data. Solving these kinds of problems requires a multidisciplinary approach concerning both applied mathematics and computer science. In applied mathematics, it is essentially the field of numerical schemes that is concerned. In computer science, parallel computing and the design of high-performance codes to be executed on today's major computing platforms are concerned (parallel computers organized as a large network of SMP nodes, production GRIDs).

Through this approach, ScAlApplix intends to contribute to all steps in the line that goes from the design of new high-performance, more robust and more precise, numerical schemes to the optimized implementation of algorithms and codes for the simulation of *physical* (fluid mechanics, inert and reactive flows, multimaterial and multiphase flows) and *chemical* (molecular dynamic simulations).

Another domain we are currently investigating is the development of distributed environments for coupling numerical codes and for steering interactively numerical simulations. The computational steering is an effort to make the typical simulation work-flow (modeling, computing, analyzing) more efficient, by providing on-line visualization and interactive steering over the on-going computational processes. On-line visualization appears very useful to monitor and detect possible errors in long-running applications, and interactive steering allows the researcher to alter simulation parameters on-the-fly and immediately receive feedback on their effects. Thus, scientists gain a better insight in the simulation regarding to the cause-effect relationship and can better grasp the complexity of the underlying models.

2.2. Highlights of the year

- Scotch 5.1 has been released, providing state-of-the-art parallel graph partitioning routines to the PT-Scotch library.
- PaStiX 5.1.1 has been released and has been successfully used by CEA/CESTA to solve a symmetric complex sparse linear system arising from a 3D electromagnetism code with more than 83 million unknowns on the TERA-10 CEA supercomputer. Solving this system required about 5 Petaflops (in double precision) and the task was completed in about 5 hours on 768 processors. To our knowledge, a system of this size and this kind has never been solved by a direct solver.

- HIPS 1.0 has been released and declared to the APP agency. HIPS provides state-of-the-art hybrid direct/iterative and fully iterative parallel solvers based on the Schur complement approach. HIPS uses a domain decomposition-like approach and is highly scalable in terms of memory and run-time.
- FluidBox has been declared at the APP agency and the first very large scale simulations have been run; first calculations with very large meshes for the ADIGMA project (3 million vertices) have been performed.
- The LibMultiScale software package and library have been released as an open source software under the CECILL-C license. Multiscale simulations with millions of atoms have been performed for 3D wave and crack propagations.

3. Scientific Foundations

3.1. Numerical schemes for fluid mechanics

Keywords: *complex physical models, non upwind schemes, unstructured meshes, upwind schemes.*

A large number of industrial problems can be translated into fluid mechanics ones. They may be coupled with one or more physical models. An example is provided by aeroelastic problems, which have been studied in details by other INRIA teams. Another example is given by flows in pipelines where the fluid (a mixture of air–water–gas) does not have well-known physical properties. One may also consider problems in aeroacoustics, which become more and more important in everyday life. In some occasions, one needs specific numerical tools because fluids have exotic equation of states, or because the amount of computation becomes huge, as for unsteady flows. Another situation where specific tools are needed is when one is interested in very specific quantities, such as the lift and drag of an airfoil, a situation where commercial tools can only provide a very crude answer.

It is a fact that there are many commercial codes. They allow users to consider some of these examples, but the quality of the solutions is far from being optimal. Moreover, the numerical tools of these codes are often not the most recent ones. An example is the noise generated by vortices crossing through a shock wave. It is, up to our knowledge, even out of reach of the most recent technologies because the numerical resources that would necessitate such simulations are tremendous ! In the same spirit, the simulation of a 3D compressible mixing layer in a complex geometry is also out of reach because very different temporal and physical scales need to be captured, thus we need to invent specific algorithms for that purpose.

In order to reach efficient simulation of complex physical problems, we are working on some fundamental aspects of the numerical analysis of non linear hyperbolic problems. Our goal is to develop schemes that can adapt to the modern computer architectures. More precisely, we are working on a class of numerical schemes specifically tuned for unstructured and hybrid meshes. They have the most possible compact stencil that is compatible with the expected order of accuracy. The order of accuracy typically ranges from two to four. Since the stencil is compact, the implementation on parallel machines becomes simple. The price to pay is that the scheme is necessarily implicit. However, the compactness of the scheme enables to use the high performance parallel linear algebra tools developed by the team for the lowest order version of these schemes. The high order versions of these schemes, that are still under development, will lead to new scientific problems at the border between numerical analysis and computer science. In parallel to these fundamental aspects, we also work on adapting more classical numerical tools to complex physical problems such as those encountered in interface flows, turbulent or multiphase flows.

Within a few years, we expect to be able to consider the physical problems that are now difficult to compute thanks to the know–how coming from our research on compact distribution schemes and the daily discussions with specialists of computer science and scientific computing. These problems range from aeroacoustic to multiphysic problems, such as the one mentioned above. We also have interest in solving compressible MHD problems in relation with the ITER project. Because of the existence of a magnetic field and the type of solutions we are seeking for, this leads to additional scientific challenges. Our research work about numerical algorithms has led to software FluidBox that is described in section 5.3. This work is supported by the EU-Strep ADIGMA, various research contracts and in part by the ANR-CIS ASTER project (see section 4.4 also).

3.2. Multiresolution and multiscale schemes and algorithms

Keywords: *adaptive mesh refinement, code coupling, molecular dynamics, multiresolution scheme, multiscale scheme, parallelism and distributed computing.*

The large scale applications which we study are built in order to capture complex phenomena such as magnetohydrodynamic instabilities, cracks and so on. Such phenomena may be very localized in the domain of simulation and a full modeling or discretization at the finest level require several tens or hundreds teraflops of computation using several terabyte of data. Many of these phenomena are multiscale both in time and in space. To capture them efficiently, we need to introduce multiresolution schemes as adaptive mesh refinement, wavelet methods or even multiscale models when the current model is no more valid at the finest level of discretization.

3.2.1. Multiresolution algorithms

Adaptive techniques based on a hierarchical decomposition could lead to a dramatic reduction of computation and memory costs in numerical simulation. It turns out that adaptive numerical methods are often difficult to parallelize, because they introduce dependencies between data at different grid levels and it is then difficult to manage data locality. Work has to be done in order to build efficient adaptive parallel methods. For example, in an wavelet framework, a detail is defined as the difference between the real value of the function of interest and an interpolated value. We have a natural criterion of compression since the detail is small where the function is regular. Compression of a wavelet decomposition is performed by removing grid points with detail smaller than a given threshold. One could write numerical schemes in which only the significant details are computed and others are skipped in order to reduce both memory usage and computation cost. Parallel simulations that use interpolets have been performed in our team. It allows us to perform simulation on very large grid sizes that are not accessible with non-adaptive simulators. The main challenge is to design parallel numerical schemes that do not lead to large overheads in terms of load balancing, sparsity management, and interprocessor communication.

In the context of ASTER project, the study of MHD instabilities is increasingly complex in the sense that they require some simulations at a very large scale; several tens or hundreds teraflops of computation using several terabyte of data are needed, and they use multiphysic or multiscale modelling giving place to fine couplings. To carry out these more and more accurate full-scale simulations without increasing the number of unknowns in a uniform way, the AMR technique consists in handling a finer grid where the solution varies abruptly and a coarser grid at other places accordingly to quite specific criteria. These new finer grids are regarded as patches to the initial grid. One can thus obtain a hierarchy of grids by repeating this procedure.

This approach is broadly used when explicit numerical schemes are considered. In the case of numerical implicit schemes (as implemented in JOREK and FLUIDBOX), the difficulty is to take into account the additional unknowns due to the hierarchy of grids in the modified linear system that must be solved at each time step of the simulation.

A specific study of the error estimators adapted to the problem of MHD instabilities needs to be carried out, in particular in connection with the discretization with cubic Hermite (Bezier) finite elements used in JOREK.

In the context of high performance computing, we have to study within a parallel framework the principal difficulties encountered in these methods. The efficiency will therefore require the use of dynamic load balancing methods over the set of processors to compensate the variability of the grid during the simulation.

3.2.2. Multiscale algorithms

Due to the increase of available computer power, new applications such as failure material simulations like crack propagation are now commonly performed by physicists. These computations simulate systems up to billion of atoms in materials, for large time scales up to several nanoseconds. The larger the simulation is, the smaller the computational cost of the potential driving the phenomena is, resulting in low precision results. Moreover, full simulations at the finest level are not computationally feasible on the whole materials. Most of the time, the finest level is only necessary where the phenomena of interest occurs, for example, in a crack

propagation simulation far from the tip, we have a macroscopic behavior of the material and then we can use a coarser model. The idea is to limit the more expensive level simulation to a subset of the domain and to combine it with a macroscopic level. But combining quantum and atomistic or atomistic and continuum simulations are still a challenge to obtain a robust model, a good scheme and an efficient implementation.

We are currently focusing on several aspects of these problems:

- The difficulty is to build an efficient scheme which couples the two different scales without any loss of precision and without introducing any numerical artifacts like wave reflections. We are working to improve our reduction operator in order to obtain a more robust one with respect to the size of the mesh and the frequency of atomistic waves.
- The algorithm we have developed is valid for crystals at zero Kelvin but heat flux is very important for realistic simulations. The main questions for us are what kind of model must be considered at the coarser level, how we transfer atomic kinetic energy in the continuum model.
- Many of such couplings are realized by putting together different codes. A main problem is to define a coupler that can efficiently exploit the parallelism of the underlying codes. In this context, we focus on the parallel algorithmic of the coupler such as data redistribution, how we can conserve the efficiency of each code in the distributed application and obtain a good efficiency for the global application.

3.3. Algorithms and high-performance solvers

Keywords: *fast multipole methods, high-performance computing, parallel sparse linear algebra.*

3.3.1. High-performance direct solvers for distributed clusters

Solving large sparse systems $Ax = b$ of linear equations is a crucial and time-consuming step, arising in many scientific and engineering applications. Consequently, many parallel techniques for sparse matrix factorization have been studied and implemented.

Sparse direct solvers are mandatory when the linear system is very ill-conditioned; such a situation is often encountered in structural mechanics codes, for example. Therefore, to obtain an industrial software tool that must be robust and versatile, high-performance sparse direct solvers are mandatory, and parallelism is then necessary for reasons of memory capability and acceptable solving time. Moreover, in order to solve efficiently 3D problems with more than 50 million unknowns, which is now a reachable challenge with new SMP supercomputers (see section 2.2), we must achieve good scalability in time and control memory overhead. Solving a sparse linear system by a direct method is generally a highly irregular problem that induces some challenging algorithmic problems and requires a sophisticated implementation scheme in order to fully exploit the capabilities of modern supercomputers.

In the ScAlApplix project, we focused first on the block partitioning and scheduling problem for high performance sparse LDL^T or LL^T parallel factorization without dynamic pivoting for large sparse symmetric positive definite systems. Our strategy is suitable for non-symmetric sparse matrices with symmetric pattern, and for general distributed heterogeneous architectures whose computation and communication performances are predictable in advance.

Research about high performance sparse direct solvers is carried on in collaboration with P. Amestoy (ENSEEIH – IRIT) and J.-Y. L'Excellent (INRIA Rhône-Alpes), and has led to software developments (see section 5.4, 5.5, 5.8) and to industrial contracts with CEA (Commissariat à l'Énergie Atomique). Those works are supported by the ANR-CIS project "SOLSTICE".

3.3.2. High-performance iterative solvers

In addition to the project activities on direct solvers, we also study some robust preconditioning algorithms for iterative methods. The goal of these studies is to overcome the huge memory consumption inherent to the direct solvers in order to solve 3D dimensional problems of huge size (several million of unknowns). Our studies focus on the building of generic parallel preconditioners based on ILU factorizations. The classical ILU

preconditioners use scalar algorithms that do not exploit well the CPU power and are difficult to parallelize. Our work aims at finding some unknown orderings and partitionings that lead to a dense block structure of the incomplete factors. Then, based on the block pattern, some efficient parallel blockwise algorithms can be devised to build robust preconditioners that are also able to exploit the full capabilities of the modern high-performance computers.

We study two approaches.

- The first approach consists in building block ILU(k) preconditioners. The main idea is to adapt the classical ILU(k) factorization in order to reuse the algorithmic ingredients that have been developed for direct methods. In this case, the ordering we use is the same than in the direct factorization and a dense block pattern (i.e. a partition of the unknowns) is obtained using an algorithm that lumps columns having few differences in their non zeros pattern. We have adapted the parallel direct solver chain in order to deal with the incomplete block factors defined by this process. Thus the preconditioner computation benefits from the breakthroughs made by the direct solver techniques studied in PaStiX (sections 5.5 and 6.4).
- The second approach we recently developed is based on the Schur complement approach. In this case, we use a partition of the adjacency graph of the system matrix into a set of small subdomains with overlap. The interior of these subdomains are treated by a direct method. Solving the whole system is then equivalent to solve the Schur complement system on the interface between the subdomains (this system has a much smaller dimension). We use the hierarchical interface decomposition (HID) that has been developed in PHIDAL to reorder and partition this system. Indeed, the HID gives a natural dense block structure of the Schur complement. Based on this partition, we define some efficient block preconditioners that allow the use of BLAS routines and a high degree of parallelism thanks to the HID properties. All these algorithms are implemented in the library HIPS. Details can be found in sections 5.6 and 6.4.

These works are also supported by the ANR-CIS project “SOLSTICE”.

3.3.3. Fast multipole methods

In most of scientific computing applications considered nowadays as computational challenges like biological systems, astrophysics or electromagnetism, the introduction of hierarchical methods based on an octree structure has dramatically reduced the amount of computation needed to simulate those systems for a given error tolerance.

Among these methods, the Fast Multipole Method (FMM) allows the computation of the interactions in, for example, a molecular dynamics system of N particles in $O(N)$ time, against $O(N^2)$ with a direct approach. The extension of these methods and their efficient implementation on current parallel architectures is still a critical issue. Moreover the use of periodic boundary conditions, or of duplications of the system in 2 out of 3 space dimensions, just as well as the use of higher approximations for integral equations are also still relevant.

They can be used in the three (quantum, molecular and continuum) models for atom-atom interactions in quantum or molecular mechanics, atom-surface interactions for the coupling between continuum and the other models, and also for fast matrix-vector products in the iterative solving of the linear system given by the integral formulation of the continuum method. Moreover, the significant experience achieved by the Scotch and PaStiX projects (see section 5.8 and 5.5) will be useful in order to develop efficient implementations of the FMM methods on parallel clusters of SMP nodes.

3.4. Computational steering for parallel and distributed numerical simulations

Keywords: *computational steering, coupling, interaction, numerical simulation, scientific visualization.*

Thanks to the constant evolution of computational capacity, numerical simulations are becoming more and more complex; it is common to couple different models in different distributed codes running on supercomputer or heterogeneous grids. Nowadays, such simulations are typically running in batch mode, and the analysis of the results is then performed on a local workstation as a post-processing step, which implies to preliminarily collect all the simulation output files. In most research fields, the post-processing step is realized thanks to 3D scientific visualization techniques. In the batch approach, there is a lack of control over the in-progress computations, and it might drastically decrease the profitability of the computational resources (repeated tests with different input files separated by excessively long waiting periods).

For years, the scientific computing community has expressed the need of new computational steering tools. The computational steering is an alternative approach to the typical simulation work-flow of performing computation and visualization sequentially. It mainly consists in coupling a remote simulation with a graphics system through the network in order to provide scientists with on-line visualization and interactive steering. On-line visualization appears very useful to monitor the evolution of the simulation by rendering the current results. It also allows us to validate the simulation codes and to detect conceptual or programming errors before the completion of a long-running application. Interactive steering allows the researcher to change the parameters of the simulation without stopping it. As the on-line visualization provides an immediate visual feedback on the effect of a parameter change, the scientist gains additional insight in the simulation, regarding to the cause-effect relationship. Such a tool might help the scientist to better grasp the complexity of the underlying models and to drive more rapidly the simulation *into the right direction*. Basically, a computational steering environment can be defined as a communication infrastructure, coupling a simulation with a remote user interface, called steering system. This interface usually provides on-line visualization and user interaction. Over the last decade, many steering environments have been developed. They distinguish themselves by their simulation integration process. A first solution is the problem solving environment (PSE) approach, like in SCIRun. This approach allows the scientist to construct a steering application according to a visual programming model. As an opposite, the majority of the steering environments, such as the well-known CUMULVS, are based on the annotation of the application source-code. This latter approach allows fine grain steering functionalities and can achieve better run-time performances.

Even though most of existing computational steering environments, such as CUMULVS, DAQV or gViz support parallel simulations, they are limited to sequential visualization systems. This leads to an important bottleneck and increased rendering time. In the gViz project, the IRIS Explorer visualization system has been extended to run the different modules (simulation, visualization, rendering) in a distributed fashion on the Grid, but the visualization and the rendering modules are still sequential. Recent works in the Uintah PSE (Problem Solving Environment) has addressed the problem of massively parallel computation connected to a remote parallel visualization module, but this latter module is only running on a shared-memory machine. Therefore, it would be particularly *valuable* for the scientist if a steering environment would be able to perform parallel visualization using a PC-based graphics cluster.

In the EPSN project, we intend to explore this latter purpose, called *M×N computational steering*. More precisely, the EPSN environment (Environment for the Steering of Parallel Numerical Simulations, see section 5.2) enables to interconnect parallel simulations with visualization systems, that can be parallel as well. In other words, we want to provide a framework that can benefit from immersive virtual reality technology (e.g. tiled display wall) and that might help scientists to better grasp the complexity of real-life simulations. Such a coupling between parallel numerical simulations and parallel visualization systems raises crucial issues we investigate in the EPSN project:

Parallel coordination of steering operations. The on-line visualization and the computational steering of parallel simulations come up against a serious coherence problem. Indeed, data distributed over parallel processes must be accessed carefully to ensure they are presented to the visualization system in a meaningful way. To overcome this problem, we introduce a Hierarchical Task Model (HTM) that represents the control-flow of the simulation, too often considered in other approaches as a “single-loop” program. Thanks to this representation, we schedule in parallel the user interaction on the simulation processes and satisfy the temporal coherence.

Parallel data redistribution. The efficient redistribution of data is a crucial issue to achieve an efficient coupling between simulation and visualization codes. Most of the previous works have limited their studies to structured data with regular distributions (e.g dense arrays with simple block-cyclic distributions). However, computational steering applications (or multiphysic simulations) frequently handle more irregular data, like particle sets, unstructured meshes or hierarchical grids. In such a context, the data transfer involves to switch from one distribution to another, and thus to “redistribute” the data. For this, we introduce a description model of distributed data, based on the notion of “complex object”. Thanks to this model, we propose redistribution algorithms that generate symbolic messages, independently from a particular communication layer. We distinguish two main approaches: the spatial approach and the placement approach.

4. Application Domains

4.1. Introduction

The main objective of the ScAlApplix project is to analyze and solve scientific computing problems coming from complex research and industrial applications that require a scalable approach. This allows us to validate the numerical schemes, the algorithms and the associated softwares that we develop. We have today three reference application domains which are fluid mechanics, material physics and the MHD simulation dedicated to the ITER project. In these three domains, we study and simulate phenomena that are by nature multiscale and multiphysic, and that requires enormous computing power. A major part of these works leads to industrial collaborations in particular with the CNES, ONERA, and with the french CEA/CESTA, CEA/Ile-de-France and CEA/Cadarache centers.

4.2. Fluid mechanics

Keywords: *MHD flows, fluid mechanics, multiphase flows, steady and unsteady flows, unstructured meshes.*

The numerical simulation of steady and unsteady flows is still a challenge since efficient schemes and efficient implementations are needed. The accuracy of schemes is still a problem nowadays. This challenge is even higher if large size problems are considered, and if the meshes are not regular.

Among the problems to be considered, one may list the computation of mixing layers, shock–vortices interactions, the noise generated by a flow. This last item clearly needs very high order schemes, and the today best schemes use regular structured meshes. Hence, one of our objectives is to construct very high order schemes for unstructured meshes.

Another example where large computer resources are needed is the simulation of multiphase flows. In that case, several difficulties must be faced: unsteady flows, complex geometries and a very complex physical model. A last example is related to unsteady MHD flow problems related to the ITER project. In addition to “standard” difficulties arising in flow simulations, we have to take into account the magnetic field that satisfies a geometrical constraint. Mesh adaption is also essential for the instability problems we consider. This is done in collaboration with physicist at CEA Cadarache via the ANR CIS 2006 ASTER project.

4.3. Material physics

Keywords: *continuum method, crack propagation, molecular dynamics, multiscale method, quantum method.*

Due to the increase of available computer power, new applications such as study of properties of new materials (photovoltaic materials, bio- and environmental sensors,...) failure material simulations, nano-indentation are now commonly performed by chemists.

These computations simulate systems up to billion of atoms in materials, for large time scales up to several nanoseconds. The larger the simulation is, the smaller the computational cost of the potential driving the phenomena is, resulting in low precision results. Moreover, full simulations at the finest level are not computationally feasible on the whole materials. Most of the time, the finest level is only necessary where the phenomena of interest occurs for example in a crack propagation simulation far from the tip, we have a macroscopic behavior of the material and then we can use a coarser model. The idea is to limit the more expensive level simulation to a subset of the domain and to combine it with a macroscopic level. This implies that atomistic simulations must be speeded up by several orders of magnitude.

The study of crack propagation understanding is important in many area like lens for laser, ... but also to understand how the crack appears. To do this, the main difficulty is to take into account the temperature in the model.

4.4. High performance simulation dedicated to ITER project

Keywords: *ITER, MHD flows and instabilities, adaptive method, mesh refinement, numerical schemes, parallel application design.*

Numerical simulation has become a major tool for the study of many physical phenomena involving charged particles, in particular beam physics, space and laboratory plasmas including fusion plasmas. Moreover, it is a subject of interest to figure out and optimize physics experiments in the present fusion devices and also to design future reactors like in the ITER project. Parallelism is required to carry on numerical simulations for realistic test cases.

On this topic, we have a strong collaboration with the physicists of the CEA/DRFC group. Moreover, we worked with the INRIA CALVI team to bring our high performance skills to some of their numerical simulators. The study of plasma turbulence requires, for example, to solve the Maxwell equations coupled to the calculation of the plasma response to the perturbed electromagnetic field. This response can be computed by using either a fluid or a kinetic description of the plasma. We have contributed to kinetic (in solving the Vlasov equation) and fluid (through the study of MHD models) simulators.

We have established an other collaboration with the physicists of the CEA/DRFC group in the context of the ANR CIS 2006 project called ASTER (Adaptive mhd Simulation of Tokamak Elms for iteR). The magnetohydrodynamic instability called ELM for Edge Localized Mode is commonly observed in the standard tokamak operating scenario. The energy losses the ELM will induce in ITER plasmas are a real concern. However, the current understanding of what sets the size of these ELM induced energy losses is extremely limited. No numerical simulations of the complete ELM instability, from its onset through its non-linear phase and its decay, are referenced in literature. Recently, encouraging results on the simulation of an ELM cycle have been obtained with the JOREK code developed at CEA but at reduced toroidal resolution. The JOREK code uses a fully implicit time evolution scheme in conjunction with the PaStiX sparse matrix library.

5. Software

5.1. Introduction

We develop two kinds of software. The first one consists in generic libraries that are used in application codes. These libraries comprise a sequential and parallel partitioner for large irregular graphs or meshes (Scotch), and high performance direct or hybrid solvers for very large sparse systems of equations (MUMPS, PaStiX, HIPS). The second kind of software corresponds to dedicated software for fluid mechanics (FluidBox), and to a platform for computational steering (EPSN).

For these parallel software developments, we use the message passing (MPI) paradigm, the OpenMP programming language, threads, and the Java and/or CORBA technologies.

5.2. EPSN

Keywords: *code coupling, computational steering, data redistribution, numerical simulation, scientific visualization.*

Participants: Olivier Coulaud, Aurélien Esnard [corresponding member], Nicolas Richart, Mathieu Souchaud, Guilhem Caramel.

The EPSN project has been partially supported by both the ACI-GRID program from the french Ministry of Research (grant number PPL02-03), the ARC RedGRID, and is now supported by the ANR program called MASSIM (grant number ANR-05-MMSA-0008-03).

EPSN is a distributed computational steering environment which allows the steering of remote parallel simulations with sequential or parallel visualization tools or graphics user interface. It is a distributed environment based on a simple client/server relationship between user interfaces (clients) and simulations (servers). The user interfaces can dynamically be connected to or disconnected from the simulation during its execution. Once a client is connected, it interacts with the simulation component through an asynchronous and concurrent request system. We distinguish three kinds of steering request. Firstly, the "control" requests (play, step, stop) allow to steer the execution flow of the simulation. Secondly, the "data access" requests (get, put) allow to read/write parameters and data from the memory of the remote simulation. Finally, the "action" requests enable to invoke user-defined routines in the simulation. In order to make a legacy simulation steerable, the end-user annotates its simulation source-code with the EPSN API. These annotations provide the EPSN environment with two kinds of information: the description of the program structure according to a Hierarchical Task Model (HTM) and the description of the distributed data that will be accessible by the remote clients.

Concerning the development of client applications, we also provide a front-end API that enables to integrate EPSN in a high-level visualization system such as *VTK* or *ParaView*. We also provide a lightweight user interface, called *SiMonE* (Simulation Monitoring for EPSN), that enables to easily connect any simulations and interact with them, by controlling the computational flow, viewing the current parameters or data on a simple data-sheet and modifying them optionally. *SiMonE* also includes simple visualization plug-ins to online display the intermediate results. Moreover, the EPSN framework offers the ability to exploit parallel visualization and rendering techniques thanks to the Visualization ToolKit (VTK). This approach allows to reduce the steering overhead of the EPSN platform and allows to process efficiently large dataset. To visualize with high resolution image and to improve the rendering time, EPSN can also exploit tiled-display wall based on ICE-T library developed at Sandia Laboratory.

As both the simulation and the visualization can be parallel applications, EPSN is based on the $M \times N$ redistribution library called RedGRID. This library is in charge of computing all the messages that will be exchanged between the two parallel components, and is also in charge of performing the data transfer in parallel. Thus, RedGRID is able to aggregate the bandwidth and to achieve high performance. Moreover, it is designed to consider a wide variety of distributed data structures usually found in the numerical simulations, such as structured grids, pints or unstructured meshes.

Both EPSN and RedGRID use a communication infrastructure based on CORBA which provides our platform with portability, interoperability and network transparency.

Web: <http://www.labri.fr/projet/epsn> and <http://gforge.inria.fr/projects/epsn>.

5.3. FluidBox

Keywords: *MHD flows, finite volume schemes, fluid mechanics, inert and reactive flows, multimaterial and multiphase flows, upwind and residual distribution scheme.*

Participants: Rémi Abgrall, Cécile Dobrzynski, Robin Huart, Adam Larat, Boniface Nkonga [corresponding member], Mario Ricchiuto.

FluidBox is a software dedicated to the simulation of inert or reactive flows. It is also able to simulate multiphase, multimaterial and MHD flows. There exist 2D and 3D dimensional versions. The 2D version is used to test new ideas that are later implemented in the 3D one. Two classes of schemes have been implemented: classical finite volume schemes and the more recent residual distribution schemes. Several low Mach preconditioning techniques are also implemented. The code has been parallelized with and without overlap of the domains. Recently, the PaStiX solver has been integrated in FluidBox. FluidBox has also been coupled with the EPSN platform. A partitioning tool exists in the package, it uses Scotch. FluidBox has also benefited from many software and functionality improvements from Rémi Butel (IMB); up to now, it is only a private project, but we expect to open some part of the code to public before the end of the year. In order to facilitate the project development, FluidBox has been uploaded to the INRIA/Gforge page.

5.4. MUMPS

Keywords: *parallel asynchronous sparse multifrontal solver.*

Participants: Patrick Amestoy, Abdou Guermouche [corresponding member].

In the context of PARASOL (Esprit IV Long Term Project, 1996-1999), CERFACS and ENSEEIHT-IRIT teams have initiated a parallel sparse solver MUMPS (“MULTifrontal Massively Parallel Solver”). Since the first public release of MUMPS (March 2000), this research and (also software) project is the context of a tight and fruitful collaboration with J. Y. L’Excellent (INRIA-LIP-ENS Lyon) and the INRIA project GRAAL. Recent work related to out-of-core executions and performance scalability has been incorporated in the new improved version of the package (release 4.8.3 available since July 2008 at <http://mumps.enseeiht.fr/> or <http://graal.ens-lyon.fr/MUMPS>).

MUMPS is a package for solving linear systems of equations $Ax = b$, where the matrix A is sparse and can be either unsymmetric, symmetric positive definite, or general symmetric. It uses a multifrontal technique which is a direct method based on either the LU or the LDL^T factorization of the matrix. The main features of the MUMPS package include numerical pivoting during factorization, solution of the transposed system, input of the matrix in assembled format (distributed or centralized) or elemental format, error analysis, iterative refinement, scaling of the original matrix, and return of a Schur complement matrix. It also offers several built-in ordering algorithms, a tight interface to some external ordering packages such as Scotch and is available in various arithmetics (real or complex, single or double).

5.5. PaStiX

Keywords: *complete and incomplete supernodal sparse parallel factorizations.*

Participants: Mathieu Faverge, Pascal Hénon, Xavier Lacoste, François Pellegrini, Pierre Ramet [corresponding member], Jean Roman.

This work is supported by the French “Commissariat à l’Energie Atomique CEA/CESTA” in the context of structural mechanics and electromagnetism applications.

PaStiX (<http://pastix.gforge.inria.fr>) (Parallel Sparse matrix package) is a scientific library that provides a high performance parallel solver for very large sparse linear systems based on block direct and block ILU(k) iterative methods. Numerical algorithms are implemented in single or double precision (real or complex): LLt (Cholesky), LDLt (Crout) and LU with static pivoting (for non symmetric matrices having a symmetric pattern). This latter version is now used in FluidBox (see section 5.3). The PaStiX library is released under INRIA CeCILL licence.

The PaStiX library uses the graph partitioning and sparse matrix block ordering package Scotch (see section 5.8). PaStiX is based on an efficient static scheduling and memory manager, in order to solve 3D problems with more than 50 million of unknowns. The mapping and scheduling algorithm handles a combination of 1D and 2D block distributions. This algorithm computes an efficient static scheduling of the block computations for our supernodal parallel solver which uses a local aggregation of contribution blocks. This can be done by taking into account very precisely the computational costs of the BLAS 3 primitives, the communication costs and the cost of local aggregations. We also improved this static computation and communication scheduling algorithm to anticipate the sending of partially aggregated blocks, in order to free memory dynamically. By doing this, we are able to reduce dramatically the aggregated memory overhead, while keeping good performances.

Another important point is that our study is suitable for any heterogeneous parallel/distributed architecture when its performances are predictable, such as clusters of SMP nodes. In particular, we propose now a high performance version with a low memory overhead for SMP node architectures, which fully exploits shared memory advantages by using an hybrid MPI-thread implementation.

Direct methods are numerically robust methods, but the very large three dimensional problems may lead to systems that would require a huge amount of memory despite any memory optimization. A studied approach consists to define an adaptive blockwise incomplete factorization that is much more accurate (and numerically more robust) than the scalar incomplete factorizations commonly used to precondition iterative solvers. Such incomplete factorization can take advantage of the latest breakthroughs in sparse direct methods and particularly should be very competitive in CPU time (effective power used from processors and good scalability) while avoiding the memory limitation encountered by direct methods.

5.6. HIPS

Keywords: *domain decomposition, multilevel method, parallel iterative solver, schur complement.*

Participants: Jérémie Gaidamour, Pascal Hénon [corresponding member], Jean Roman, Yousef Saad.

HIPS (Hierarchical Iterative Parallel Solver) is a scientific library that provides an efficient parallel iterative solver for very large sparse linear systems. HIPS has been built on top of the PHIDAL library; it is based on the ordering and partitioning that were developed in PHIDAL and it proposes some new hybrid direct iterative algorithms. The keypoint of the methods implemented in HIPS is to define an ordering and a partition of the unknowns that relies on a form of nested dissection ordering in which cross points in the separators play a special role (Hierarchical Interface Decomposition ordering). The subgraphs obtained by the nested dissection correspond to the unknowns that are eliminated using a direct method and the Schur complement system on the remaining of the unknowns (that correspond to the interface between the subgraphs viewed as subdomains) is solved using an iterative method (GMRES or Conjugate Gradient at the time being). This special ordering and partitioning allow the use of dense block algorithms both in the direct and iterative part of the solver and provides a high degree of parallelism to these algorithms.

HIPS can be used as a stand-alone program that reads a sparse linear system from a file ; it also provides an interface to be called from any C, C++ or Fortran code.

It handles symmetric, unsymmetric, real or complex matrices. The code provides a hybrid method which blends direct and iterative solvers. HIPS exploits the partitioning and multistage ILU techniques developed in the PHIDAL algorithm to enable a highly parallel scheme whereby several subdomains can be assigned a process. It also provides the scalar preconditioner based on the multistage ILUT factorization that were developed in PHIDAL. Thus HIPS is a software library that provides several methods to build an efficient preconditioner in almost all situations [42], [51].

Since august 2008, HIPS is publicly available at <http://hips.gforge.inria.fr> under INRIA CeCILL licence.

5.7. LibMultiScale

Keywords: *code coupling, dynamic molecular simulation, elasticity, multiscale, parallelism.*

Participants: Guillaume Anciaux [corresponding member], Olivier Coulaud.

The multiscale coupling methods are a powerful tool to study the local phenomenon that occurs at the atomic scale. Such methods are generally employed to study crack propagations for example, dislocations or nano indentations. With these approaches, one couples an atomic description which models the material at the finest scale - such as molecular dynamics - with a macroscopic model of continuum mechanics. The use of a macroscopic model allows to reduce considerably the number of unknowns to handle. Moreover, it is easier to apply force fields or complex boundary conditions to macroscopic models. Thus, this type of coupling can help to provide complex boundary conditions to the molecular dynamics domain.

The LibMultiScale is a C++ parallel framework for the multiscale coupling methods dedicated to material simulations. This framework is designed with the form of a library providing an API which makes it possible to program coupled simulations. At the present time, the stable implemented coupling method is based on the bridging method from T. Belytschko and S. Xiao.

The coupled parts can be provided by existing projects. In such a manner, the API gives C++ templated interfaces to reduce to the maximum the cost of integration taking the form of plugins or alike. Such codes have been integrated to provide a functional prototype of the framework. For example, molecular dynamics software that have been integrated in Stamp (a code of the CEA) and Lammmps (Sandia laboratories). The unique software of continuum mechanics, discretized by finite elements, is based on the libMesh framework.

This software is the result of a collaboration between ScAlApplix and the CEA/DPTA Ile de France. The LibMultiScale is now distributed with a CECILL-C open-source licence. The LibMultiScale is available at (<http://libmultiscale.gforge.inria.fr/>)

5.8. Scotch

Keywords: *mesh partitioning, parallel graph partitioning, parallel sparse matrix block ordering, sparse matrix block ordering, static mapping.*

Participants: Sébastien Fourestier, Jun-Ho Her, François Pellegrini [corresponding member].

Scotch (<http://www.labri.fr/~pelegrin/scotch/>) is a software package for parallel and sequential sparse matrix ordering, parallel and sequential graph partitioning, as well as sequential static mapping, and mesh and hypergraph partitioning.

The initial purpose of Scotch was to compute high-quality partitions and static mappings of valuated graphs representing parallel computations and target architectures of arbitrary topologies. The original contribution consisted in developing a “*divide and conquer*” algorithm in which processes are recursively mapped onto processors by using graph bisection algorithms that are applied both to the process graph and to the architecture graph. This allows the mapper to take into account the topology and heterogeneity of the valuated graph which models the interconnection network and its resources (processor speed, link bandwidth). As new multicore, multinode parallel machines tend to be less uniform in terms of memory latency and communication bandwidth, this feature may possibly regain interest.

The software has then been extended in order to produce vertex separators instead of edge separators, using a multilevel framework. Recursive vertex separation is used to compute orderings of the unknowns of large sparse linear systems, which both preserve sparsity when factorizing the matrix and exhibit concurrency for computing and solving the factorized matrix in parallel. The original contribution has been to study and implement a tight coupling between the nested dissection and the approximate minimum degree methods; this work was carried out in collaboration with Patrick Amestoy, of ENSEEIHT-IRIT.

In version 4.0, released in February 2006, new data structures and methods have been added to the libScotch library, which allow it to compute efficient orderings of native meshes, resulting in the handling of larger problems than with standard sequential graph partitioners. Meshes are represented as bipartite graphs, in which node vertices are connected to element vertices only, and vice versa. Since this structure is equivalent to a hypergraph, where nodes are connected to hyper-edges only, and vice versa, the mesh partitioning routines of

Scotch turn it into a sequential hypergraph partitioner. This version was released as free/libre software under LGPL license, in order to encourage members of the community to use it as a testbed for the quick and easy development of new partitioning and ordering methods. It has been downloaded more than 2400 times from the InriaGForge repository.

Version 5.0 of Scotch [98], [97], which has been released last August, was the first version to comprise parallel routines. This extension, called PT-Scotch (for “*Parallel Threaded Scotch*”), is based on a distributed memory model, and makes use of the MPI and, optionally, Posix thread APIs. A distributed graph structure has been defined, which allows users to reserve vertex indices on each processor for future local adaptive refinement. Its parallel graph ordering routine provides orderings which are of the same quality as the ones yielded by the sequential Scotch ordering routine, while competing software ParMETIS experiences a severe loss of quality when the number of processors increase. Scotch 5.0 has been released under the CeCILL-C free/libre software license, and is also registered at APP (“*Agence pour la Protection des Programmes*”).

Scotch can be called from MUMPS, PaStiX and HIPS as an external ordering library. It is also part of the latest release of CODE_ASTER, a GPLed thermal and mechanical analysis software developed by French state-owned electricity producer EDF.

6. New Results

6.1. Numerical schemes and algorithms for fluid mechanics.

Keywords: *MHD flows, compact distribution schemes, mesh adaptation, multimaterial and multiphase flows.*

Participants: Rémi Abgrall, Cécile Dobrzynski, Robin Huart, Adam Larat, Boniface Nkonga, Mario Ricchiuto.

6.1.1. Residual distribution schemes

This year many developments have been conducted and implemented in the FluidBox software after [1] which has open up many doors.

This year we have extended the 3rd order RD scheme to Navier Stokes problems, as well as for unsteady problems. In the latter case some difficulties appear due to the non positive nature of the Lagrange basis functions. In order to overcome this problems, we have shown how to extend the method to non Lagrange basis, for example Bézier approximation. Some very first results for unsteady problems have been given by J. Triflick and confirms our expectations. The scheme has been parallelized and is scalable at least on a moderate number of procesors.

Mario Ricchiuto is conducting an analysis of the mass matrix for second order unsteady problems. The aim is to lower the number of operation by constructing an approximation of the scheme, which remains second order, but with a diagonal mass matrix.

Guillaume Baurin has started his PhD. The goal is to extend our current (3rd order) methodology to multicomponents flows for SNECMA.

6.1.2. Uncertainty quantification

In collaboration with Tim Barth (Nasa Ames Rc, USA), R. Abgrall has started to develop a strategy for computing some statistical parameters that need to be introduced because some elements of a physical model are unknown. For example, the boundary might be uncertain because of imperfections, or the inflow boundary conditions, or some parameters describing the equation of state or a turbulent model. In the approach we are working on, the main parameters will be the conditional expectancy of the fluid description (density, velocity, pressure). The approach is non intrusive. For now, some encouraging but preliminary results have been obtained for scalar hyperbolic and parabolic models.

6.1.3. Finite volume schemes for complex applications

Under this denomination, we cover relaxation schemes and MUSCL type schemes.

In relaxation schemes, one modifies the problem by adding a “relaxation” term that depends on a parameter that should be very small, if not zero. Hence it is a family of schemes that is indexed by this parameter. The advantage of this formal derivation (that can be justified theoretical in many cases), is to put the non linearities in this relaxation term, and a very simple PDE in the rest of the problem. This permits an easier analysis, and from a more practical view point, it is an avenue for developing schemes with provable properties, such as positivity of the density and the pressure.

This year, we have employed this technique for solving: multifluid flow problems [78], [75], and radiative transfert with CELIA [72], [76], [70]. In both cases, the physical stiffness (positivity of partial volume for example) could be easily resolved.

For MUSCL schemes, we have developed a numerical procedure for which we can prove that mass and pressure remain positive under a standard CFL type property. The method has been validated against many standard test cases [5], [73], [71].

6.1.4. Interface, multifluid problems and phase transition

In collaboration with CEA–CESTA, we are developing an interface tracking method using the level set method with the Ghost fluid technique. The method has been implemented in 2D and validated against several benchmark tests, in particular for interface between compressible and incompressible fluids [74], [77].

B. Braconnier has finished his thesis on the development of numerical procedures for the simulation of low mach number flows [96]. The method use relaxation solvers. The scheme, which is implicit, uses the PaStiX library, and this strategy seems to be the most efficient. In the near future, we expect to be able to consider larger problems by using the iterative method library that is developed in the team.

In the case of phase transition, V. Perrier has been studying the structure of discontinuities in the seven equations model and the five equations model using traveling wave techniques. This topic is done in collaboration with H. Guillard from the Smash project in Sophia Antipolis.

6.1.5. Parallelization and fluid mechanics

We have developed a parallel strategy by domain decomposition which is tuned for aerodynamics problems (THOT code of CEA/CESTA). The problem was to construct domain decomposition criteria which preserve the multiblock structure of the mesh. The code is fully operational and will be delivered to CEA/CESTA in the near future.

6.1.6. Mesh adaptation

C. Dobrzynski has worked on fully parallel mesh adaptation procedure that uses standard sequential mesh adaptation codes. The idea is to adapt the mesh on each processor without changing the interfaces. Then the interfaces are modified. The main advantage is the simplicity because there is no need to parallelize mesh generation tools (insert/delete, swap, etc). The main techniques are described in [82], [83].

C Dobrzynski has also developed an efficient tool for handling moving 2D and 3D meshes. Here, contrarily to most ALE methods, the connectivity of the mesh is changing in time as the objects within the computational domain are moving. The objective is to garanty a high quality mesh in term of minimum angle for example. Other criteria, which depend on the physical problem under consideration, can also been handled. Currently this meshing tool is beeing coupled with FluidBox in order to produce CFD applications. One target example is the simulation of the 3D flow over helicopter blades. A publication is in preparation see also [81].

We also have started to work on the definition of an anisotropic metric which is computed from the output of a Residual distribution code. Once this will be done, standard mesh adaptation method will be used so that the numerical error of the solution is controlled.

6.2. Multiscale schemes and algorithms

Keywords: *coupling model, crack propagation, elasticity, hierarchical methods, molecular dynamics.*

Participants: Guillaume Anciaux, Olivier Coulaud, Pierre Fortin, Jean Roman.

6.2.1. Crack simulation

A study of crack propagation in silica glasses with a coupling method between molecular dynamics and elasticity began in collaboration with the CEA Ile-de-France in December 2003. Simulations which follow crack propagation at atomistic level lead to huge number of atoms on a small domain. The coupling between two length scales allows us to treat larger domains with smaller number of atoms. Nevertheless 3D atomistic simulations involve several million atoms; they must be parallel and use a coupling with elasticity codes based on finite element approximation.

Our algorithm to couple such models is based on the Bridging method introduced by T. Belytschko. We have extended our previous work on 1D analysis of the model to higher dimension and we have developed a parallel framework to compute and visualize the coupling algorithms. This framework allows us to couple finite element technique with molecular dynamics. We validated the approach based on the Bridging Method on several multi-dimensional cases like wave propagation and crack propagation. The coupling algorithm solves a coupling linear equation and redistributes the corrections among degrees of freedom (atoms and finite elements nodes). Optimized data structures have been used in several parts of the coupling process. For example we build an efficient algorithm based on an initial computing of the finite element shape functions in order to accelerate the field's interpolation at atom positions. One other crucial service of the framework is the ability to control and forward the information on dynamic load balance strategies. Those strategies migrate atoms between processors; thus the communication scheme to update the variables attached to the coupling system (like dofs) needs to be updated. Moreover, this framework integrates EPSN that allows a powerful monitoring. All these works was presented at the IMA Workshop [60].

We are now studying the effect of the coupling method on acoustic waves. More precisely, we compared the wave reflexion rates, as a numerical results issued from simulations, with respect to various parameters such as the finite element size, the coupling zone size, etc.. We have observed that thoses rates are not constant with respect to the studied wave frequency. We have developed a spectral formulation of the coupling scheme in order to explain the reasons of such phenomenon. The numerical results as well as the spectral formulation was submitted [57].

6.3. High performance simulation for plasma and beam physics

Keywords: *ITER, MHD flows and instabilities, adaptive method, mesh refinement, numerical schemes, parallel application design.*

Participants: Rémi Abgrall, Olivier Coulaud, Robin Huart, Pascal Hénon, Guillaume Latu, Boniface Nkonga, Pierre Ramet.

6.3.1. MHD instabilities

The goal of this work is to develop and implement methods to improve the MHD simulation code in order to enable high-resolution MHD simulations of ELMs (Edge Localized Mode).

We have started to upgrade the FluidBox code so that it can run MHD flows. Starting from a "standard" Residual Distribution scheme for ideal gases, we have adapted the scheme to handle the magnetic field and its coupling to the other flow variables. This has been done via the formulation we have developed and that starts from a Rusanov like version of the first order distribution. The rest is standard once the eigen-structure of the system has been computed. The novelty is how to handle the geometrical $\vec{\text{div}} \vec{B} = 0$ constraint. For this we have used a trick which is originally due to Munz et al [95]. It consists in introducing a fictitious variable in the system (a Lagrange multiplier) and to introduce a fictitious time τ . The first step is to solve the steady modified problem in τ , and then to march in (true) time. Several validations are beeing conducted in order to validate the method. In the future, we wish to adapt the high order strategy as well as the technique with discontinuous elements we are developing in other context to reach some of the goals assigned to the ASTER project. We also have to take into account the resistivity terms.

The improvements also include adaptive mesh refinement (see section 3.2), a robust numerical MHD scheme and refinable cubic Hermite finite elements. These developments need to be consistent with the implicit time evolution scheme and the PaStiX solver [46]. The implicit scheme is essential due to the large variety of time scales in the MHD simulations.

The new methods have to be implemented and evaluated in the code FluidBox and the JOREK code to optimize the exchange of expertise on numerical methods and MHD simulations. This work is supported by the ANR CIS 2006 project called ASTER (Adaptive MHD Simulation of Tokamak Elms for iter).

6.3.2. Very scalable semi-Lagrangian Vlasov solver

One of our contribution was to propose an efficient parallel implementation of the semi-Lagrangian method using a *local* cubic spline interpolation method. Coupled with a time splitting procedure, the *standard* cubic spline interpolation is a good compromise between accuracy and simplicity. Nevertheless, the standard method does not provide the locality to find the spline coefficients. So previous semi-Lagrangian simulators used frequent global redistribution of the main data structure on the parallel machine. Thus, the scalability of the parallel application was very limited by these costly collective communications. Our local spline method [80],[39] reduces dramatically the communication overhead. It was used successfully into two codes: LOSS4D and GYSELA5D. Notably, the GYSELA code, developed by the CEA Cadarache, is the first *full-f* semi-Lagrangian code in the world [99], [85], [87], [86],[43]. Our contribution to this code has led to high scalability [94], shown on test case with up to 4096 processors. This software provides an important alternative to classical PIC and pure Eulerian simulators available in the international plasma community.

6.3.3. Adaptive Vlasov solver with low memory consumption

We have contributed to an adaptive method to solve the multi-dimensional Vlasov equation on a uniform grid. Although mathematicians have already studied wavelet transforms and adaptive numerical scheme, coupling wavelets and non-linear partial differential equations approximation represents a great deal of interest. The wavelet decomposition gives a sparse representation and a natural criterion to perform local grid refinements. The algorithmic complexity and the memory consumption of such an adaptive simulator is asymptotically less than classical dense simulator. The parallelization of such method is interesting in order to deal with applications that manipulate very large data having several dimensions. Our work focused on the parallelization with OpenMP of a semi-Lagrangian code OBIWAN4D that considers 4D data [89],[39]. We have been able to achieve simulation on a 1024^4 uniform grid with this tool. Such big grids are typically not within the scope of classical non-adaptive Vlasov simulator.

6.3.4. Dynamic load balancing for PIC solvers

Today's and future massively parallel supercomputers allow to envision the simulation of realistic problems involving complex geometries and multiple scales. One aim of the HOUPIK project is to develop massively parallel Particle-In-Cell simulator that can handle such problems. The Particle-In-Cell method consists in solving the Vlasov equation using a particle method (advancing numerically the particle thanks to motion equation). In order to achieve this goal, new numerical methods need to be investigated in the field solver part. This includes the use of hybrid grids that combine structured and unstructured mesh types. The mixing of different mesh types leads to numerical complexities and parallelization problems. A fine analysis of load balancing issues needs to be done in order to enable the scalability up to hundreds of processors. The aim of our on-going work is to compare different field solvers and investigate their efficient coupling to the particles. These self consistent relativistic PIC solvers will be the first of this kind in this context and promise to have an impact for the simulation of realistic problems in accelerator and plasma physics.

6.4. Algorithms and high-performance solvers

Keywords: *algorithms for GPGPU, direct and hybrid direct-iterative solvers, domain decomposition, graphs and irregular meshes, high performance computing, parallel fast multipole methods, parallel sparse matrix ordering, scalable parallel graph partitioning, sparse matrix ordering.*

Participants: Rached Abdelkhalek, Patrick Amestoy, Charles-Edmond Bichot, Mathieu Chanaud, Olivier Coulaud, Fabrice Dupros, Mathieu Faverge, Pierre Fortin, Jérémie Gaidamour, Abdou Guermouche, Pascal Hénon, Jun-Ho Her, Xavier Lacoste, François Pellegrini, Pierre Ramet, Jean Roman, Sébastien Fourestier.

6.4.1. *Parallel domain decomposition and sparse matrix reordering*

The work carried out within the *Scotch* project (see section 5.8) focused on two main axes.

The first one regards the parallelization of the parallel graph partitioning algorithms of *Scotch*. Basing on the parallel multi-level framework already developed in the context of parallel sparse matrix ordering [21], a set of routines for parallel graph bipartitioning has been designed. Then, a k -way parallel graph partitioning routine has been created, by transposing in parallel the recursive bisection algorithms used in the sequential *Scotch*. Parallel refinement algorithms to be used during the uncoarsening step of the multi-level method have also been investigated. Last year, a new diffusion-based refinement algorithm, called the "jug of the Danaïdes", had been designed and successfully tested in the sequential case, and showed all the more promising as it was intrinsically parallel [13]. Consequently, this year, we developed a parallel formulation of this algorithm, still in the bipartitioning case, which confirmed its efficiency as a replacement for state-of-the-art sequential local optimization algorithms such as the Fiduccia-Mattheyses algorithm.

All of these advances have been implemented in version 5.1 of *Scotch*, which was released late August [67], [68]. Although parallel recursive bipartitioning is slower than direct parallel k -way partitioning, it yields results of very good quality. On large graphs of complex shapes, the parallel graph partitioning routine of *PT-Scotch* computed cuts which were smaller on average than the ones provided by its competitor *ParMETIS*. These results have been the subject of several workshop and conference presentations [31], [30], [34], [52], and a journal paper has been submitted.

The second axis of research within the *Scotch* project regards the evaluation of the gains of static mapping, versus plain graph partitioning, on heterogeneous architectures. Indeed, with the advent of many-core and NUMA architectures, heterogeneity in speed, bandwidth and/or latency has once again to be taken into account when assigning tasks to processing elements. The advantage of *Scotch* against classical graph partitioners is that it computes static mappings, and not only partitions. The static mapping problems amounts to assigning the vertices of a source graph to vertices of a target (or architecture) graph, such that the loads of source vertices are evenly distributed with respect to the loads of the target vertices, and such that source vertices which are neighbors are assigned to target vertices which are as close to each other as possible, so as to reduce network congestion in the target architecture. In fact, graph partitioning is a subproblem of static mapping, where all target nodes have same compute power and are uniformly interconnected.

On the one hand, we are trying to extend the capabilities of *PT-Scotch* so as to offer parallel static mapping capabilities. This is an on-going work, which requires the development of new algorithms, because the approach implemented in the sequential *Scotch* possesses many unwanted synchronization bottlenecks. This is the work of Jun-Ho Her for his second year as a post-doc student. On the other hand, we want to provide users with simple means to perform static mapping and not only graph partitioning. This includes means for users to discover the underlying topology onto which they are running their jobs, and to provide them with restrictions of the global target graph describing their machine to the subset of processors they are currently using. By using such tools, designed by Sébastien Fourestier during his Master Recherche [84], and finalized during the first months of his PhD, users can expect a significant reduction of running time on heterogeneous systems, up to 25 percent in some of our experimental cases. A journal paper has been submitted which exposes these results.

A third research axis regards the design of specific graph partitioning algorithms. Several applications, such as quantum chemistry or hierarchical interface decomposition (see Section 5.6), need k -way partitions where load balance should take into account not only vertices belonging to the subdomains, but also boundary vertices, which lead to computations on each of the subdomains which share them. Some preliminary work, started last year, has been pursued by Charles-Edmond Bichot during the beginning of his post-doc, which was interrupted as he got an assistant professor position last July. This work is therefore still in progress but lacks manpower. All of these works are partially supported by the ANR-CIS project "SOLSTICE".

6.4.2. High-performance direct solvers on multi-plateforms

In order to solve linear systems of equations coming from 3D problems and with more than 50 million of unknowns, which is now a reachable challenge for new SMP supercomputers, the parallel solvers must keep good time scalability and must control memory overhead caused by the extra structures required to handle communications.

Static parallel supernodal approach. In the context of new SMP node architectures, we proposed to fully exploit shared memory advantages. A relevant approach is then to use an hybrid MPI-thread implementation. This not yet explored approach in the framework of direct solver aims at solving efficiently 3D problems with much more than 50 million of unknowns. The rationale that motivated this hybrid implementation was that the communications within a SMP node can be advantageously substituted by direct accesses to shared memory between the processors in the SMP nodes using threads. In addition, the MPI communications between processes are grouped by SMP node. We have shown that this approach allows a great reduction of the memory required for communications [92]. Many factorization algorithms are now implemented in real or complex variables, for single or double precision: LLt (Cholesky), LDLt (Crout) and LU with static pivoting (for non symmetric matrices having a symmetric pattern). This latter version is now integrated in the FluidBox software [79]. A survey article on those techniques is under preparation and will be submitted to the SIAM journal on Matrix Analysis and Applications. It will present the detailed algorithms and the most recent results. We have to add numerical pivoting technique in our processing to improve the robustness of our solver. Moreover, in collaboration with the MUMPS developers (see section 5.4), we adapt Out-of-Core techniques to overcome the physical memory constraints.

Dynamic parallel multifrontal approach.

The memory usage of sparse direct solvers can be the bottleneck to solve large-scale problems involving sparse systems of linear equations of the form $Ax=b$. If memory is not large enough to treat a given problem, disks must be used to store data that cannot fit in memory (*out-of-core* storage). In a previous work, we proposed a first out-of-core extension of a parallel multifrontal approach based on the solver MUMPS, where only the computed factors were written to disk during the factorization. We have studied in details the minimum memory requirements of this parallel multifrontal approach and proposed several mechanisms to decrease further those memory requirements. This year, we were interested by the minimization of the volume of I/O needed to perform the out-of-core factorization. Indeed, we showed in a first step that the state-of-the-art algorithms for minimizing the memory requirements of the multifrontal method are not well adapted to volume of I/O minimization. We then proposed an optimal I/O volume minimization algorithm, for the sequential multifrontal method, which reduces the volume of I/O by a factor up to two when compared with the existing memory minimization algorithms. Having this result, we are now working on the reduction of the volume of I/O when using different memory allocation schemes, that were introduced in a previous work that was concerned by the memory minimization [88] for the multifrontal method. Finally, we started a work related to the improvement of the memory scalability of the method in the context of a parallel out-of-core solver. We proposed, new scheduling techniques (either static or dynamic) to achieve a good memory “efficiency”. This is done by trying to find the best trade-off between memory-oriented techniques (i.e. sequential like depth-first traversal of the task graph) and performance oriented approaches (proportional mapping for example). This work is performed in the context of the PhD of Emmanuel Agullo, in collaboration with Jean-Yves L’Excellent (INRIA project GRAAL). Note that this PhD has been defended in november 2008.

Once the factorization is finished (i.e. factors are written to disk), we have to process the solution step which needs to read back the factors produced during the factorization. This step is very critical in the out-of-core context since it becomes a time consuming step. In addition, since the amount of computation of a single right-hand side solution step is low: the performance of the out-of-core solution step is bounded by the disk bandwidth. Thus, the objective of an out-of-core solution step is to be as close as possible to this lower bound. In the context of the thesis of Tzvetomila Slavova (PhD at CERFACS) we proposed several scheduling strategies for the out-of-core solution step that aim at ensuring that factors are accessed in a good way in the sense of I/O performance (i.e. data must be accessed in a sequential way for a given file).

Adaptation to NUMA architectures. In the context of distributed NUMA architectures, a work has recently begun, in collaboration with the INRIA RUNTIME team, to study optimization strategies, and to improve the scheduling of communications, threads and I/O. Our solvers will use *Madeleine* and *Marcel* libraries in order to provide an experimental application to validate those strategies. M. Faverge has started a Ph.D. (since october 2006) to study these aspects in the context of the NUMASIS ANR CIGC project. It has been proved that NUMA allocation can significantly improve the efficiency. In the Out-of-Core context, new problems linked to the scheduling and the management of the computational tasks may arise (processors may be slowed down by I/O operations). Thus, we have to design and study specific algorithms for this particular context (by extending our work on scheduling for heterogeneous platforms). The NUMASIS ANR CIGC project also deals with the numerical study of seismic wave propagation in geological structures, and a Ph.D. has started in 2007 (Fabrice Dupros from BRGM in collaboration with the Magique-3D INRIA project) concerning efficient parallel 3D simulations on NUMA architectures [40], [49], [62]. We will use the new version of our direct solver PaStiX adapted for these architectures for these challenge simulations. First results have been published this year [41], [50].

6.4.3. High-performance hybrid direct-iterative solvers for large sparse systems

The resolution of large sparse linear systems is often the most consuming step in scientific applications. Parallel sparse direct solver are now able to solve efficiently real-life three-dimensional problems having in the order of several million of equations, but anyway they are limited by the memory requirement. On the other hand, the iterative methods require less memory, but they often fail to solve ill-conditioned systems. We have developed two approaches in order to find some trade-off between these two classes of methods.

In these work, we consider an approach which, we hope, will bridge the gap between direct and iterative methods. The goal is to provide a method which exploits the parallel blockwise algorithmic used in the framework of the high performance sparse direct solvers. We want to extend these high-performance algorithms to develop robust parallel incomplete factorization based preconditioners for iterative methods such as GMRES or Conjugate Gradient. Those works are supported by the ANR-CIS project "SOLSTICE".

Block ILUK preconditioner. The first idea is to define an adaptive blockwise incomplete factorization that is much more accurate (and numerically more robust) than the scalar incomplete factorizations commonly used to precondition iterative solvers. Such incomplete factorization can take advantage of the latest breakthroughs in sparse direct methods and particularly should be very competitive in CPU time (effective power used from processors and good scalability) while avoiding the memory limitation encountered by direct methods. By this way, we expect to be able to solve systems in the order of hundred million of unknowns and even one billion of unknowns. Another goal is to analyse and justify the chosen parameters that can be used to define the block sparse pattern in our incomplete factorization.

The driving rationale for this study is that it is easier to incorporate incomplete factorization methods into direct solution software than it is to develop new incomplete factorizations. Our main goal at this point is to achieve a significant diminution of the memory needed to store the incomplete factors (with respect to the complete factors) while keeping enough fill-in to make the use of BLAS3 (in the factorization) and BLAS2 (in the triangular solves) primitives profitable.

In [90] and [91], we have shown the benefit of this approach over classic scalar implementation and also over direct factorisations. Indeed, on the AUDI problem (that is a reference 3D test case for direct solver with about one million of unknowns), we are able to solve the system in half the time required by the direct solver while using only one tenth of the memory needed (for a relative residual precision of 10^{-7}). We now expect to improve the convergence of our solver that fails on more difficult problems.

Recently, we have focused on the critical problem to find approximate supernodes of ILU(k) factorizations. The problem is to find a coarser block structure of the incomplete factors. The "exact" supernodes that are exhibited from the incomplete factor non zero pattern are usually very small and thus the resulting dense blocks are not large enough for an efficient use of the BLAS3 routines. A remedy to this problem is to merge supernodes that have nearly the same structure. The benefits of this approach has been shown in [24].

Hybrid direct-iterative solver based on a Schur complement approach. In recent years, a few Incomplete LU factorization techniques were developed with the goal of combining some of the features of standard ILU preconditioners with the good scalability features of multilevel methods. The key feature of these techniques is to reorder the system in order to extract parallelism in a natural way. Often a number of ideas from domain decomposition are utilized and mixed to derive parallel factorizations.

Under this framework, we developed in collaboration with Yousef Saad (University of Minnesota) algorithms that generalize the notion of “faces” and “edge” of the “wire-basket” decomposition. The interface decomposition algorithm is based on defining a “hierarchical interface structure” (HID). This decomposition consists in partitioning the set of unknowns of the interface into components called connectors that are grouped in “classes” of independent connectors [93].

In the context of robust preconditioner technique, we have developed an approach that uses the HID ordering to define a new hybrid direct-iterative solver. The principle is to build a decomposition of the adjacency matrix of the system into a set of small subdomains (the typical size of a subdomain is around a few hundreds or thousand nodes) with overlap. We build this decomposition from the nested dissection separator tree obtained using a sparse matrix reordering software as *Scotch*. Thus, at a certain level of the separator tree, the subtrees are considered as the interior of the subdomains and the union of the separators in the upper part of the elimination tree constitutes the interface between the subdomains.

The interior of these subdomains are treated by a direct method. Solving the whole system is then equivalent to solve the Schur complement system on the interface between the subdomains which has a much smaller dimension. We use the hierarchical interface decomposition (HID) to reorder and partition this system. Indeed, the HID gives a natural dense block structure of the Schur complement. Based on this partition, we define some efficient block preconditioners that allow the use of BLAS routines and a high degree of parallelism thanks to the HID properties.

We propose several algorithmic variants to solve the Schur complement system that can be adapted to the geometry of the problem: typically some strategies are more suitable for systems coming from a 2D problem discretization and others for a 3D problem; the choice of the method also depends on the numerical difficulty of the problem. We now have a parallel version of HIPS that provides full iterative methods as well as hybrid methods that mixes a direct factorization inside the domain and an iterative method in the Schur complement [33], [51], [42]. The library provide many options that allow to deal with real or complex arithmetics and symmetric or unsymmetric matrices. In particular, the very interesting feature of HIPS is that it allow to find some good trade-off between memory, robustness and time consumption in almost every cases.

Hybrid solver based on a combination of a multigrid method and a direct method. In the context of a new collaboration with the CEA/CESTA center, M. Chanaud has started this year a Ph.D. concerning a tight combination between multigrid methods and direct methods for the efficient solution of challenge 3D irregular finite element problems arising in Maxwell or Helmholtz frameworks. A sequential prototype is under validation and a parallel solver dedicated to the ODYSSEE challenge (electromagnetism) of CEA/CESTA will be implemented and validated during 2009.

6.4.4. *Parallel fast multipole method*

The Fast Multipole Method (FMM) is a hierarchical method which computes interactions for the N-body problem in $O(N)$ time for any given precision. In order to compute energy and forces on large systems, we need to improve the computation speed of the method.

This has been realized thanks to a matrix formulation of the main operator in the far field computation : this matrix formulation is indeed implemented with BLAS routines (Basic Linear Algebra Subprograms). Even if it is straightforward to use level 2 BLAS (corresponding to matrix-vector operations), the use of level 3 BLAS (that corresponds to matrix-matrix operations) is interesting because much more efficient. So, thanks to a careful data memory storage, we have rewritten the algorithm in order to use level 3 BLAS, thus greatly improving the overall runtime. Other enhancements of the Fast Multipole Method, such as the use of Fast Fourier Transform (the block FFT and the FFT with polynomial scaling), the use of rotations or the use of plane

wave expansions, allow the reduction of the theoretical operation count. Comparison tests have shown that, for the required precisions in astrophysics or in molecular dynamics, our approach is either faster (compared to rotations and plane waves) or as fast and without any numerical instabilities (compared to the FFT based method), hence justifying our BLAS approach. These results have been accepted for publication in the Journal of Computational Physics [22].

Our BLAS version has then been extended to non uniform distributions, requiring therefore a new octree data structure named octree with indirection, that is efficient for both uniform and non uniform distributions. We have also designed an efficient algorithm that detects uniform areas in structured non uniform distributions, since these areas are more suitable for BLAS computations. These results have been submitted for publication in [61]. An efficient parallel code of our BLAS version, based on an hybrid MPI-thread programming, has finally been developed and validated on shared and distributed memory architectures.

Thanks to a post-doctoral position within the ScAIApplix project, located in the *Laboratoire d'Astrophysique de Marseille* and supported by the *HALOBAR* ANR project (ANR-06-BLAN-0172), we have been able to compare our FMM code to other serial and parallel codes dedicated to astrophysical comparisons. This will result in a forthcoming submission to an international journal, as a research note or as a full paper.

6.4.5. Domain decomposition to solve neutron transport equations

This work started with a collaboration between the EDF/SINETICS team and the ScAIApplix project to design and develop techniques to optimize the efficiency of the codes used to simulate the physics of nuclear reactors. G. Caramel was recruited, during one year, as an associate engineer, to work on this contract.

This collaboration has given rise to a new study for a large-scale parallel computing code simulation. Bruno Lathuilère started in November 2006 a thesis on domain decomposition methods applied for solving neutron transport equations. Many scientific fields are involved like reactor physics, numerical analysis and computer science. For a stationary problem (critically), one has to solve a generalized eigenvalue problem using the power algorithm preconditioned by three recursive levels of deteriorated iterative solvers.

At the end of the second year of his thesis, Bruno has been able to implement, in an industrial code, algorithms he has proposed and the initial results, obtained on a the simulation of the IAEA benchmark, have been published this year [48], [47], [37].

However, this approach to parallelize the SP_n approximation, used to solve the Boltzmann equation, will not be sufficient to obtain a complete simulation. We therefore propose in a second step to improve the efficiency of parallel code from this first work by taking into account the heterogeneity of the reactor (that increases the number of energy groups and the order of anisotropy). Mesh refinement techniques and the coupling of different models should be adapted to neutron simulations.

6.4.6. High performance algorithms for depth imaging

In the context of a new collaboration with TOTAL, Rached Abdelkhalek has started a Ph.D. concerning massive parallelism and new algorithms for taking advantage of large clusters with additional GPU devices for seismic depth imaging problems. We work on data distribution, thread scheduling on a large number of computational cores and for the efficient joint use of CPUs and GPUs. Promising results have been achieved for the Reverse Time Migration algorithm and a first paper will be submitted to ICS'2009.

6.5. Computational steering environment for distributed numerical simulations

Keywords: *code coupling, computational steering, data redistribution, numerical simulation, scientific visualization.*

Participants: Olivier Coulaud, Aurélien Esnard, Nicolas Richart, Mathieu Souchaud, Guilhem Caramel.

Redistribution of hierarchical grid. In the context of the ANR MASSIM, we are now considering more complex data structures such as *hierarchical grids* for the purpose of parallel data redistribution. In previous works, we have proposed a description model of distributed data, that can represent scalar parameters, structured grids, points or unstructured meshes. We have recently extended this model with hierarchical grids, that are designed in the RedGRID framework as a coarse grid (by inheritance of classical grid) with multiple recursive subgrids. We have validated our model with the LOSS4D simulation that uses a compression scheme based on hierarchical finite element basis. Concerning the redistribution of hierarchical grids, we have proposed two different approaches based on previous algorithms already available in the RedGRID framework. We distinguish two main approaches: the spatial approach and the placement approach. In the first approach, one uses spatial information on the coarse grid to compute the pieces of blocks (with refined cells) that should be exchanged through the network. In the the placement approach, one assigns a weight to each cell of the coarse grid that depends on the amount of data in sub-levels. And then, one uses this information to calculate the placement of blocks from one code to another with a strategy inspired from load-balancing algorithms. These algorithms have been implemented in the RedGRID framework and integrated to the EPSN steering environment.

Model for the steering of parallel-distributed simulations. The model that we have proposed in the EPSN framework can only steer efficiently SPMD simulations. A natural evolving is to consider more complex simulations such as coupled SPMD codes called M-SPMD (Multiple SPMD like multiscale simulation for “crack-propagation”) and client/server simulation codes. In order to steer these kinds of simulation, we have designed an extension to the Hierarchical Task Model (HTM), that affords to solve the coherency problem for such complex applications. This model has been published in [15]. The EPSN framework has been extended to handle this new kind of simulations. In the context of the ANR MASSIM, we have recently validated our works with a multi-scale simulation for “crack-propagation” (LibMultiScale). In this case-study, EPSN is able to pause/resume the whole coupled simulation, to coherently get and visualize the complex distributed data: a distributed unstructured mesh at the continuum scale, mixed with distributed atoms at the atomic scale.

Compression and Visualisation of 4D data. The main output of our most recent Vlasov simulators is typically a 4D+t data. Thanks to the recent extension in the EPSN framework, one can now perform the steering and visualization of the 4D data during simulations. In order to reach this goal, we have plugged a compression scheme (that uses Hierarchical Finite element basis) in the LOSS4D semi-Lagrangian code. Indeed, the full 4D data structure is most of the time too large to be sent quickly on the network to the visualization process. So, we provide at each time step of the simulator a compressed hierarchical data structure that can be sent to a remote visualization tool, called PlasmaViz (owned by the INRIA CALVI project). To perform efficient online visualization, PlasmaViz can directly manipulate hierarchical data structure, extract slices, and plot them without decompression. We have also imported the 4D compression algorithm into the EPSN framework, in order to transfer compressed data even if the simulation uses dense data structure like the standard version of Loss4D. Finally, we expect to steer another simulation code called Yoda4D that natively uses hierarchical data structure.

EPSN framework. Many improvements have been brought in the software both in the EPSNkernel and in the client part. Firstly, low layers of RedGRID have been deeply modified and more precisely the ColCOWS layer. The communication algorithms have been optimized by using a hierarchical scheme. Currently, ColCOWS can deploy a distributed application on a cluster of 256 nodes. It takes now less than 1 second, against more than 4 seconds before. Secondly, we have developed QT4 plugins to help the integration of EPSN in QT application (client) such as SiMonE (Simulation Monitoring for EPSN), PlasmaViz (used in MASSIM Project) and ParaView. Finally, we have merged both the RedGRID and EPSNframeworks into a single platform publicly available on INRIA Gforge (<http://gforge.inria.fr/projects/epsn>).

7. Contracts and Grants with Industry

7.1. Research and development contracts

CEA research and development contracts:

- Relaxation method for Laser/Plasma interaction: Ti-Te model (Aubin Bellue, Boniface Nkonga);
- Numerical tools for interfaces flows: compressible/incompressible (Marie Billaud, Boniface Nkonga);
- Feasability study of the use of the graph reorderer and partitioner PT-Scotch and of the hybrid direct-iterative solver HIPS for large 3D electromagnetism problems (J er mie Gaidamour, Pascal H enon, Fran ois Pellegrini);
- Conception of an hybrid solver combining multigrid and direct methods (Mathieu Chanaud, David Goudin, Jean-Jacques Pesqu e, Jean Roman).

EDF research and development contract:

- Application of a domain decomposition method to the neutronic SPn equations (Bruno Lathuili re, Pierre Ramet, Jean Roman).

TOTAL research and development contract:

- Massive parallelism and use of GPU devices for seismic depth imaging problems (Rached Abdelkhalek, Olivier Coulaud, Guillaume Latu, Jean Roman).

8. Other Grants and Activities

8.1. National initiatives

8.1.1. *AEROCAP: Experimental and numerical study of the aeroacoustic in cylindrical cavities*

Participant: R emi Abgrall.

Grant: SNECMA

Dates: 2006-2009

Partners: Ecole Centrale Lyon, ONERA-DSNA, ENSAM, Universit  Aix-Marseille

Overview: The AEROCAP project goal is to study the noise produced by the circulation of air around elliptic or cylindrical cavities. This kind of noises are particularly intense around aircraft wing at take-off and landing phase. Our task is to analyse new schemes that are high order accurate and can be used on general unstructured meshes. This is done within the framework of residual distribution schemes.

8.1.2. *ASTER: Adaptive MHD Simulation of Tokamak Elms for ITER*

Participants: R emi Abgrall, Olivier Coulaud, Robin Huart, Pascal H enon, Boniface Nkonga, Pierre Ramet.

Grant: ANR-06-CIS

Dates: 2006 – 2009

Partners: CEA Cadarache.

Overview: The magnetohydrodynamic instability called ELM for Edge Localized Mode is commonly observed in the standard tokamak operating scenario. The energy losses the ELM will induce in ITER plasmas are a real concern. However, the current understanding of what sets the size of these ELM induced energy losses is extremely limited. No numerical simulations of the complete ELM instability, from its onset through its non-linear phase and its decay, exist in literature. Recently, encouraging results on the simulation of an ELM cycle have been obtained with the JOREK code developed at CEA but at reduced toroidal resolution. The JOREK code uses a fully implicit time evolution scheme in conjunction with the PaStiX sparse matrix library. In this project it is proposed to develop and implement methods to improve the MHD simulation code to enable high-resolution MHD simulations of ELMs. The ELM simulations are urgently needed to improve our understanding of ELMs and to evaluate possible mechanism to control the energy losses. The improvements include adaptive mesh refinement, a robust numerical MHD scheme and refinable cubic Hermite finite elements. These developments need to be consistent with the implicit time evolution scheme and the PaStiX solver. The implicit scheme is essential due to the large variety of time scales in the MHD simulations. The new methods will be implemented and evaluated in the code FluidBox, developed by the ScAlApplix team and the JOREK code to optimize the exchange of expertise on numerical methods and MHD simulations.

The project is a collaboration between the Departement de Recherche sur la Fusion Controlee (DRFC, CEA/Cadarache) and the Laboratoire Bordelais de Recherche en Informatique (LaBRI) and Mathématiques Appliquées de Bordeaux (IMB) at the University of Bordeaux I.

Web: <http://aster.gforge.inria.fr/>

8.1.3. *HALOBAR: Dark matter and the dynamical evolution of barred galaxies*

Participants: Olivier Coulaud, Pierre Fortin.

Grant: ANR - Programme blanc

Dates: 2007 – 2010

Partners: Laboratoire d'Astrophysique de Marseille (leader of the project), Astroparticules et Cosmologie

Overview: We want to study the secular evolution of barred galaxies in the presence of both a gaseous component and a live and responsive halo. Although considerable steps have been made in understanding the two influences independently, little, if anything, has been done about understanding their coupled effect and the outcome of the internal competition between angular momentum emitters and absorbers when all partners are active. This is partly due to the fact that, so far, understanding one of the two influences was already a sufficient challenge. It is also partly due to the fact that an all-round approach is necessary at this stage. Simulations should be more performing, necessitating deep knowledge of N-body and hydrodynamic codes, as well as of the ways to link them. The effects of the resonances need to be fully understood, necessitating expertise in analytical work on resonances, but also in orbital structure theory including periodic orbits as well as chaotic ones. Software for the analysis of the resonances and the study of the orbital structure in the simulations is necessary. Finally, links with observations need to be very tight, since they will be the only way to ensure that the solutions found are realistic and relevant to real galaxies, rather than abstract mathematical models.

8.1.4. *MASSIM: development of a software environment for interactively processing and visualizing complex large scale data coming from numerical simulations*

Participants: Olivier Coulaud, Aurélien Esnard, Guillaume Latu, Nicolas Richart.

Grant: ANR MMSA - ARA MAsses de données

Dates: 2005 – 2008

Partners: IRMA (Strasbourg, UMR 7001), LSIIT (leader of the project, Strasbourg, UMR 7005)

Overview: Numerical simulation is a continuously growing area, especially with the increasing computational power of current computer technology, thus covering larger and larger scientific application fields. But at these days, monitoring tools are still seriously lacking, since developers and users desire more and more to get faster and faster feedbacks of the simulation results. In this project, we are interested in large scale simulations dealing with complex data (multivariate and multidimensional). Our aim is to realize a plate-form / framework to couple parallel and distributed simulations, like in GRID'5000, with an interactive monitoring and visualization system. This plate-form will be validated on two types of large scale applications: plasma and crack propagation simulation using multiscale approaches. For these applications, the simulation codes are definitely very complex and need some highly efficient tools to represent the large amount of data, to redistribute the data using visualization and to control and validate the corresponding computation algorithms. Since results may be multivariate and multidimensional, they need also specific data exploration and visualization tools.

Web: <http://dpt-info.u-strasbg.fr/~dischler/massim/massim.html>

8.1.5. *NOSSI: New platform for parallel, hybrid quantum/classical simulations*

Participants: Olivier Coulaud, Aurélien Esnard, Damien Genet, Jean Roman.

Grant: ANR 2007 – CIS

Dates: 2008 – 2010

Partners: CPMOH (Bordeaux, UMR 5098), DRIMM, IMPREM (leader of the project, Pau, UMR 5254), Institut Néel (Grenoble, UPR2940)

Overview: Physicists, chemists and computer scientists join forces in this project to further high performance numerical simulation of materials, by developing and deploying a new platform for parallel, hybrid quantum/classical simulations. The platform synthesizes established functions and performances of two major European codes, SIESTA and DL_POLY, with new techniques for the calculation of the excited states of materials, and a graphical user interface allowing steering, visualization and analysis of running, complex, parallel computer simulations.

The platform couples a novel, fast TDDFT (Time dependent density functional theory) route for calculating electronic spectra with electronic structure and molecular dynamics methods particularly well suited to simulation of the solid state and interfaces.

The software will be capable of calculating the electronic spectra of localized excited states in solids and at interfaces. Applications of the platform include hybrid organic-inorganic materials for sustainable development, such as photovoltaic materials, bio- and environmental sensors, photocatalytic decontamination of indoor air and stable, non-toxic pigments.

Web: <http://nossi.gforge.inria.fr/index.html>

8.1.6. *NUMASIS: adaptation and optimization of the applicative performance on NUMA architectures*

Participants: Mathieu Faverge, Abdou Guermouche, Pierre Ramet, Jean Roman.

Grant: ANR-05-CIGC-002

Dates: 2006 – 2009

Partners: Bull, Total, BRGM, CEA, ID-Imag (leader of the project), PARIS (IRISA), Runtime (INRIA Futurs Bordeaux).

Overview: The multiprocessor machines of tomorrow will rely on an NUMA architecture introducing multiple levels of hierarchy into computers (multimodules, chips multibody, multithreading material, etc). To exploit these architectures, parallel applications must use powerful runtime supports making possible the distribution of execution and data streams without compromising their portability. Project NUMASIS proposes to evaluate the functionalities provided by the current systems, to apprehend the limitations, to design and implement new mechanisms of management of the processes, data and communications within the basic softwares (operating system, middleware, libraries). The target algorithmic tools that we retained are parallel linear sparse solvers with application to seismology.

Web: <http://numasis.gforge.inria.fr/>

8.1.7. *PETAL*

Participants: Pascal Hénon, François Pellegrini, Pierre Ramet.

Grant: ANR Cosinus 2008

Dates: 2009-2010

Partners: INRIA Saclay-Ile de France (leader of the project), Paris 6, IFP (Rueil-Malmaison), CEA Saclay

Overview: In this collaborative effort, we propose to develop parallel preconditioning techniques for the emergent hierarchical models of clusters of multi-core processors, as used for example in future petascale machines. The preconditioning techniques are based on recent progress obtained in combining the well known incomplete LU (ILU) factorization with the tangential filtering, another incomplete factorization where a filtering condition is satisfied. The goal of this project is to transform these preconditioners into black box parallel preconditioners that could be as usable as standard and popular methods such as ILU. For this, we address several issues related to the quality of the combined preconditioner. We also aim to make the connection of these methods with the domain decomposition methods. To obtain a preconditioner suitable for parallelism, we will study associated graph partitioning and reordering techniques.

8.1.8. *SOLSTICE: High performance solvers and simulations*

Participants: Patrick Amestoy, Jérémie Gaidamour, Abdou Guermouche, Pascal Hénon, François Pellegrini, Pierre Ramet, Jean Roman.

Grant: ANR-06-CIS

Dates: 2006 – 2009

Partners: CERFACS, EADS IW, EDF R&D SINETICS, INRIA Rhône-Alpes and LIP, INPT/IRIT, CEA/CESTA, CNRS/GAME/CNRM.

Overview: New advances in high-performance numerical simulation require the continuing development of new algorithms and numerical methods. These technologies must then be implemented and integrated into real-life parallel simulation codes in order to address critical applications that are at the frontier of our know-how. The solution of sparse systems of linear equations of (very) large size is one of the most critical computational kernel in terms of both memory and time requirements. Three-dimensional partial differential equations (3D-PDE) are particularly concerned by the availability of efficient sparse linear algorithms since the numerical simulation process often leads to linear systems of 10 to 100 million variables that need to be solved many times. In a competitive environment where numerical simulation becomes extremely critical compared to physical experimentation, very precise models involving a very accurate discretisation are more and more critical. The objective of our project is thus both to design and develop high-performance parallel linear solvers that will be efficient to solve complex multiphysic and multiscale problems of very large size. To demonstrate the impact of our research, the work produced in the project will be integrated in real simulation codes to perform simulations that could not be considered with today's technologies.

Web: <http://solstice.gforge.inria.fr/>

8.1.9. SYNERGHY: Simulate the beam, gravity and hydrodynamic phenomenon

Participants: Aubin Bellue, Boniface Nkonga.

Grant: ANR-06-CIS6-009

Dates: 2006 – 2009

Partners: CEA/DSM(leader of the project), CELIA.

Overview: This research project aims at satisfying simulation needs in the fields of Astrophysics, Hot Dense Matter and Inertial Confinement Fusion. A large part of the scientific production in these fields rely upon simulations of complex unsteady hydrodynamic flows, coupled to non equilibrium transport and chemical kinetics. As the characteristic time scales of transport may be much shorter than the fluid time scale, implicit numerical tools are required. Some of the numerical codes developed in the presently co working institutions (CEA/DSM, INRIA, CELIA) may be now regarded as mature. Nevertheless, large scale exploitation of these codes present issues the resolution of which exceed the isolated capabilities of the developing teams. A coordinated action is thus necessary. These issues are :

- Numerical methods for multimaterial, three dimensional compressible, unsteady flows.
- Parallel algorithms. Specifically, efficient implicit solvers for non linear transport or diffusion equations.
- Data management and visualisation.
- Constitution and efficient access to large shared data bases for equations of state and transport coefficients.
- Source code management and software design in order to ensure: Cross validation of codes, Stability and long term maintenance of software, Interoperability (portability of physical packages, exchange of numerical data, code linking).

Web: <http://www-dapnia.cea.fr/Projets/SINERGHY/>

8.2. International initiatives

8.2.1. Associated team PHyLeaS

Participants: Olivier Coulaud, Jérémie Gaidamour, Pascal Hénon, Pierre Ramet, Jean Roman.

Grant: INRIA

Dates: 2008-2010.

Partners: University of Minnesota, INRIA Sophia-Antipolis Méditerranée, Institute of Computational Mathematics Brunswick, LIMA-IRIT (UMR CNRS 5505).

Overview: New advances in high performance scientific computing require continuing the development of innovative algorithmic and numerical techniques, their efficient implementation on modern massively parallel computing platforms and their integration in application software in order to perform large-scale numerical simulations currently out of reach. The solution of sparse linear systems is a basic kernel which appears in many academic and industrial applications based on partial differential equations (PDEs) modeling physical phenomena of various nature. In most of the applications, this basic kernel is used many times (numerical optimization procedure, implicit time integration scheme, etc.) and often accounts for the larger part of the computing time. In a competitive environment where the numerical simulation tends to replace the experiment, the modeling calls for PDEs of ever increasing complexity. Furthermore, realistic applications involve multiple space and time scales, and non-trivial geometrical features. In this context, a common trend is to discretize the underlying PDE models using arbitrary high-order finite element methods designed on unstructured grids. As a consequence, the resulting algebraic systems are irregularly structured and very large in size. The aim of this project is the design and efficient implementation of parallel hybrid linear system solvers which combine the robustness of direct methods with the implementation flexibility of iterative schemes. These approaches are candidate to get scalable solvers on massively parallel computers.

Web: <http://www.labri.fr/projet/scalapplix/phyleas.html>

8.2.2. *ADIGMA: a european project on the development of adaptive higher order variational methods for aerospace applications*

Participant: Rémi Abgrall.

Grant: European Commission

Dates: 2006-2009

Partners: AIRBUS F et AIRBUS D, DASSAULT, ALENIA, DLR, ONERA, NLR, ARA, VKI, INRIA, Nankin University, Universities of Stuttgart, Bergamo, Twente, Nottingham, Swansea, Charles (Prague), Varsovie, CENAERO, ENSAM Paris)

Overview: Computational Fluid Dynamics is a key enabler for meeting the strategic goals of future air transportation. However, the limitations of today numerical tools reduce the scope of innovation in aircraft development, keeping aircraft design at a conservative level. Within the 3rd Call of the 6th European Research Framework Programme, the strategic target research project ADIGMA has been initiated. The goal of ADIGMA is the development and utilization of innovative adaptive higher-order methods for the compressible flow equations enabling reliable, mesh independent numerical solutions for large-scale aerodynamic applications in aircraft design. A critical assessment of the newly developed methods for industrial aerodynamic applications will allow the identification of the best numerical strategies for integration as major building blocks for the next generation of industrial flow solvers. In order to meet the ambitious objectives, a partnership of 22 organizations from universities, research organizations and aerospace industry from 10 countries with well proven expertise in CFD has been set up guaranteeing high level research work with a clear path to industrial exploitation.

Web: http://www.dlr.de/as/en/Desktopdefault.aspx/tabid-2035/2979_read-4582/

8.2.3. *France-Berkeley Fund Award*

Participants: Patrick Amestoy, Abdou Guermouche.

Grant: France-Berkeley Fund

Dates: 2007-2009.

Partners: CERFACS, ENSEEIHT-IRIT and GRAAL (INRIA project)

Overview: In the framework of the France-Berkeley Fund, we have been awarded a research grant to enable an exchange program involving both young and confirmed scientists. The collaboration will focus on massively parallel solvers for large sparse matrices and will reinforce the collaboration initiated by Emmanuel Agullo. On the French side, this project also involves I. Duff (CERFACS) and J. L'Excellent (GRAAL, INRIA project).

8.2.4. Collaboration project with JAEA (Japan)

Participant: Pierre Ramet.

Partners: ENSEEIHT-IRIT and GRAAL (INRIA project)

Overview: The JAEA (Japan Atomic Energy Agency) develops its own solvers in sparse linear algebra for its applications in numerical simulations. The collaboration, initiated by Prof. M. Dayde (ENSEEIHT-IRIT), started by the organization of the REDIMPS workshop at Tokyo (may 2007), mainly around the TLSE platform and the DIET middleware. However, during this workshop, researchers from JAEA have highlighted the need to compare their approaches with the solvers we are developing.

9. Dissemination

9.1. Participation to the Scientific Community

Rémi Abgrall is scientific associate editor of the international journals “Mathematical Modelling and Numerical Analysis”, “Computer and Fluids”, “Journal of Computational Physics”, “Journal of Scientific Computing” and “Journal of Computing Science and Mathematics”. He is member of the scientific committee of the international conference ICCFD. He is member of the CFD committee of ECOMAS. He is also member of the scientific committee of CERFACS and that of the ANR “Intensive Computation and Simulation” theme. He is member of the Comité National des Universités, section 26. He is member of the board of the GAMNI group of SMAI.

Olivier Coulaud is member of the scientific committee of the international conference VECPAR’08 and member of the thematic CP6 committee for IDRIS, CINES and CCRT.

François Pellegrini has been member for four years, until this year, of the program committee of the High Performance and Grid Computing (HPGC) conference.

Pierre Ramet is member of the committee of Researchers at CINES for topic number 6 (mathematics).

Jean Roman is president of the Project Committee of INRIA Bordeaux - Sud-Ouest and member of the National Evaluation Committee of INRIA. He has been member of the scientific committee of the international conferences ADVCOMP08, EuroMicroPDP08 (IEEE), PMAA08 (SIAM) and of the national conference Renpar’08. He is member of the “Strategic Comity for Intensive Computation” of the French Research Ministry.

9.2. Teaching

In complement of the normal teaching activity of the university members and of ENSEIRB members, Olivier Coulaud and Pascal Hénon teach at ENSEIRB (computer science engineering school).

Mathieu Faverge, Jérémie Gaidamour, Abdou Guermouche, Pascal Hénon and Pierre Ramet have given lectures in the CNRS “autumn school on scientific computing” held in Sète (France) from september 29th to october 3rd 2008.

Mario Ricchiuto has given lectures in the “Mastère Spécialisé en Ingénierie Aéronautique et Spatiale” organised by ENSAM, MatMéca, ENSEIRB, Institu de Cognitique and several local industrial partners. He also teaches at MatMéca.

10. Bibliography

Major publications by the team in recent years

- [1] R. ABGRALL. *Essentially non-oscillatory residual distribution schemes for hyperbolic problems*, in "J. Comp. Phys.", vol. 214, n^o 2, 2006, p. 773–808.

- [2] R. ABGRALL, M. MEZINE. *Construction of second order accurate monotone and stable residual distribution schemes for unsteady flow problems.*, in "J. Comput. Phys.", vol. 188, n^o 1, 2003, p. 16-55.
- [3] E. AGULLO, A. GUERMOUCHE, J.-Y. L'EXCELLENT. *A Parallel Out-of-core Multifrontal Method: Storage of Factors on Disk and Analysis of Models for an Out-of-core Active Memory*, in "Parallel Computing, Special Issue on Parallel Matrix Algorithms", vol. 34, n^o 6-8, 2008, p. 296-317.
- [4] G. ANCIAUX, O. COULAUD, J. ROMAN. *High Performance Multiscale Simulation for Crack Propagation*, in "Proceedings of HPSEC'2006", vol. ICPPW'06, IEEE Computer Society Press, August 2006, p. 473-480.
- [5] C. BERTHON. *Robustness of MUSCL schemes for 2D unstructured meshes*, in "J. Comp. Phys.", vol. 218, 2006, p. 495-509.
- [6] N. BESSE, G. LATU, A. GHIZZO, E. SONNENDRÜCKER, P. BERTRAND. *A wavelet-MRA-based adaptive semi-Lagrangian method for the relativistic Vlasov-Maxwell system*, in "Journal of Computational Physics", vol. 227, n^o 16, august 2008, p. 7889–7916.
- [7] O. COULAUD, P. FORTIN, J. ROMAN. *High performance BLAS formulation of the multipole-to-local operator in the Fast Multipole Method*, in "Journal of Computational Physics", vol. 227, 2008, p. 1836-1862, <http://hal.inria.fr/inria-00000957/en/>.
- [8] H. DECONINCK, M. RICCHIUTO. *19*, in "Residual Distribution schemes : foundations and analysis.", John Wiley and Sons, Ltd, october 2007.
- [9] C. DOBRZYNSKI, P. FREY, B. MOHAMMADI, O. PIRONNEAU. *Fast and accurate simulations of air-cooled structures*, in "Comput. Methods Appli. Mech. Engrg.", vol. 195, 2006, p. 3168-3180.
- [10] A. ESNARD, N. RICHART, O. COULAUD. *A Steering Environment for Online Parallel Visualization of Legacy Parallel Simulations*, in "Proceedings of the 10th International Symposium on Distributed Simulation and Real-Time Applications (DS-RT 2006), Torremolinos, Malaga, Spain", vol. ds-rt, IEEE Press, October 2006, p. 7-14.
- [11] J. GAIDAMOUR, P. HÉNON. *A parallel direct/iterative solver based on a Schur complement approach.*, in "IEEE 11th International Conference on Computational Science and Engineering, Sao Paulo, Brazil", july 2008, p. 98–105.
- [12] P. HÉNON, P. RAMET, J. ROMAN. *On finding approximate supernodes for an efficient ILU(k) factorization*, in "Parallel Computing", vol. 34, 2008, p. 345–362.
- [13] F. PELLEGRINI. *A parallelisable multi-level banded diffusion scheme for computing balanced partitions with smooth boundaries*, in "Proc. Euro-Par, Rennes", LNCS, vol. 4641, August 2007, p. 191–200.
- [14] M. RICCHIUTO, N. VILLEDIEU, H. DECONINCK, R. ABGRALL. *On uniformly high-order accurate residual distribution schemes for advection-diffusion*, in "Journal of Computational Applied Mathematics", vol. 215, 2008, p. 547-556, <http://hal.inria.fr/inria-00333998/en/>.

- [15] N. RICHART, A. ESNARD, O. COULAUD. *Toward a Computational Steering Environment for Legacy Coupled Simulations*, in "Proceedings of 6th International Symposium on Parallel and Distributed Computing (ISPDC 2007), Hagenberg, Austria", IEEE Press, July 2007, p. 319–326.

Year Publications

Articles in International Peer-Reviewed Journal

- [16] R. ABGRALL, A. LARAT, M. RICCHIUTO, C. TAVÉ. *A simple construction of very high order non oscillatory compact schemes on unstructured meshes.*, in "Computers and Fluids", 2008, <http://hal.inria.fr/inria-00333756/en/>.
- [17] R. ABGRALL, C. SHU. *Development of residual distribution schemes for discontinuous Galerkin method: the scalar case with linear elements*, in "Communications in Computational Physics", to appear, vol. 5, 2009, p. 376–390.
- [18] E. AGULLO, A. GUERMOUCHE, J.-Y. L'EXCELLENT. *A Parallel Out-of-core Multifrontal Method: Storage of Factors on Disk and Analysis of Models for an Out-of-core Active Memory*, in "Parallel Computing, Special Issue on Parallel Matrix Algorithms", vol. 34, n^o 6-8, 2008, p. 296-317.
- [19] P. AMESTOY, I. DUFF, A. GUERMOUCHE, T. SLAVOVA. *Analysis of the Out-of-Core Solution Phase of a Parallel Multifrontal Approach*, in "Parallel Computing", To appear, 2008.
- [20] N. BESSE, G. LATU, A. GHIZZO, E. SONNENDRÜCKER, P. BERTRAND. *A wavelet-MRA-based adaptive semi-Lagrangian method for the relativistic Vlasov-Maxwell system*, in "Journal of Computational Physics", vol. 227, n^o 16, august 2008, p. 7889–7916.
- [21] C. CHEVALIER, F. PELLEGRINI. *PT-Scotch: A tool for efficient parallel graph ordering*, in "Parallel Computing", vol. 34, n^o 6–8, 2008, p. 318–331.
- [22] O. COULAUD, P. FORTIN, J. ROMAN. *High performance BLAS formulation of the multipole-to-local operator in the Fast Multipole Method*, in "Journal of Computational Physics", vol. 227, 2008, p. 1836-1862, <http://hal.inria.fr/inria-00000957/en/>.
- [23] S. GENAUD, P. GANCARSKI, G. LATU, D. BLANSCHOT. *Exploitation of a parallel clustering algorithm on commodity hardware with P2P-MPI*, in "Journal of Supercomputing", doi:10.1007/s11227-007-0136-2, vol. 43, n^o 1, january 2008, p. 21–41.
- [24] P. HÉNON, P. RAMET, J. ROMAN. *On finding approximate supernodes for an efficient ILU(k) factorization*, in "Parallel Computing", vol. 34, 2008, p. 345–362.
- [25] M. RICCHIUTO, A. BOLLERMANN. *Stabilized residual distribution for shallow water simulations*, in "J. Comput. Phys.", doi:10.1016/j.jcp.2008.10.020, 2008.
- [26] M. RICCHIUTO, N. VILLEDIEU, H. DECONINCK, R. ABGRALL. *On uniformly high-order accurate residual distribution schemes for advection-diffusion*, in "Journal of Computational Applied Mathematics", vol. 215, 2008, p. 547-556, <http://hal.inria.fr/inria-00333998/en/>.

[27] R. SAUREL, F. PETITPAS, R. ABGRALL. *Modelling phase transition in metastable liquids: Application to cavitating and flashing flows.*, in "Journal of Fluid Mechanics", vol. 607, 2008, p. 313-350, <http://hal.inria.fr/inria-00333908/en/>.

[28] N. VILLEDIEU, T. QUINTINO, R. ABGRALL, H. DECONINCK. *High-order residual distribution schemes on quadrilateral meshes.*, in "Int. J. Numer. Methods Fluids", 2008, <http://hal.inria.fr/inria-00334003/en/>.

Articles in National Peer-Reviewed Journal

[29] A. ESNARD. *RedGRID: un environnement pour la redistribution d'objets complexes*, in "Technique et Science Informatiques (TSI)", vol. 27, n^o 3-4, 2008, p. 427-455.

Invited Conferences

[30] C. CHEVALIER, J.-H. HER, F. PELLEGRINI. *Efficient and scalable parallel graph partitioning*, in "37th SpeedUp workshop on high-performance computing, ETH Zürich", September 2008.

[31] C. CHEVALIER, J.-H. HER, F. PELLEGRINI. *The PT-Scotch project: purpose, algorithms, current results*, in "IBM seminar, Zürich", June 2008.

[32] J. GAIDAMOUR, P. HÉNON. *HIPS : a parallel hybrid direct/iterative solver based on a Schur complement approach*, in "Sparse Days at CERFACS, Workshop of VECPAR 08, France", Invited presentation, June 2008.

[33] J. GAIDAMOUR, P. HÉNON, J. ROMAN, Y. SAAD. *Parallel resolution of sparse linear systems by mixing direct and iterative methods*, in "International Symposium on Iterative Methods in Scientific Computing (IMACS), Lille, France", invited presentation, march 2008.

[34] J.-H. HER, F. PELLEGRINI. *Towards efficient and scalable graph partitioning methods*, in "Proc. CSC'2008, Atlanta", March 2008.

International Peer-Reviewed Conference/Proceedings

[35] R. ABGRALL, S. KARNI. *A relaxation scheme for the two-layer shallow water system.*, in "Hyperbolic problems. Theory, numerics and applications Proceedings of the 11th international conference on hyperbolic problems., France Lyon", Springer, Benzoni-Gavage, Sylvie (ed.) et al, 2008, <http://hal.inria.fr/inria-00334007/en/>.

[36] E. AGULLO, A. GUERMOUCHE, J.-Y. L'EXCELLENT. *On the I/O volume in Out-of-Core Multifrontal Methods with a Flexible Allocation Scheme*, in "VECPAR'08 International Meeting on High Performance Computing for Computational Science", To appear, 2008.

[37] M. BARRAULT, B. LATHUILIÈRE, P. RAMET, J. ROMAN. *A domain decomposition method applied to the simplified transport equations*, in "IEEE 11th International Conference on Computational Science and Engineering, Sao Paulo, Brazil", July 2008, p. 91-97.

[38] Y. CANIOU, J.-S. GAY, P. RAMET. *Tunable parallel experiments in a GridRPC framework: application to linear solvers*, in "To appear in LNCS 5336 proceedings of VECPAR'08, Toulouse, France", June 2008.

[39] N. CROUSEILLES, M. GUTNIC, G. LATU, E. SONNENDRÜCKER. *Comparison of Two Eulerian solvers for the four dimensional Vlasov equation (Parts I and II)*, in "Proc. of the 2nd international conference of

Vlasovia", *Communications in Nonlinear Science and Numerical Simulation*, vol. 13(1), february 2008, p. 88-99.

- [40] F. DUPROS, H. AOCHI, A. DUCELLIER, D. KOMATITSCH, J. ROMAN. *Exploiting intensive multithreading for the efficient simulation of 3D seismic wave propagation.*, in "IEEE 11th International Conference on Computational Science and Engineering, Sao Paulo, Brazil", july 2008, p. 253–260.
- [41] M. FAVERGE, P. RAMET. *Dynamic Scheduling for sparse direct Solver on NUMA architectures*, in "LNCS proceedings of PARA'2008, Trondheim, Norway", to appear, May 2008.
- [42] J. GAIDAMOUR, P. HÉNON. *A parallel direct/iterative solver based on a Schur complement approach.*, in "IEEE 11th International Conference on Computational Science and Engineering, Sao Paulo, Brazil", july 2008, p. 98–105.
- [43] V. GRANDGIRARD, Y. SARAZIN, X. GARBET, G. DIF-PRALADIER, P. GHENDRIH, N. CROUSEILLES, G. LATU, E. SONNENDRÜCKER, N. BESSE. *Computing ITG turbulence with a full-f semi-Lagrangian code*, in "Proc. of the 2nd international conference of Vlasovia", *Communications in Nonlinear Science and Numerical Simulation*, vol. 13(1), Elsevier, february 2008, p. 81-87.
- [44] M. RICCHIUTO, R. ABGRALL. *Stable and convergent residual distribution for time-dependent conservation laws*, in "ICCFD4, 4th International Conference on Computational Fluid Dynamics", to appear, Springer Verlag, 2008.
- [45] N. VILLEDIEU-LIGOUT, M. RICCHIUTO, H. DECONINCK. *High-order residual distribution : discontinuity capturing crosswind dissipation and diffusion*, in "ICCFD4, 4th International Conference on Computational Fluid Dynamics", to appear, Springer Verlag, 2008.

Workshops without Proceedings

- [46] R. ABGRALL, O. COULAUD, P. HÉNON, R. HUART, G. HUYSMANS, G. LATU, B. NKONGA, S. PAMELA, P. RAMET. *Numerical simulation of tokamak plasmas*, in "7th PAMIR International Conference on Fundamental and Applied MHD, Presqu'île de Giens, France", September 2008.
- [47] M. BARRAULT, B. LATHUILLIÈRE, P. RAMET, J. ROMAN. *A Domain Decomposition Method Applied to Large Eigenvalue Problems in Neutron Physics*, in "Proceedings of PMAA'2008, Neuchâtel, Suisse", June 2008.
- [48] M. BARRAULT, B. LATHUILLIÈRE, P. RAMET, J. ROMAN. *A domain decomposition method for the resolution of an eigenvalue problem in neutron physics*, in "International Symposium on Iterative Methods in Scientific Computing (IMACS), Lille, France", March 2008.
- [49] F. DUPROS, F. DE MARTIN, E. FOERSTER, D. KOMATITSCH, J. ROMAN. *High-performance finite element simulations of seismic wave propagation in 3D inelastic media*, in "Proceedings of PMAA'2008, Neuchâtel, Suisse", june 2008.
- [50] M. FAVERGE, X. LACOSTE, P. RAMET. *A NUMA Aware Scheduler for a Parallel Sparse Direct Solver*, in "Proceedings of PMAA'2008, Neuchâtel, Suisse", June 2008.
- [51] J. GAIDAMOUR, P. HÉNON. *HIPS : a parallel hybrid direct/iterative solver based on a Schur complement approach*, in "Proceedings of PMAA'2008, Neuchâtel, Suisse", june 2008.

- [52] J.-H. HER, F. PELLEGRINI. *Efficient and scalable parallel graph partitioning*, in "Proc. PMAA'08, Neuchâtel", June 2008.
- [53] G. HUYSMANS, R. ABGRALL, M. BECOULET, R. HUART, B. NKONGA, S. PAMELA, P. RAMET. *Non-Linear MHD code development for ELM simulations, Hersonissos, Crete, Greece*, in "35th EPS Plasma Physics Conference, poster session", June 2008.
- [54] M. RICCHIUTO, A. BOLLERMANN. *Accuracy of stabilized residual distribution for shallow water flows including dry beds*, in "HYP2008, 12th International Conference on Hyperbolic Problems : Theory, Numerics, Applications. Maryland (USA)", June 2008.

Research Reports

- [55] R. ABGRALL. *A simple, flexible and generic deterministic approach to uncertainty quantifications in non linear problems: application to fluid flow problems*, Rapport de recherche, INRIA, 2008, <http://hal.inria.fr/inria-00325315/en/>.
- [56] R. ABGRALL. *Une méthode de type distributive employant des éléments discontinus pour le calcul d'écoulements avec choc*, RR-6439, Rapport de recherche, INRIA, 2008, <http://hal.inria.fr/inria-00218209/en/>.
- [57] G. ANCIAUX, O. COULAUD, J. ROMAN, G. ZERAH. *Ghost force reduction and spectral analysis of the 1D bridging method*, RR-6582, submitted to JCP, Rapport de recherche, INRIA, 2008, <http://hal.inria.fr/inria-00300603/en/>.

Other Publications

- [58] R. ABGRALL. *Construction of simple stable and convergent high order schemes for steady first order Hamilton Jacobi equations*, in revision, 2008.
- [59] E. AGULLO, A. GUERMOUCHE, J.-Y. L'EXCELLENT. *On Reducing the I/O Volume in Sparse Out-of-core Multifrontal Methods*, in revision, 2008.
- [60] G. ANCIAUX, O. COULAUD, J. ROMAN. *Parallel multiscale simulation for crack propagation*, november 2008, IMA Workshop: Development and analysis of multiscale methods, poster session.
- [61] O. COULAUD, P. FORTIN, J. ROMAN. *High-performance BLAS formulation of the Adaptive Fast Multipole Method*, Submitted, Elsevier, 2008.
- [62] F. DUPROS, F. DE MARTIN, E. FOERSTER, D. KOMATITSCH, J. ROMAN. *High-performance finite element simulations of seismic wave propagation in 3D non linear geological media*, Submitted, 2008.
- [63] S. FOURESTIER, F. PELLEGRINI, F. LEDOUX, J.-C. WEILL. *Algorithms and tools for static mapping on NUMA architectures*, Submitted, 2008.
- [64] J. GAIDAMOUR, P. HÉNON. *A hybrid direct/iterative solver based on a Schur complement approach.*, in "Mini-Workshop on parallel iterative solvers and domain decomposition techniques, Minneapolis, USA", July 2008.

- [65] J.-H. HER, F. PELLEGRINI. *Efficient and scalable parallel graph partitioning*, Submitted, 2008.
- [66] P. HÉNON, P. RAMET, J. ROMAN. *A supernode amalgamation algorithm for an efficient block incomplete factorization*, in "Mini-Workshop on parallel iterative solvers and domain decomposition techniques, Minneapolis, USA", July 2008.
- [67] F. PELLEGRINI. *PT-SCOTCH 5.1 User's guide*, August 2008.
- [68] F. PELLEGRINI. *SCOTCH 5.1 User's Guide*, August 2008.
- [69] M. RICCHIUTO, A. BOLLERMANN. *Stabilized Residual Distribution for Shallow Water simulations.*, in press (on line), 2008.

References in notes

- [70] C. BERTHON. *A relaxation scheme to approximate the 10-moment Gaussian closure*, in "European Conference on Computational Fluid dynamics, ECCOMAS CFD 2006", P. WESSELING, E. ONATE, J. PERIAUX (editors), 2006.
- [71] C. BERTHON. *A stable gradient reconstruction for the MUSCL schemes applied to systems of conservation laws*, in "Flow, Turbulence and Combustion, Special Issue: Numerical Simulation of Complex and Multiphase Flows", vol. 76, 2006, p. 343-351.
- [72] C. BERTHON. *Numerical approximations of the 10-moment Gaussian closure*, in "Math. Comp.", vol. 75, 2006, p. 1809-1831.
- [73] C. BERTHON. *Why the MUSCL-Hancock scheme is L^1 -stable*, in "Numer. Math.", vol. 104, 2006, p. 27-46.
- [74] C. BERTHON, B. BRACONNIER, J. CLAUDEL, B. NKONGA. *A relaxation method for a two phase flow with surface tension*, in "Numerical mathematics and advanced applications", Springer Verlag, 2006, p. 843-851.
- [75] C. BERTHON, M. BREUSS, M.-O. TITEUX. *A relaxation scheme for pressureless gas dynamics*, in "Num. Method Partial Diff. Equations", vol. 22, 2006, p. 484-505.
- [76] C. BERTHON, B. NKONGA. *Behavior of the finite volumes schemes in material and numerical interfaces*, in "Third International Symposium on Finite Volumes for Complex Applications", Hermes Science Publications, 2002, p. 139-146.
- [77] C. BERTHON, B. NKONGA. *Numerical model of a compressible multi-fluid fluctuating flow*, in "Inter. J. finite volumes", 2005.
- [78] C. BERTHON, B. NKONGA. *Multifluid numerical approximations based on multipressure formulation*, in "Comp. Fluids", vol. 36, 2006, p. 467-479.
- [79] B. BRACONNIER, B. NKONGA, M. PAPIN, P. RAMET, M. RICCHIUTO, J. ROMAN, R. ABGRALL. *Efficient solution technique for low Mach number compressible multiphase problems*, in "Proceedings of PMAA'2006, Rennes, France", 2006.

- [80] N. CROUSEILLES, G. LATU, E. SONNENDRÜCKER. *Hermite Spline Interpolation on Patches for a parallel implementation of beam focusing problems*, in "Proceedings of HIF 2006", Nuclear instruments and Methods in Physics Research A., vol. 557(1), Elsevier, July 2007, p. 129-132.
- [81] C. DOBRZYNSKI. *3D local anisotropic remeshing for rigid bodies displacement*, October 2007, http://www-rocq.inria.fr/who/Frederic.Alauzet/tetra_eng.html, Second Workshop on Grid Generation for Numerical Computations, INRIA Rocquencourt.
- [82] C. DOBRZYNSKI, P. FREY, B. MOHAMMADI, O. PIRONNEAU. *Fast and accurate simulations of air-cooled structures*, in "Comput. Methods Appl. Mech. Engrg.", vol. 195, 2006, p. 3168-3180.
- [83] C. DOBRZYNSKI, P. FREY, O. PIRONNEAU. *Couplage et adaptation de maillage anisotrope pour des simulations de flux d'air dans des géométries complexes*, in "Revue européenne de mécanique numérique", vol. 16, 2007, p. 749-773.
- [84] S. FOURESTIER. *Placement statique de graphes sur architecture NUMA*, Masters thesis, Université Bordeaux 1, Jun 2008.
- [85] X. GARBET, Y. SARAZIN, V. GRANDGIRARD, G. DIF-PRALADIER, G. DARMET, P. GHENDRIH, P. BERTRAND, N. BESSE, E. GRAVIER, P. MOREL, E. SONNENDRÜCKER, N. CROUSEILLES, J.-M. DISCHLER, G. LATU, E. VIOLARD, M. BRUNETTI, S. BRUNNER, X. LAPILLONNE, T.-M. TRAN, L. VILLARD. *Beyond scale separation in gyrokinetic turbulence*, in "21st IAEA Fusion Energy Conference, Chengdu, China", October 2006.
- [86] V. GRANDGIRARD, Y. SARAZIN, P. ANGELINO, B. ALBERTO, N. CROUSEILLES, G. DARMET, G. DIF-PRALADIER, X. GARBET, P. GHENDRIH, S. JOLLIET, G. LATU, L. VILLARD, E. SONNENDRÜCKER. *Global full-f gyrokinetic simulations of plasma turbulence*, in "Plasma physics and controlled fusion", doi:10.1088/0741-3335/49/12B/S16, vol. 49, December 2007, p. 173-182.
- [87] V. GRANDGIRARD, Y. SARAZIN, X. GARBET, G. DIF-PRALADIER, P. GHENDRIH, N. CROUSEILLES, G. LATU, E. SONNENDRÜCKER, N. BESSE, P. BERTRAND. *GYSELA, a full-f global gyrokinetic Semi-Lagrangian code for ITG turbulence simulations*, in "Proceedings of Theory of Fusion Plasmas, Varenna", 2006.
- [88] A. GUERMOUCHE, J.-Y. L'EXCELLENT. *Constructing Memory-minimizing Schedules for Multifrontal Methods*, in "ACM Transactions on Mathematical Software", vol. 32, n^o 1, 2006, p. 17-32.
- [89] M. GUTNIC, G. LATU, E. SONNENDRÜCKER. *Adaptive two-dimensional Vlasov simulations of heavy ion beams*, in "Proceedings of HIF 2006", Nuclear instruments and Methods in Physics Research A., vol. 557(1), Elsevier, February 2007, p. 125-128.
- [90] P. HÉNON, F. PELLEGRINI, P. RAMET, J. ROMAN, Y. SAAD. *Applying parallel direct solver skills to build robust and highly performant preconditioners*, in "PARA'04, Workshop on state-of-the-art in scientific computing, Copenhagen, Denmark", June 2004.
- [91] P. HÉNON, P. RAMET, J. ROMAN. *A Blockwise Algorithm for Parallel Incomplete Cholesky Factorization*, in "Proceedings of PMAA'2004, Marseille, France", October 2004.

-
- [92] P. HÉNON, P. RAMET, J. ROMAN. *On using an hybrid MPI-Thread programming for the implementation of a parallel sparse direct solver on a network of SMP nodes*, in "Proceedings of Sixth International Conference on Parallel Processing and Applied Mathematics, Workshop HPC Linear Algebra, Poznan, Pologne," , September 2005.
- [93] P. HÉNON, Y. SAAD. *A Parallel Multistage ILU Factorization based on a Hierarchical Graph Decomposition*, in "SIAM J. Sci. Comput.", vol. 28, 2006, p. 2266–2293.
- [94] G. LATU, N. CROUSEILLES, V. GRANDGIRARD, E. SONNENDRÜCKER. *Gyrokinetic Semi-Lagrangian Parallel Simulation using a Hybrid OpenMP/MPI Programming*, in "Recent Advances in PVM and MPI", LNCS, vol. 4757, Springer, 2007, p. 356-364.
- [95] C.-D. MUNZ, P. OMNES, R. SCHNEIDER, E. SONNENDRÜCKER, U. VOSS. *Divergence correction techniques for Maxwell solvers based on a hyperbolic model*, in "J. Comput. Phys.", vol. 161, n^o 2, 2000, p. 484-511.
- [96] B. NKONGA. *Superconsistency property for nonconservative operator approximation : application to compressible multifluid flow.*, in "Proceedings of The 2nd International Conference on Scientific Computing and Partial Differential Equations and The First East Asia SIAM Symposium, Hong Kong.", 2006.
- [97] F. PELLEGRINI. *PT-SCOTCH 5.0 User's guide*, August 2007.
- [98] F. PELLEGRINI. *SCOTCH 5.0 User's Guide*, August 2007.
- [99] Y. SARAZIN, V. GRANDGIRARD, G. DIF-PRALADIER, E. FLEURANCE, X. GARBET, P. GHENDRIH, P. BERTRAND, N. BESSE, N. CROUSEILLES, E. SONNENDRÜCKER, G. LATU, E. VIOLARD. *Impact of large scale flows on turbulent transport*, in "Plasma Phys. Control Fusion", vol. 48, december 2006, p. B179-B188.