



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Team Bacchus

*Parallel tools for Numerical Analysis and
Resolution of essentially Hyperbolic
problems*

Bordeaux - Sud-Ouest

Theme : Stochastic Methods and Models

Activity
R *eport*

2009

Table of contents

1. Team	1
2. Overall Objectives	1
2.1. Overall Objectives	1
2.2. Highlights	2
3. Scientific Foundations	2
3.1. Numerical schemes for fluid mechanics	2
3.2. Algorithms and high-performance solvers	3
3.2.1. High-performance direct solvers for distributed clusters	3
3.2.2. High-performance iterative and hybrid direct/iterative solvers	3
3.2.3. Meshes and graph partitioning	5
3.2.3.1. Parallel graph partitioning and static mapping	5
3.2.3.2. Adaptive dynamic mesh partitioning	5
4. Application Domains	6
4.1. Panorama	6
4.2. Fluid mechanics	6
4.2.1. Steady transonic and supersonic flows	6
4.2.2. Unsteady transonic flows	6
4.3. High performance simulation dedicated to ITER project	6
5. Software	7
5.1. Introduction	7
5.2. FluidBox	7
5.3. PaStiX	7
5.4. HIPS	8
5.5. Scotch	8
5.6. MMG3D	9
5.7. Montjoie	9
6. New Results	10
6.1. Numerical schemes and algorithms for fluid mechanics.	10
6.1.1. Residual distribution schemes	10
6.1.2. Uncertainty quantification	10
6.1.3. Discontinuous Galerkin schemes, New elements in DG schemes	10
6.1.4. Mesh adaptation	11
6.2. High performance simulation for plasma physics	11
6.3. Algorithms and high-performance solvers	12
6.3.1. Parallel domain decomposition and sparse matrix reordering	12
6.3.2. High-performance direct solvers on multi-platforms	12
6.3.3. Hybrid direct-iterative solver based on a Schur complement approach.	13
6.3.4. MURGE a common interface to sparse linear solvers.	13
7. Contracts and Grants with Industry	13
7.1. EDF	13
7.2. SNECMA	14
7.3. DASSAULT	14
8. Other Grants and Activities	14
8.1. National initiatives	14
8.1.1. AEROCav: Experimental and numerical study of the aeroacoustic in cylindrical cavities	14
8.1.2. ASTER: Adaptive MHD Simulation of Tokamak Elms for iteR	14
8.1.3. NUMASIS: adaptation and optimization of the applicative performance on NUMA architectures	15
8.1.4. PETAL: PETascaling ALgorithms for preconditioning used in scientific applications	15

8.1.5.	SOLSTICE: High performance solvers and simulations	15
8.1.6.	MACAO	16
8.2.	Actions Funded by the EC	16
8.2.1.	ADIGMA: a European project on the development of adaptive higher order variational methods for aerospace applications	16
8.2.2.	ADDECCO : ADaptive schemes for DEterministic and stoChastiC Flow PrOblems	17
9.	Dissemination	17
9.1.	Participation to the Scientific Community	17
9.2.	Teaching	18
10.	Bibliography	18

BACCHUS is an joint INRIA Team with University of Bordeaux (UB and IPB) and CNRS (IMB, UMR 5251 and LaBRI, UMR 5800), created on January 1st 2009.

1. Team

Research Scientist

Pascal Hénon [Research Associate]

Mario Ricchiuto [Research Associate]

Faculty Member

Rémi Abgrall [Team Leader, INRIA, Research Director (on secondment), Professor at IPB, HdR]

Cécile Dobrzynski [Assistant Professor, IPB]

Marc Duruflé [Assistant Professor, IPB]

Pierre Ramet [Assistant Professor, University Bordeaux 1]

François Pellegrini [Assistant Professor, IPB, HdR]

Technical Staff

Pascal Jacq [INRIA, Expert Engineer, funding from world competitiveness cluster AESE, MACAO Project]

Xavier Lacoste [INRIA, Associate Engineer]

PhD Student

Guillaume Baurin [SNECMA, CIFRE, since Oct. 2008]

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]

Jérémie 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]

Cedric Lachat [INRIA, CORDIS S]

Arnaud Krust [INRIA, funding from ADECCO]

Algiane Froehly [INRIA, funding from ADECCO]

Bruno Lathuilière [EDF, CIFRE, since Nov. 2006]

Jiri Treflik [INRIA, funding from AEROCV and CTU, Prag, since Oct. 2008, up to November 1st 2009]

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

Post-Doctoral Fellow

Jun-Ho Her [INRIA, since Sep. 2007, now visiting from National Research Foundation of Korea]

Visiting Scientist

ZJ Wang [Professor, Iowa State University, Iowa, 2 weeks, july 2009]

Smadar Karni [Professor, University of Michigan, Ann Arbor, one week, september 2009]

Administrative Assistant

Josy Baron [Secretary INRIA]

2. Overall Objectives

2.1. Overall Objectives

BACCHUS is a joint team of INRIA Bordeaux - Sud-Ouest, LaBRI (Laboratoire Bordelais de Recherche en Informatique – CNRS UMR 5800, University of Bordeaux and IPB) and IMB (Institut Mathématique de Bordeaux – CNRS UMR 5251, University of Bordeaux). BACCHUS has been created on the first of January, 2009 (<http://bacchus.bordeaux.inria.fr>).

The purpose of the BACCHUS project is to analyze and solve efficiently 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 involving both applied mathematics and computer science.

Our major focus is fluid problems, and especially the simulation of *physical wave propagation problems* including fluid mechanics, inert and reactive flows, multimaterial and multiphase flows, acoustic, etc. BACCHUS intends to solve these problems by bringing contributions to all steps of the development chain that goes from the design of new high-performance, more robust and more precise, numerical schemes, to the creation and implementation of optimized parallel algorithms and high-performance codes.

By taking into account architectural and performance concerns from the early stages of design and implementation, the high-performance software which will implement our numerical schemes will be able to run efficiently on most of today's major parallel computing platforms (UMA and NUMA machines, large networks of SMP nodes, production GRIDs).

2.2. Highlights

- 2009 has been the starting year of the ERC Advanced Grant ADDECCO. Two PhDs and one engineer have been hired.
- We have defined a common API to easily call any sparse linear solvers from a parallel or a sequential code. This project is open and we have released the interface specification as well as documentation and code examples at (<http://murge.gforge.inria.fr>). This interface has been implemented in HIPS and PaStiX.
- We participated in the aeronautic "Salon du Bourget" in June 2009.

3. Scientific Foundations

3.1. Numerical schemes for fluid mechanics

Participants: Rémi Abgrall, Marc Duruflé, Mario Ricchiuto.

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 compute many flow realizations, but the quality of the results is far from being optimal in many cases. 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. Consequently, 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 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, though some progress have been made recently so that this is not anymore a constraint. We are also interested in Discontinuous Galerkin type schemes. However, the compactness of the scheme enables us 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, which 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 which are now difficult to compute thanks to the know-how coming from our research on compact distribution schemes and the daily discussions with specialists in computer science and scientific computing. These problems range from aeroacoustic to multiphysic problems, such as the ones 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 which is described in section 5.2. This work is supported by the EU-Strep ADIGMA, various research contracts and in part by the ANR-CIS ASTER project (see section 4.3 also), and also by the ERC grant ADDECCO.

3.2. Algorithms and high-performance solvers

Participants: Cécile Dobrzynski, Pascal Hénon, François Pellegrini, Pierre Ramet.

3.2.1. 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 BACCHUS 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 the computation and communication performance of which are predictable in advance. This has led to software developments (see sections 5.3, 5.5)

3.2.2. High-performance iterative and hybrid direct/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 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 CPU power and are difficult to parallelize. Our work aims at finding some unknown orderings and partitioning 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 fully exploit the capabilities of modern high-performance computers.

In this context, we study two approaches.

- 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 this approach, we focus 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 have been shown in [37]. These algorithms are implemented in the PaStiX library.

- The second approach is based on the 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 [38].

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 sub-domains (the typical size of a sub-domain 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 sub-trees are considered as the interior of the sub-domains and the union of the separators in the upper part of the elimination tree constitutes the interface between the sub-domains.

The interior of these sub-domains are treated by a direct method. Solving the whole system is then equivalent to solve the Schur complement system on the interface between the sub-domains 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 discretisation and others for a 3D problem; the choice of the method also depends on the numerical difficulty of the problem. In the HIPS library, we provide full iterative methods (very low memory consumption) as well as hybrid methods that mixes a direct factorization inside the domain and an iterative method in the Schur complement [34]. The library provides many options that allow one to deal with real or complex arithmetic, and symmetric or unsymmetric matrices. In particular, the very interesting feature of HIPS is that it allows one to find some good trade-off between memory, robustness and time consumption in almost every case.

Details can be found in sections 5.4 and 6.3.

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

3.2.3. Meshes and graph partitioning

3.2.3.1. Parallel graph partitioning and static mapping

Finding vertex separators for sparse matrix ordering is only one of the many uses of generic graph partitioning tools. For instance, finding balanced and compact domains in problem graphs is essential to the efficiency of parallel iterative solvers. Here again, because of the size of the problems at stake, parallel graph partitioning tools are mandatory to provide good load balance and minimal communication cost.

The execution of parallel applications implies communication between processes executed on the different cores. On NUMA architectures which are strongly heterogeneous in terms of latency and capacity, communication cost strongly depends on the repartition of tasks among cores. Architecture-aware load balancing must take into account both the characteristics of the parallel applications (including for instance task processing costs and the amount of communication between tasks) and the topology of the target architecture (providing the powers of cores and the costs of communication between all of them). When processes are assumed to coexist simultaneously for all the duration of the program, this optimization problem is called mapping. A mapping is called static if it is computed prior to the execution of the program and is never modified at run-time.

The sequential Scotch tool was able to perform static mapping since its first version, but this feature was not widely known nor used by the community. With the increasing need to map very large problem graphs onto very large and strongly heterogeneous parallel machines (whether hierarchical NUMA clusters or GPU-based systems), there is an increasing demand for parallel static mapping tools.

3.2.3.2. Adaptive dynamic mesh partitioning

Many simulations which model the evolution of a given phenomenon along with time (turbulence and unsteady flows, for instance) need to re-mesh some portions of the problem graph in order to capture more accurately the properties of the phenomenon in areas of interest. This re-meshing is performed according to criteria which are closely linked to the undergoing computation and can involve large mesh modifications: while elements are created in critical areas, some may be merged in areas where the phenomenon is no longer critical.

Performing such re-meshing in parallel creates additional problems. In particular, splitting an element which is located on the frontier between several processors is not an easy task, because deciding when splitting some element, and defining the direction along which to split it so as to preserve numerical stability most, require shared knowledge which is not available in distributed memory architectures. Ad-hoc data structures and algorithms have to be devised so as to achieve these goals without resorting to extra communication and synchronization which would impact the running speed of the simulation.

Most of the works on parallel mesh adaptation attempt to parallelize in some way all the mesh operations: edge swap, edge split, point insertion, etc. It implies deep modifications in the (re)mesher and often leads to bad performance in term of CPU time. An other work [30] proposes to base the parallel re-meshing on

existing mesher and load balancing to be able to modify the elements located on the frontier between several processors.

In addition, the preservation of load balance in the re-meshed simulation requires dynamic redistribution of mesh data across processing elements. Several dynamic repartitioning methods have been proposed in the literature [40], [39], which rely on diffusion-like algorithms and the solving of flow problems to minimize the amount of data to be exchanged between processors. However, integrating such algorithms into a global framework for handling adaptive meshes in parallel has yet to be done.

4. Application Domains

4.1. Panorama

The main objective of the BACCHUS 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 software 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 multiphysics, and which require 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

Participants: Rémi Abgrall [Corresponding member], Marc Duruflé, Mario Ricchiuto.

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. The schemes developed in for fluid mechanics problems use `Scotch`, `HIPS` and `PaStiX` when the type of problems and the CPU requirements make this useful.

4.2.1. Steady transonic and supersonic flows

One of our application fields is the one of steady subsonic, transonic and supersonic flow problems when the equation of state is the one of standard air. This class of physical problems corresponds to “standard” aerodynamics and the models are those of the Euler equations and the Navier Stokes ones, possibly with turbulent effects. Here we consider the residual distribution and SUPG schemes.

4.2.2. Unsteady transonic flows

Another field of application is the one of *unsteady* problems with the same physics, or in the case of the linearized Euler equations. The schemes we develop are the Residual distribution schemes 5.2 and Discontinuous Galerkin schemes 5.7. Specific modifications, with respect to their steady counter parts, are done in order to reduce dramatically the computational time.

4.3. High performance simulation dedicated to ITER project

Participants: Rémi Abgrall, Pierre Ramet [Corresponding member].

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.

We have established a 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 magneto-hydrodynamic 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 within 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 and HIPS). The second kind of software corresponds to dedicated software for fluid mechanics (FluidBox).

For these parallel software developments, we use the message passing paradigm (basing on the MPI interface), sometimes combined with threads so as to exploit multi-core architectures at their best: in some computation kernels such as solvers, when processing elements reside on the same compute node, message buffer space can be saved because the aggregation of partial results can be performed directly in the memory of the receiving processing element. Memory savings can be tremendous, and help us achieve problem sizes which could not be reached before (see Section 5.3).

5.2. FluidBox

Participant: Rémi Abgrall [corresponding member].

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. A partitioning tool exists in the package, which 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.

5.3. PaStiX

Participant: Pierre Ramet [corresponding member].

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 5.2). 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 5.5). 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 performance.

Another important point is that our study is suitable for any heterogeneous parallel/distributed architecture when its performance is predictable, such as clusters of SMP nodes. In particular, we now offer a high performance version with a low memory overhead for SMP node architectures, which fully exploits the advantage of shared memory 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 in defining 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.

5.4. HIPS

Participant: Pascal Hénon [corresponding member].

HIPS (Hierarchical Iterative Parallel Solver) is a scientific library that provides an efficient parallel iterative solver for very large sparse linear systems. The key point 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 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 sub-domains) is solved using an iterative method (GMRES or Conjugate Gradient at the time being). This special ordering and partitioning allows for 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. The code provides a hybrid method which blends direct and iterative solvers. HIPS exploits the partitioning and multistage ILU techniques (see [38]) to enable a highly parallel scheme where several subdomains can be assigned to the same process. It also provides a scalar preconditioner based on the multistage ILUT factorization.

HIPS can be used as a standalone 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. Thus, HIPS is a software library that provides several methods to build an efficient preconditioner in almost all situations.

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

5.5. Scotch

Participant: François Pellegrini [corresponding member].

Scotch (<http://www.labri.fr/~pelegri/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 is regaining 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 factored matrix in parallel.

Version 5.0 of Scotch, released on August 2007, 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 was released under the CeCILL-C free/libre software license, and has been registered at APP (“*Agence pour la Protection des Programmes*”).

Version 5.1 of Scotch, released on September 2008, extended the parallel features of PT-Scotch, which can now compute graph partitions in parallel by means of a parallel recursive bipartitioning framework.

Scotch has been integrated in numerous third-party software, which indirectly contribute to its diffusion. For instance, it is used by the ZOLTAN module of the TRILINOS software (SANDIA Labs), by CODE_ASTER LIBRE, a GPLed thermal and mechanical analysis software developed by French state-owned electricity producer EDF, by the parallel solvers MUMPS (ENSEEITH/IRIT, LIP and LaBRI), SuperLUDist (U.C. Berkeley), PaStiX (LaBRI) and HIPS (LaBRI), as well as by several other scientific computing software.

5.6. MMG3D

Participant: Cécile Dobrzynski [corresponding member].

MMG3D is a tetrahedrally automatic remesher. Starting with a tetrahedral mesh, it produces quasi-uniform meshes with respect to a metric tensor field. This tensor prescribes a length and a direction for the edges, thus the resulting meshes will be anisotropic meshes. The software is based on local mesh modifications and an anisotropic version of Delaunay kernel is implemented to insert vertex on the mesh. Moreover, MMG3D allows to deal with rigid body motion and moving meshes. When a displacement is prescribed on a part of the boundary, a final mesh is generated such as the surface points will be moved according this displacement. MMG3D is used in particular in GAMMA for their mesh adaptation developments, but also at EPFL (math department), Dassault Aviation, Lemma (a french SME), etc. More details can be found on <http://www.math.u-bordeaux1.fr/~dobj/logiciels/mmg3d.php>.

5.7. Montjoie

Participant: Marc Duruflé [corresponding member].

Montjoie is a finite element code initially handling only quadrilateral/hexahedral elements. Because of the tensorization of these elements, efficient algorithms can be written for the computation of finite element matrices. It can now also handle prisms and pyramids.

6. New Results

6.1. Numerical schemes and algorithms for fluid mechanics.

Participants: Rémi Abgrall [Corresponding member], Guillaume Baurin, Cécile Dobrzynski, Marc Duruflé, Algyane Froehly, Robin Huart, Arnaud Krust, Pascal Jacq, Cédric Lachat, Luc Mieussens (IMB/IPB), Pierre-Henri Maire (CELIA/CEA), Adam Larat, Mario Ricchiuto, Jirka Treffilick.

6.1.1. Residual distribution schemes

This year, many developments have been conducted and implemented in the FluidBox software after [28] which has opened up many doors.

We have extended the 3rd order RD scheme to Navier Stokes problems, as well as to unsteady problems. Pascal Jacq and Cédric Lachat have extended the communication scheme in FluidBox so as to handle high order schemes.

Some difficulties have appeared due to the non positive nature of the Lagrange basis functions, in particular for unsteady problems. In order to overcome this problem, we have shown how to extend the method to non Lagrange basis, for example by means of Bézier approximation. Some very first results for unsteady problems have been given by J. Treflik and confirm our expectations.

Mario Ricchiuto is conducting an analysis of the mass matrix for second order unsteady problems. The aim is to lower the number of operations by constructing an approximation of the scheme, which remains second order, but with a diagonal mass matrix. This should provide a much more efficient method than the one which was used before.

Guillaume Baurin has started his PhD. The goal is to extend our current (3rd order) methodology to multicomponent flows for SNECMA. Arnaud Krust has started his PhD, and, with G. Baurin, is examining several algorithmic solutions to the discretisation of the viscous terms vis RD scheme. Algyane Froehly has started her PhD and is studying other than Lagrange elements in the context of RD schemes for compressible fluid flow problems.

Adam Larat has finished and defended his PhD on high order RD schemes with the first applications to 3D and Navier Stokes applications. This has been done in the context of the ADIGMA project.

Algyane Froehly has started her PhD wicth topic is the conception of RD schemes using non lagrange elements like Bezier elements or NURBS. Arnaud Krust has started his PhD in studying the approximation of the navier Stokes equations using solution-dependant elements.

Mario Ricchiuto and Luc Mieussens have started to study how to combine Residual Distribution schemes and Asymptotic Preserving scheme methodologies.

6.1.2. Uncertainty quantification

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. Discontinuous Galerkin schemes, New elements in DG schemes

The earlier versions of Montjoie could only handle hex and quadrangles. However the generation of purely hexahedral meshes is still challenging, that's why we have studied, in collaboration with Morgane Bergot, finite elements methods able to handle hybrid meshes, including hexahedra, prisms, pyramids and tetrahedra. Contrary to the other elements, finite element space for pyramids is non-polynomial, and in [29], the optimal finite element space is discussed for pyramids with an exhaustive comparison of all previous works about this issue.

These elements have been implemented in Montjoie and applied to wave equation and time-domain Maxwell equations. The results showed the advantage to use hybrid meshes when no nice hexahedral mesh was available. These results have been conducted for continuous finite elements and with discontinuous Galerkin method as well.

Rémi Abgrall and Pierre-Henri Maire, with François Vilar (PhD at CELIA funded by a CEA grant started in october 2009), have started to work on Lagrangian schemes within the Discontinuous Galerkin schemes.

6.1.4. 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, after which interfaces are modified. The main advantage is simplicity, because there is no need to parallelize mesh generation tools (insert/delete, swap, etc). The main techniques are described in [31], [32].

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 guaranty a high quality mesh in term of minimum angle for example. Other criteria, which depend on the physical problem under consideration, can also be handled. Currently this meshing tool is being coupled with FluidBox in order to produce CFD applications. One target example is the simulation of the 3D flow over helicopter blades.

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.

Moreover, a work on high order mesh generation has begun. We are modifying the classic mesh operators to take into account the curve edges. Beginning with a derefined valid curve mesh, we would be able to generate an uniform refined curve mesh and also to adapt the mesh density in certain region (boundary layer).

6.2. High performance simulation for plasma physics

Participants: Rémi Abgrall, Robin Huart, Pascal Hénon, Pierre Ramet [Corresponding member], Hocine Sellama.

We have been involved in two tasks : in the first one we work on novel numerical schemes for solving the compressible resistive MHD equations. In the second one, in connecting with the JOREK code developed in CEA Cadarache, we work on adaptive mesh refinement problems and their connection with the solution of large linear systems to be solved in parallel. This has led to two publications [14], [21].

The aim of our work in the ASTER project is to provide an efficient numerical method for solving the MHD equations, more especially in the form they are used for the ITER model. Here we want to improve the ability of Residual Distribution schemes to solve this hyperbolic system. Once fixed the choice of the full set of equations and the behavior of physical parameters, and with a validated numerical solver for these equations, we should be able to simulate plasma instabilities like those encountered in the ITER tokamak configuration. The step to ELMs simulations would then be achieved. This is a global view of the context, and it may be seen as a framework. However, one should notice that, contrarily to the work on JOREK at the CEA Cadarache, our purpose is a more general and academic code (FluidBox).

Due to the localized nature of the ELMs at the boundary of the plasma the use of mesh refinement is ideally suited to minimize the number of elements required for a given accuracy. The high resolution is only required where large gradients develop which is on a surface which is deforming in time. At a later stage during the ELM evolution, blobs of plasma are disconnected from the main plasma for which a mesh refinement also appears to be an optimal solution. We first adapt the mesh at the beginning of the simulation during the initialisation phase, in order to guaranty the equilibrium, then we will apply the modifications on the mesh in order to get the adaptative mesh refinement procedure during the whole simulation.

6.3. Algorithms and high-performance solvers

Participants: Mathieu Faverge, Sébastien Fourestier, Jérémie Gaidamour, Pascal Hénon [Corresponding member], Jun-Ho Her, Cédric Lachat, Xavier Lacoste, François Pellegrini, Pierre Ramet.

6.3.1. *Parallel domain decomposition and sparse matrix reordering*

The work carried out within the Scotch project (see section 5.5) focused on four main axes.

The first one regards the parallelization of the static mapping routines already available in the sequential version of Scotch. Since its version 5.1, released last year, Scotch provides parallel graph partitioning capabilities, but graph partitions are computed to date by means of a parallel multilevel recursive bisection framework. This framework provides partitions of very high quality for a moderate number of parts (about under 512), but load imbalance dramatically increases for larger numbers of parts. Also, the more parts the user wants, the more expensive it is to compute them, because of the recursive bisection process. Consequently, efforts have been put this year on designing a direct k-way parallel graph partitioning framework. In fact, the problem which has been considered in this respect is not plain graph partitioning, but static mapping, because of the increasing need to take into account the topology of the target machine when assigning computations to processing elements. Preliminary results have been achieved [36], during the second post-doc year of Jun-Ho Her, but much has yet to be done, as the cost of parallel direct k-way static mapping algorithms is still extremely high compared to sequential methods.

The second axis concerns dynamic repartitioning. Since graphs may now comprise more than one billion vertices, distributed on machines having more than one hundred thousand processing elements, it is important to be able to compute partitions which create as few data movements as possible with respect to a prior partition. The integration of repartitioning features into the sequential version of Scotch is currently under way, in the context of the PhD of Sébastien Fourestier, and will be extended to the parallel domain afterwards. These two axes were partially supported by the ANR-CIS project “SOLSTICE”.

A third research axis regards the design of specific graph partitioning algorithms. Several applications, such as Schur complement methods for hybrid solvers (see Section 5.4), need k-way partitions where load balance should take into account not only vertices belonging to the sub-domains, but also boundary vertices, which lead to computations on each of the sub-domains which share them. This work, which had been temporarily set aside by lack of manpower, is now being resumed by Jun-Ho Her, in the context of the ANR project “PETAL”.

The fourth axis is the design of efficient and scalable software tools for parallel dynamic remeshing. This is a joint work with Cécile Dobrzynski, which took form this fall with the start of the PhD of Cédric Lachat. Cédric started his work by devising cache-oblivious orderings of the unknowns in the FLuidBox and MMG3D3d software, in order to speed-up computations.

6.3.2. *High-performance direct solvers on multi-platforms*

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 motivate 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. 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. A survey article on these 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.

Adaptation to NUMA architectures. New supercomputers incorporate many microprocessors which include themselves one or many computational cores. These new architectures induce strongly hierarchical topologies. These are called NUMA architectures. In the context of distributed NUMA architectures, a work has begun, in collaboration with the INRIA RUNTIME team, to study optimization strategies, and to improve the scheduling of communications, threads and I/O. Sparse direct solvers are a basic building block of many numerical simulation algorithms. We propose to introduce a dynamic scheduling designed for NUMA architectures in the PaStiX solver. The data structures of the solver, as well as the patterns of communication have been modified to meet the needs of these architectures and dynamic scheduling. We are also interested in the dynamic adaptation of the computation grain to use efficiently multi-core architectures and shared memory. Experiments on several numerical test cases have been performed to prove the efficiency of the approach on different architectures. M. Faverge defended his Ph.D. [1] on these aspects in the context of the NUMASIS ANR CIGC project.

6.3.3. *Hybrid direct-iterative solver based on a Schur complement approach.*

In HIPS, 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 discretisation and others for a 3D problem; the choice of the method also depends on the numerical difficulty of the problem. We 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.

In [34], we have presented an hybrid version of the solver where the Schur complement preconditioner was built using parallel scalar ILUT algorithm. This year we have also developed a parallel version of the algorithms where the Schur complement incomplete factorization is done using a dense block structure. That is to say that there is no additional term dropping in the Schur complement preconditioner other than the ones prescribed by the block pattern defined by the HID graph partitioning. This variant of the preconditioner is more expensive in term of memory but for some difficult test cases they are the only alternative to direct solvers. A general comparison of all the hybrid methods in HIPS has been presented in [35].

This year, J. Gaidamour has defended his Ph.D. [33] on the hybrid solver techniques developed in HIPS.

6.3.4. *MURGE a common interface to sparse linear solvers.*

This year we have also defined a general programming interface for sparse linear solvers. Our goal is to normalize the API to sparse linear solvers and to provide some very simple ways of doing some fastidious tasks such as the parallel matrix assembly for instance. We have thus proposed a generic API specifications called MURGE (<http://murge.gforge.inria.fr>) and also provided some test programs and documentation. This interface has been coded in HIPS and PaStiX. We have also tested this interface in FluidBox and JOREK for HIPS and PaStiX.

7. Contracts and Grants with Industry

7.1. EDF

Participants: Pierre Ramet [Corresponding member], Bruno Lathuilière, Jean Roman [Team Hiepac].

Dates: 2006-2009

Application of a domain decomposition method to the neutronic SPn equations

7.2. SNECMA

Participants: Guillaume Baurin, Rémi Abgrall, Mario Ricchiuto [Corresponding member].

Dates: 2008-2011

Transfert and development of the Residual Distribution schemes in the Natur code (in collaboration with INCKA).

7.3. DASSAULT

Participants: Rémi Abgrall [Corresponding member], Pierre-Elie Normand (PhD student, DGA grant, Dassault-Aviation).

Dates: 2008-2011

Study and validation of very high order SUPG schemes in AETHER.

8. Other Grants and Activities

8.1. National initiatives

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

Participants: Rémi Abgrall [Corresponding member], Mario Ricchiuto, Jiri Treflik.

Grant: SNECMA

Dates: 2006-2009

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

Overview: The AEROCAV 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: Rémi Abgrall, Robin Huart, Pascal Hénon, Pierre Ramet [Corresponding member], Hocine Sellama.

Grant: ANR-06-CIS

Dates: 2006 – 2009

Partners: CEA Cadarache.

Overview: The magneto-hydrodynamic 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 BACCHUS 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 Controlée (DRFC, CEA/Cadarache) and the Laboratoire Bordelais de Recherche en Informatique (LaBRI) and Mathématiques Appliquées de Bordeaux (IMB) at the University of Bordeaux 1.

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

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

Participants: Mathieu Faverge, Abdou Guermouche, Pierre Ramet, Jean Roman [Team Hiepac, corresponding member].

Grant: ANR-05-CIGC-002

Dates: 2006 – 2009

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

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/>

8.1.4. *PETAL: PETascaling ALgorithms for preconditioning used in scientific applications*

Participants: Pascal Hénon [Corresponding member], Jun-Ho Her, François Pellegrini, Pierre Ramet.

Grant: ANR Cosinus 2008

Dates: 2009–2011

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.

Web: <http://petal.saclay.inria.fr/>

8.1.5. *SOLSTICE: High performance solvers and simulations*

Participants: Patrick Amestoy, Sébastien Fourestier, Jérémie Gaidamour, Abdou Guermouche, Pascal Hénon, Jun-Ho Her, François Pellegrini, Pierre Ramet, Jean Roman [Team Hiepac, corresponding member].

Grant: ANR-06-CIS

Dates: 2006 – 2009

Partners: CERFACS, EADS IW, EDF R&D SINETICS, INRIA Rhone-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/>

8.1.6. MACAO

Participants: Rémi Abgrall [Corresponding member], Pascal Jacq.

Grant: Competitivity cluster AESE

Dates: 2006 – 2009

Partners: CERFACS, ONERA, TURBOMECA, CEA, INRIA, etc

Our task in Macao was to couple FluidBox and EISa. FluidBox is a code using unstructured meshes, the unknowns are localised at the vertices (in its the second order finite volume version), while EISa is a structured multibloc solver. The coupling has been done using a special module tuned to that purpose, written in FluidBox and subroutines written in python. From a mathematical point of view, the coupling is realised using Discontinuous Galerkin type of methods, but at the interface structured/unstructured only. Euler and Navier Stokes simulations have been conducted.

8.2. Actions Funded by the EC

8.2.1. ADIGMA: a European project on the development of adaptive higher order variational methods for aerospace applications

Participants: Rémi Abgrall [Corresponding member], Pascal Hénon, Pascal Jacq, Cédric Lachat, Adam Larat, François Pellegrini, Mario Ricchiuto.

Grant: European Commission

Dates: 2006-2009

Partners: AIRBUS F et AIRBUS D, DASSAULT, ALENIA, DLR, ONERA, NLR, ARA, VKI, INRIA, Nanjing 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/

8.2.2. ADDECCO : ADaptive schemes for DEterministic and stoChastiC Flow PrOblems

Participant: Rémi Abgrall [Corresponding member].

Grant: European Commission

Dates: 2009-2014

The numerical simulation of complex compressible flow problem is still a challenge nowadays, even for the simplest physical model such as the Euler and Navier Stokes equations for perfect gases. Researchers in scientific computing need to understand how to obtain efficient, stable, very accurate schemes on complex 3D geometries that are easy to code and to maintain, with good scalability on massively parallel machines. Many people work on these topics, but our opinion is that new challenges have to be tackled in order to combine the outcomes of several branches of scientific computing to get simpler algorithms of better quality without sacrificing their efficiency properties. In this proposal, we will tackle several hard points to overcome for the success of this program.

We first consider the problem of how to design methods that can handle easily mesh refinement, in particular near the boundary, the locations where the most interesting engineering quantities have to be evaluated. CAD tools enable to describe the geometry, then a mesh is generated which itself is used by a numerical scheme. Hence, any mesh refinement process is not directly connected with the CAD. This situation prevents the spread of mesh adaptation techniques in industry and we propose a method to overcome this even for steep problems.

Second, we consider the problem of handling the extremely complex patterns that occur in a flow because of boundary layers: it is not always sufficient to only increase the number of degrees of freedom or the formal accuracy of the scheme. We propose to overcome this with class of very high order numerical schemes that can utilise solution dependant basis functions.

Our third item is about handling unsteady uncertainties in the model, for example in the geometry or the boundary conditions. This need to be done efficiently: the amount of computation increases a priori linearly with the number of uncertain parameters. We propose a non-intrusive method that is able to deal with general probability density functions (pdf), and also able to handle pdfs that may evolve during the simulation via a stochastic optimisation algorithm, for example. This will be combined with the first two items of this proposal. Many random variables may be needed, the curse of dimensionality will be dealt thanks to multiresolution method combined with sparse grid methods.

The aim of this proposal is to design, develop and evaluate solutions to each of these challenges. Currently, and up to our knowledge, none of these problems have been dealt with for compressible flows with steep patterns as in many moderns aerodynamics industrial problems. We propose a work program that will lead to significant breakthroughs for flow simulations with a clear impact on numerical schemes and industrial applications. Our solutions, though developed and evaluated on flow problems, have a wider potential and could be considered for any physical problem that are essentially hyperbolic.

9. Dissemination

9.1. Participation to the Scientific Community

Rémi Abgrall is scientific associate editor of the international journals “Mathematical Modeling 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, and of the “commission d’évaluation de la direction Simulation Numérique en aérodynamique” of ONERA. 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 (Programme COSINUS). He is member of the Comité National du CNRS, section 01. He is member of the board of the GAMNI group of SMAI and is its current responsible.

Pierre Ramet and Rémi Abgrall are members of the GENCI scientific committee (Mathematics and Computer Sciences).

Pascal Hénon and François Pellegrini has been members of the “commission consultative” for the LaBRI in 2009.

Pascal Hénon and Pierre Ramet has been requested as experts for the ANR program COSINUS 2009.

Pascal Hénon has organized the workshop on “GP-GPU computing” at INRIA - Bordeaux Sud-Ouest (around 50 persons). He was also in the decision board for the “Platfrim” project.

Mario Ricchiuto, Cécile Dobrzynski and Rémi Abgrall are, in collaboration with the team–project MC2, preparing the CANUM 2010 in Carcan Maubuisson (june 2010).

9.2. Teaching

Rémi Abgrall has given one master 2 course on high order methods in CFD at ENSEIRB-MATMECA, as well as a master 1 course on aerodynamics.

In complement of the normal teaching activity of the university members and of IPB members, Pascal Hénon teaches at IPB (computer science engineering school).

François Pellegrini gives a master 2 class on the architecture of high-performance systems and ways to exploit them, in the context of practical projects.

Pierre Ramet gives a master 2 class on parallel numerical algorithms at IPB (computer science engineering school).

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

10. Bibliography

Year Publications

Doctoral Dissertations and Habilitation Theses

- [1] M. FAVERGE. *Ordonnancement hybride statique-dynamique en algèbre linéaire creuse pour de grands clusters de machines NUMA et multi-coeurs*, LaBRI, Université Bordeaux I, Talence, Talence, France, December 2009, Ph. D. Thesis.
- [2] J. GAIDAMOUR. *Conception d'un solveur linéaire creux parallèle hybride direct-itératif*, Université Bordeaux I, 351 cours de la Libération, 33405 Talence, France, December 2009, Thèse.
- [3] A. LARAT. *Conception et analyse de schémas d'ordre très élevé distribuant le résidu. Application à la mécanique des fluides.*, IMB, Université Bordeaux I, Talence, 2009, Ph. D. Thesis.
- [4] F. PELLEGRINI. *Contributions au partitionnement de graphes parallèle multi-niveaux / Contributions to parallel multilevel graph partitioning*, Université Bordeaux I, 351 cours de la Libération, 33405 Talence, France, December 2009, Habilitation à Diriger des Recherches.

Articles in International Peer-Reviewed Journal

- [5] R. ABGRALL. *A residual distribution method using discontinuous elements for the computation of possibly non smooth flows.*, in "Advances in Applied Mathematics and Mechanics", 2009, in press.

- [6] R. ABGRALL. *A simple, flexible and generic deterministic approach to uncertainty quantifications in non linear problems: application to fluid flow problems*, in "J. Comput. Physics", 2009, submitted.
- [7] R. ABGRALL. *Construction of simple stable and convergent high order schemes for steady first order Hamilton Jacobi equations*, in "SIAM J. Sci. Comput.", vol. 31, n^o 4, 2009, p. 2419–2446.
- [8] R. ABGRALL. *A residual distribution method using discontinuous elements for the computation of possibly non smooth flows.*, in "Applied Mathematics and Mechanics", 2010, in press.
- [9] R. ABGRALL, S. KARNI. *A Comment on the Computation of Non Conservative Products*, in "J. Comput. Physics", 2010, in press.
- [10] 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", vol. 38, n^o 7, 2009, p. 1314–1323, <http://hal.inria.fr/inria-00333756/en/>.
- [11] R. ABGRALL, C. SHU. *Development of residual distribution schemes for discontinuous Galerkin method: the scalar case with linear elements*, in "Communications in Computational Physics", vol. 5, 2009, p. 376–390.
- [12] R. ABGRALL, J. TREFILÍCK. *An example of high order Residual Distribution scheme using non Lagrange elements*, in "Journal of Scientific Computing", 2010, in revision.
- [13] M. DURUFLE, P. GROB, P. JOLY. *Influence of Gauss and Gauss-Lobatto quadrature rules on the accuracy of a quadrilateral finite element method in the time domain.*, in "Numer. Methods Partial Differ. Equations", vol. 25, n^o 3, 2009, p. 526-551.
- [14] R. HUART, R. ABGRALL, P. RAMET. *Numerical simulation of unsteady MHD flows and applications*, in "Magnetohydrodynamics", vol. 45, n^o 2, 2009, p. 225-232, <http://mhd.sal.lv>.
- [15] C.-Y. KUO, B. NKONGA, M. RICCHIUTO, Y.-C. TAI, B. BRACONNIER. *Dry granular flows with erosion/deposition process.*, in "ESAIM, Proc.", vol. 28, 2009, p. 135-149.
- [16] P.-H. MAIRE, B. NKONGA. *Multi-scale Godunov-type method for cell-centered discrete Lagrangian hydrodynamics.*, in "J. Comput. Phys.", vol. 228, n^o 3, 2009, p. 799-821.
- [17] M. RICHIUTO, A. BOLLERMAN. *Stabilized residual distribution for shallow water simulations.*, in "J. Comput. Physics", vol. 228, n^o 4, 2009, p. 1071–1115.

International Peer-Reviewed Conference/Proceedings

- [18] R. ABGRALL, C. TAVÉ. *Construction of high order residual distribution schemes*, in "Computational Fluid Dynamics 2006", E. DICK, H. DECONINCK (editors), Springer Verlag, 2009, p. 89–94.
- [19] J. GAIDAMOUR, P. HÉNON. *Algorithms To Build Robust Hybrid Direct/Iterative Sparse Linear Solvers*, in "International Conference On Preconditioning Techniques For Scientific And Industrial Applications, Hong-Kong Hong-Kong", 08 2009, <http://hal.archives-ouvertes.fr/hal-00414600/en/>, Abstract long publié dans les actes de la conférences..

- [20] J.-H. HER, F. PELLEGRINI. *Towards efficient and scalable parallel static mapping*, in "CSC'2009", SIAM, October 2009.
- [21] G. HUYSMANS, S. PAMELA, E. VAN DER PLAS, P. RAMET. *Non-Linear MHD simulations of Edge Localised Modes*, in "36th EPS Plasma Physics Conference, Sofia, Bulgaria", June 2009.
- [22] F. PELLEGRINI. *Distillating knowledge about Scotch*, in "Combinatorial Scientific Computing, Dagstuhl, Germany", U. NAUMANN, O. SCHENK, H. D. SIMON, S. TOLEDO (editors), Dagstuhl Seminar Proceedings, n^o 09061, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, February 2009, <http://drops.dagstuhl.de/opus/volltexte/2009/2091>, 12 pages.
- [23] M. RICCHIUTO, R. ABGRALL. *Stable and convergent residual distribution for time-dependent conservation laws*, in "Computational Fluid Dynamics 2006", E. DICK, H. DECONINCK (editors), Springer Verlag, 2009, p. 89–94.
- [24] N. VILLEDIEU-LIGOUT, M. RICCHIUTO, H. DECONINCK. *High-order residual distribution : discontinuity capturing crosswind dissipation and diffusion*, in "Computational Fluid Dynamics 2006", E. DICK, H. DECONINCK (editors), Springer Verlag, 2009, p. 71–76.

Workshops without Proceedings

- [25] M. BARRAULT, B. LATHUILIÈRE, P. RAMET, J. ROMAN. *A Non Overlapping Parallel Domain Decomposition Method Applied to The Simplified Transport Equations*, in "International Conference on Mathematics, Computational Methods & Reactor Physics, New-York, USA", May 2009.

Scientific Books (or Scientific Book chapters)

- [26] R. ABGRALL, A. LARAT, M. RICCHIUTO. *Construction of high order residual distribution scheme*, in "CFD Review", M. HAFEZ (editor), Wiley, 2009.

Research Reports

- [27] M. RICCHIUTO, R. ABGRALL. *Explicit Runge-Kutta Residual Distribution schemes for Time Dependent Problems : second order case*, n^o RR-6998, INRIA, 2009, <http://hal.inria.fr/inria-00406958/en/>, Research Report.

References in notes

- [28] R. ABGRALL. *Essentially non-oscillatory residual distribution schemes for hyperbolic problems.*, in "J. Comput. Phys.", vol. 214, n^o 2, 2006, p. 773-808.
- [29] M. BERGOT, G. COHEN, M. DURUFLÉ. *Higher-Order finite elements for hybrid meshes using new nodal pyramidal elements*, in "Journal of Scientific Computing", 2009, to appear.
- [30] T. COUPEZ, H. DIGONNET, R. DUCLOUX. *Parallel meshing and remeshing*, in "Appl. Math. Modelling", vol. 25, 2000, p. 153-175.
- [31] 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.

-
- [32] 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.
- [33] J. GAIDAMOUR. *Conception d'un solveur linéaire creux parallèle hybride direct-itératif*, Université Bordeaux I, 351 cours de la Libération, 33405 Talence, France, December 2009, Thèse.
- [34] 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.
- [35] J. GAIDAMOUR, P. HÉNON. *Algorithms To Build Robust Hybrid Direct/Iterative Sparse Linear Solvers*, in "International Conference On Preconditioning Techniques For Scientific And Industrial Applications, Hong-Kong Hong-Kong", 08 2009, <http://hal.archives-ouvertes.fr/hal-00414600/en/>, Abstract long publié dans les actes de la conférences..
- [36] J.-H. HER, F. PELLEGRINI. *Towards efficient and scalable parallel static mapping*, in "CSC'2009", SIAM, October 2009.
- [37] 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.
- [38] 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.
- [39] K. SCHLOEGEL, G. KARYPIS, V. KUMAR. *Dynamic repartitioning of adaptively refined meshes*, in "Supercomputing '98: Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)", IEEE, 1998, p. 1–8.
- [40] C. WALSHAW, M. CROSS, M. G. EVERETT. *Parallel Dynamic Graph Partitioning for Adaptive Unstructured Meshes*, in "J. Parallel and Distributed Computing", vol. 47, 1997, p. 102–108.