



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Team HiePACS

*High-End Parallel Algorithms for
Challenging Numerical Simulations*

Bordeaux - Sud-Ouest

Theme : Distributed and High Performance Computing

Activity
R *eport*

2009

Table of contents

| | |
|--|-----------|
| 1. Team | 1 |
| 2. Overall Objectives | 1 |
| 2.1. Introduction | 1 |
| 2.2. Highlights of the year | 3 |
| 3. Scientific Foundations | 3 |
| 3.1. Introduction | 3 |
| 3.2. High-performance computing on next generation architectures | 3 |
| 3.3. High performance solvers for large linear algebra problems | 5 |
| 3.3.1. Hybrid direct/iterative solvers based on algebraic domain decomposition techniques | 6 |
| 3.3.2. Hybrid solver based on a combination of a multigrid method and a direct method | 6 |
| 3.3.3. Linear Krylov solvers | 6 |
| 3.4. High performance Fast Multipole Method for N-body problems | 8 |
| 3.4.1. Improvement of calculation efficiency | 8 |
| 3.4.2. Non uniform distributions | 9 |
| 3.4.3. Fast Multipole Method for dislocation operators | 9 |
| 3.4.4. Fast Multipole Method for boundary element methods | 9 |
| 3.5. Efficient algorithmics for code coupling in complex simulations | 9 |
| 3.5.1. Efficient schemes for multiscale simulations | 9 |
| 3.5.2. Coupling of complex simulations based on the hypergraph model | 10 |
| 3.5.3. Steering and Interacting with complex coupled simulations | 10 |
| 4. Application Domains | 11 |
| 4.1. Introduction | 11 |
| 4.2. Material physics | 11 |
| 4.2.1. Hybrid materials | 12 |
| 4.2.2. Material failures | 13 |
| 4.3. Application framework customers of high performance linear algebra solvers | 13 |
| 5. Software | 14 |
| 5.1. Introduction | 14 |
| 5.2. MaPHyS | 14 |
| 5.3. EPSN | 15 |
| 5.4. Other software | 16 |
| 6. New Results | 17 |
| 6.1. Algorithms and high-performance solvers | 17 |
| 6.1.1. Hybrid direct/iterative solvers based on algebraic domain decomposition techniques | 17 |
| 6.1.2. Hybrid solver based on a combination of a multigrid method and a direct method | 17 |
| 6.1.3. Domain decomposition methods to solve neutron transport equations | 17 |
| 6.1.4. High performance algorithms for wave propagation | 18 |
| 6.2. Efficient algorithmics for code coupling in complex simulations | 18 |
| 6.3. Computational steering environment for distributed numerical simulations | 18 |
| 6.3.1. Model for the steering of parallel-distributed simulations | 18 |
| 6.3.2. Distributed Shared Memory approach for the steering of parallel simulations | 19 |
| 6.4. Material physics | 19 |
| 7. Contracts and Grants with Industry | 20 |
| 8. Other Grants and Activities | 20 |
| 8.1. National initiatives | 20 |
| 8.1.1. NOSSI: New platform for parallel, hybrid quantum/classical simulations | 20 |
| 8.1.2. NUMASIS: adaptation and optimization of the applicative performance on NUMA architectures | 21 |
| 8.1.3. SOLSTICE: High performance solvers for frontier simulations | 21 |

| | |
|--|-----------|
| 8.2. International initiatives | 21 |
| 9. Dissemination | 22 |
| 9.1. Participation to the Scientific Community | 22 |
| 9.2. Teaching | 22 |
| 10. Bibliography | 22 |

HIEPACS is an INRIA Team joint with University of Bordeaux and CNRS (LaBRI, UMR 5800) and will be a Research Initiative of the joint Laboratory INRIA-CERFACS on High-Performance Computing (<http://hiepacs.bordeaux.inria.fr/>).

1. Team

Research Scientist

Jean Roman [INRIA, Team Leader, Research Director (on secondment), Professor at IPB, HdR]
Olivier Coulaud [INRIA, Research Director (DR)]
Luc Giraud [INRIA, Research Director (DR), HdR]
Xavier Vasseur [CERFACS, ALGO Project, Senior Scientist]

Faculty Member

Aurélien Esnard [University Bordeaux 1, Assistant Professor (MdC)]
Abdou Guermouche [University Bordeaux 1, Assistant Professor (MdC)]
Nicolas Richart [University Bordeaux 1 and IPB, ATER]

External Collaborator

Iain Duff [CERFACS, ALGO Project, Senior Scientist , HIEPACS **Scientific Advisor**]
Philip Avery [Stanford University, Farhat Research Group, Senior Scientist]
Pierre Fortin [Paris 6 University, LIP6, Assistant Professor (MdC)]
Guillaume Latu [Strasbourg University, LSIT, Assistant Professor (MdC), on secondment at CEA Cadarache]

Technical Staff

Damien Genet [LaBRI, funding from ANR ARA CIS NOSSI, since Sep. 2008]
Yohan Lee-Tin-Yien [INRIA, funding from ADT ParScaLi, since Dec. 2009]

PhD Student

Rached Abdelkhalek [TOTAL, CIFRE, since Jan. 2008]
Mathieu Chanaud [INRIA, funding from INRIA and CEA, since Dec. 2007]
Fabrice Dupros [BRGM, funding from BRGM and ANR CIGC NUMASIS, since Feb. 2007]
Jérôme Soumagne [CSCS, funding from Europe FP7/ICT/FET NextMuSE STREP since Apr. 2009]

Post-Doctoral Fellow

Azzam Haidar [INPT, since Sep. 2009]
Pavel Jiranek [CERFACS, ALGO Project, since Oct. 2008]
Peter Koval [LaBRI, funding from ANR ARA CIS NOSSI, since Sep. 2009]

Administrative Assistant

Josy Baron [INRIA]

Other

Mohamed Amine El Afrit [INRIA, Internship from Feb. to Sep. 2009]
Amal Khabou [INRIA, Internship from Feb. to Sep. 2009]

2. Overall Objectives

2.1. Introduction

Over the last few decades, there have been innumerable science, engineering and societal breakthroughs enabled by the development of high performance computing (HPC) applications, algorithms and architectures. These powerful tools have provided researchers with the ability to computationally find efficient solutions for some of the most challenging scientific questions and problems in medicine and biology, climatology, nanotechnology, energy and environment. It is admitted today that *numerical simulation is the third pillar for the development of scientific discovery at the same level as theory and experimentation*. Numerous reports and

papers also confirmed that very high performance simulation will open new opportunities not only for research but also for a large spectrum of industrial sectors (see for example the documents available on the web link <http://www.sc.doe.gov/ascr/ProgramDocuments/ProgDocs.html>).

An important force which has continued to drive HPC has been to focus on frontier milestones which consist in technical goals that symbolize the next stage of progress in the field. In the 1990s, the HPC community sought to achieve computing at a teraflop rate and currently we are able to compute on the first leading architectures at a petaflop rate. Generalist petaflop supercomputers are likely to be available in 2010-2012 and some communities are already in the early stages of thinking about what computing at the exaflop level would be like.

For application codes to sustain a petaflop and more in the next few years, hundreds of thousands of processor cores or more will be needed, regardless of processor technology. Currently, a few HPC simulation codes easily scale to this regime and major code development efforts are critical to achieve the potential of these new systems. Scaling to a petaflop and more will involve improving physical models, mathematical modelling, super scalable algorithms that will require paying particular attention to acquisition, management and visualization of huge amounts of scientific data.

In this context, the purpose of the HiεPACS project is to perform efficiently frontier simulations arising from challenging research and industrial *multiscale* applications. The solution of these challenging problems require a multidisciplinary approach involving applied mathematics, computational and computer sciences. In applied mathematics, it essentially involves advanced numerical schemes. In computational science, it involves massively parallel computing and the design of highly scalable algorithms and codes to be executed on future petaflop (and beyond) platforms. Through this approach, HiεPACS intends to contribute to all steps that go from the design of new high-performance more scalable, robust and more accurate numerical schemes to the optimized implementations of the associated algorithms and codes on very high performance supercomputers. This research will be conducted on close collaboration in particular with European and US initiatives or projects such as PRACE (Partnership for Advanced Computing in Europe – <http://www.prace-project.eu/>) or IESP (International Exascale Software Project – http://icl.cs.utk.edu/iesp/Main_Page).

In order to address these research challenges, some of the researchers of the former ScAlApplix INRIA Project-Team and some researchers of the Parallel Algorithms Project from CERFACS have joined HiεPACS in the framework of the joint INRIA-CERFACS Laboratory on High Performance Computing. The director of the joint laboratory is J. Roman while I.S. Duff is the senior scientific advisor. HiεPACS will be the first research initiative of this joint Laboratory. Because of his strong involvement in RAL and his outstanding action in other main initiatives in UK and worldwide, I.S. Duff appears as an external collaborator of the HiεPACS project while his contribution will be significant. There are two other external collaborators. Namely, P. Fortin who will be mainly involved in the activities related to the parallel fast multipole development and G. Latu who will contribute to research actions related to the emerging new computing facility.

The methodological part of HiεPACS covers several topics. First, we address generic studies concerning massively parallel computing, the design of high-end performance algorithms and software to be executed on future petaflop (and beyond) platforms. Next, several research perspectives in scalable parallel linear algebra techniques are addressed, in particular hybrid approaches for large linear systems. Then we consider research plans for N-body interaction computations based on efficient parallel fast multipole methods and finally, we address research tracks related to the algorithmic challenges for complex code couplings in multiscale simulations.

Currently, we have one major multiscale application that is in *material physics*. We will contribute to all steps of the design of the parallel simulation tool. More precisely, our applied mathematics skill will contribute to the modelling and our advanced numerical schemes will help in the design and efficient software implementation for very large parallel multiscale simulations. Moreover, the robustness and efficiency of our algorithmic research in linear algebra will be validated through industrial and academic collaborations with different partners involved in various application fields.

Our high performance software packages will be integrated in several academic or industrial complex codes and will be validated on very large scale simulations. For all our software developments, we will use first the various (very) large parallel platforms available through CERFACS and GENCI in France (CCRT, CINES and IDRIS Computational Centers), and next the high-end parallel platforms that will be available via European and US initiatives or projects such that PRACE.

2.2. Highlights of the year

- Luc Giraud and Jean Roman have organized in collaboration with Stéphane Lanteri (INRIA Sophia Antipolis - Méditerranée) and Yousef Saad (University of Minnesota) the minisymposium entitled *Toward robust hybrid parallel sparse solvers for large scale applications* at SIAM Conference on Computational Science and Engineering (CSE09), March 2-6, 2009, Miami, Florida.
- Luc Giraud and Jean Roman have organized in collaboration with Victorita Dolean and Stéphane Lanteri (INRIA Sophia Antipolis - Méditerranée) the CEA-EDF-INRIA School about *Robust methods and algorithms for solving large algebraic systems on modern high performance computing systems*, March 30 to April 3 2009 - INRIA Sophia Antipolis- Méditerranée.
- The paper of Rached Abdelkhalek, Henri Calendra, Olivier Coulaud, Guillaume Latu and Jean Roman entitled *Fast seismic modeling and reverse time migration on a GPU cluster* has got the Best Paper Award at HPCS 2009 (Leipzig, 21-24 June, IEEE).

3. Scientific Foundations

3.1. Introduction

The methodological component of HiePACS concerns the expertise for the design as well as the efficient and scalable implementation of highly parallel numerical algorithms to perform frontier simulations. In order to address these computational challenges a hierarchical organization of the research is considered. In this bottom-up approach, we first consider in Section 3.2 generic topics concerning high performance computational science. The activities described in this section are transversal to the overall project and its outcome will support all the other research activities at various levels in order to ensure the parallel scalability of the algorithms. The aim of this activity is not to study general purpose solution but rather to address these problems in close relation with specialists of the field in order to adapt and tune advanced approaches in our algorithmic designs. The next activity, described in Section 3.3, is related to the study of parallel linear algebra techniques that currently appear as promising approaches to tackle huge problems on millions of cores. We highlight the linear problems (linear systems or eigenproblems) because they are in many large scale applications the main computational intensive numerical kernels and often the main performance bottleneck. These parallel numerical techniques will be the basis of both academic and industrial collaborations described in Section 4.2 and Section 4.3, but will also be closely related to some functionalities developed in the parallel Fast Multipole activity described in Section 3.4. Finally, as the accuracy of the physical models increases, there is a real need to go for parallel efficient algorithm implementation for multiphysics and multiscale modelling in particular in the context of code coupling. The challenges associated with this activity will be addressed in the framework of the activity described in Section 3.5.

3.2. High-performance computing on next generation architectures

Participants: Rached Abdelkhalek, Olivier Coulaud, Iain Duff, Pierre Fortin, Luc Giraud, Abdou Guer-mouche, Guillaume Latu, Jean Roman.

The research directions proposed in HiePACS are strongly influenced by both the applications we are studying and the architectures that we target (i.e., massively parallel architectures, ...). Our main goal is to study the methodology needed to efficiently exploit the new generation of high-performance computers with all the constraints that it induces. To achieve this high-performance with complex applications we have to study both algorithmic problems and the impact of the architectures on the algorithm design.

From the application point of view, the project will be interested in multiresolution, multiscale and hierarchical approaches which lead to multi-level parallelism schemes. This hierarchical parallelism approach is necessary to achieve good performance and high-scalability on modern massively parallel platforms. In this context, more specific algorithmic problems are very important to obtain high performance. Indeed, the kind of applications we are interested in are often based on data redistribution for example (e.g. code coupling applications). This well-known issue becomes very challenging with the increase of both the number of computational nodes and the amount of data. Thus, we have both to study new algorithms and to adapt the existing ones. In addition, some issues like task scheduling have to be restudied in this new context. It is important to note that the work done in this area will be applied for example in the context of code coupling (see Section 3.5).

Considering the complexity of modern architectures like massively parallel architectures (i.e., Blue Gene-like platforms) or new generation heterogeneous multicore architectures, task scheduling becomes a challenging problem which is central to obtain a high efficiency. Of course, this work requires the use/design of scheduling algorithms and models specifically to tackle our target problem. This has to be done in collaboration with our colleagues from the scheduling community like for example O. Beaumont (INRIA CEPAGE Project-Team). It is important to note that this topic is strongly linked to the underlying programming model. Indeed, considering multicore architectures, it has appeared, in the last five years, that the best programming model is an approach mixing multi-threading within computational nodes and message passing between them. In the last five years, a lot of work has been developed in the high-performance computing community to understand what is critic to efficiently exploit massively multicore platforms that will appear in the near future. It appeared that the key for the performance is firstly the grain of computations. Indeed, in such platforms the grain of the parallelism must be small so that we can feed all the processors with a sufficient amount of work. It is thus very crucial for us to design new high performance tools for scientific computing in this new context. This will be done in the context of our solvers, for example, to adapt to this new parallel scheme. Secondly, the larger the number of cores inside a node, the more complex the memory hierarchy. This remark impacts the behaviour of the algorithms within the node. Indeed, on this kind of platforms, NUMA effects will be more and more problematic. Thus, it is very important to study and design data-aware algorithms which take into account the affinity between computational threads and the data they access. This is particularly important in the context of our high-performance tools. Note that this work has to be based on an intelligent cooperative underlying runtime (like the `marcel` thread library developed by the INRIA RUNTIME Project-Team) which allows a fine management of data distribution within a node.

Another very important issue concerns high-performance computing using “heterogeneous” resources within a computational node. Indeed, with the emergence of the GPU and the use of more specific co-processors (like clearspeak cards, ...), it is important for our algorithms to efficiently exploit these new kind of architectures. To adapt our algorithms and tools to these accelerators, we need to identify what can be done on the GPU for example and what cannot. Note that recent results in the field have shown the interest of using both regular cores and GPU to perform computations. Note also that in opposition to the case of the parallelism granularity needed by regular multicore architectures, GPU requires coarser grain parallelism. Thus, making both GPU and regular cores work all together will lead to two types of tasks in terms of granularity. This represents a challenging problem especially in terms of scheduling. Our final goal would be to have high performance solvers and tools which can efficiently run on all these types of complex architectures by exploiting all the resources of the platform (even if they are heterogeneous).

In order to achieve an advanced knowledge concerning the design of efficient computational kernels to be used on our high performance algorithms and codes, we will develop research activities first on regular frameworks before extending them to more irregular and complex situations. In particular, we will work first on optimized dense linear algebra kernels and we will use them in our more complicated hybrid solvers for sparse linear algebra and in our fast multipole algorithms for interaction computations. In this context, we will participate to the development of those kernels in collaboration with groups specialized in dense linear algebra. In particular, we intend develop a strong collaboration with the group of Jack Dongarra at the University Of Tennessee. The objectives will be to develop dense linear algebra algorithms and libraries for multicore architectures

in the context the PLASMA project (<http://icl.cs.utk.edu/plasma/>) and for GPU and hybrid multicore/GPU architectures in the context of the MAGMA project (<http://icl.cs.utk.edu/magma/>).

The applications targeting massively parallel architectures are very sensitive to communication or I/O management schemes. This observation becomes particularly true, when we consider applications dealing with a huge amount of data like very large scale simulations that may produce petaBytes of data. Thus, in the continuation of the work we did around out-of-core extensions of our former sparse linear solvers, we will study how we can efficiently deal with this huge amount of data. Obtaining performance when relying on I/O operations or on data transfers is mainly constrained by the capacity to overlap as much as much possible these operations with computations. Another key feature is prefetching in the context of I/O intensive applications. Even, if the problem is a well-known issue which has been studied in the past decade, it remains very complex regarding the complexity of our target platforms were we already need prefetching and asynchronism to efficiently exploit the platform (this is particularly true in the case of GPU).

A more prospective objective is to study the fault tolerance in the context of large-scale scientific applications for massively parallel architectures. Indeed, with the increase of the number of computational cores per node, the probability of a hardware crash on a core is dramatically increased. This represents a crucial problem that needs to be addressed. However, we will only study it at the algorithmic/application level even if it needed lower-level mechanisms (at OS level or even hardware level). Of course, this work can be done at lower levels (at operating system) level for example but we do believe that handling faults at the application level provides more knowledge about what has to be done (at application level we know what is critical and what is not). The approach that we will follow will be based on the use of a combination of fault-tolerant implementations of the run-time environments we use (like for example FT-MPI) and an adaptation of our algorithms to try to manage this kind of faults. This topic represents a very long range objective which needs to be addressed to guaranty the robustness of our solvers and applications.

Finally, it is important to note that the main goal of HiePACS is to design tools and algorithms that will be used within complex simulation frameworks on next-generation parallel machines. Thus, we intend with our partners to use the proposed approach in complex scientific codes and to validate them within very large scale simulations.

3.3. High performance solvers for large linear algebra problems

Participants: Philip Avery, Mathieu Chanaud, Olivier Coulaud, Iain Duff, Fabrice Dupros, Luc Giraud, Abdou Guermouche, Azzam Haidar, Pavel Jiranek, Yohan Lee-Tin-Yien, Jean Roman, Xavier Vasseur.

Starting with the developments of basic linear algebra kernels tuned for various classes of computers, a significant knowledge on the basic concepts for implementations on high-performance scientific computers has been accumulated. Further knowledge has been acquired through the design of more sophisticated linear algebra algorithms fully exploiting those basic intensive computational kernels. In that context, we still look at the development of new computing platforms and their associated programming tools. This enables us to identify the possible bottlenecks of new computer architectures (memory path, various level of caches, inter processor or node network) and to propose ways to overcome them in algorithmic design. With the goal of designing efficient scalable linear algebra solvers for large scale applications, various tracks will be followed in order to investigate different complementary approaches. Sparse direct solvers have been for years the methods of choice for solving linear systems of equations, it is nowadays admitted that such approaches are not scalable neither from a computational complexity nor from a memory view point for large problems such as those arising from the discretization of large 3D PDE problems. Some initiatives exist to further improve existing parallel packages, those efforts are mainly related to advanced software engineering. Although we will not contribute directly to this activity, we will use parallel sparse direct solvers as building boxes for the design of some of our parallel algorithms such as the hybrid solvers described in the sequel of this section. Our activities in that context will mainly address preconditioned Krylov subspace methods; both components, preconditioner and Krylov solvers, will be investigated.

3.3.1. Hybrid direct/iterative solvers based on algebraic domain decomposition techniques

One route to the parallel scalable solution of large sparse linear systems in parallel scientific computing is the use of hybrid methods that combine direct and iterative methods. These techniques inherit the advantages of each approach, namely the limited amount of memory and natural parallelization for the iterative component and the numerical robustness of the direct part. The general underlying ideas are not new since they have been intensively used to design domain decomposition techniques; those approaches cover a fairly large range of computing techniques for the numerical solution of partial differential equations (PDEs) in time and space. Generally speaking, it refers to the splitting of the computational domain into sub-domains with or without overlap. The splitting strategy is generally governed by various constraints/objectives but the main one is to express parallelism. The numerical properties of the PDEs to be solved are usually intensively exploited at the continuous or discrete levels to design the numerical algorithms so that the resulting specialized technique will only work for the class of linear systems associated with the targeted PDE.

In that context, we attempt to apply to general unstructured linear systems domain decomposition ideas. More precisely, we will consider numerical techniques based on a non-overlapping decomposition of the graph associated with the sparse matrices. The vertex separator, built by a graph partitioner, will define the interface variables that will be solved iteratively using a Schur complement techniques, while the variables associated with the internal sub-graphs will be handled by a sparse direct solver. Although the Schur complement system is usually more tractable than the original problem by an iterative technique, preconditioning treatment is still required. For that purpose, the algebraic additive Schwarz technique initially developed for the solution of linear systems arising from the discretization of elliptic and parabolic PDE's will be extended. Linear systems where the associated matrices are symmetric in pattern will be first studied but extension to unsymmetric matrices will be latter considered. The main focus will be on difficult problems (including non-symmetric and indefinite ones) where it is harder to prevent growth in the number of iterations with the number of subdomains when considering massively parallel platforms. In that respect, we will consider algorithms that exploit several sources and grains of parallelism to achieve high computational throughput. This activity may involve collaborations with developers of sparse direct solvers and will lead to the development to the library MaPHyS (see Section 5.2).

3.3.2. Hybrid solver based on a combination of a multigrid method and a direct method

The multigrid methods are among the most promising numerical techniques to solve large linear system of equations arising from the discretization of PDE's. Their ideal scalabilities, linear growth of memory and floating-point operations with the number of unknowns, for solving elliptic equations make them very appealing for petascale computing and a lot of research works in the recent years has been devoted to the extension to other types of PDE.

In this work (Ph. D. of Mathieu Chanaud in collaboration with CEA/CESTA), we consider a specific methodology for solving large linear systems arising from Maxwell equations discretized with first-order Nédelec elements. This solver combines a parallel direct solver and full multigrid cycles. The goal of this method is to compute the solution for problems defined on fine irregular meshes with minimal overhead costs when compared to the cost of applying a classical direct solver on the coarse mesh.

The direct solver can handle linear systems with up to 100 million unknowns, but this size is limited by the computer memory, so that finer problem resolutions that often occur in practice cannot be handled by this direct solver. The aim of the new method is to provide a way to solve problems with up to 1 billion unknowns, given an input coarse mesh with up to 100 million unknowns. The input mesh is used as the initial coarsest grid. Fine meshes, where only smoothing sweeps are performed, will be generated automatically. Such large problem size can be solved because the matrix is assembled only on the coarsest mesh, while on the finer meshes, multigrid cycles are performed in a matrix-free manner.

3.3.3. Linear Krylov solvers

Preconditioning is the main focus of the two activities described above. They aim at speeding up the convergence of a Krylov subspace method that is the complementary component involved in the solvers of

interest for us. In that framework, we believe that various aspects deserve to be investigated; we will consider the following ones:

Preconditioned block Krylov solvers for multiple right-hand sides. In many large scientific and industrial applications, one has to solve a sequence of linear systems with several right-hand sides given simultaneously or in sequence (radar cross section calculation in electromagnetism, various source locations in seismic, parametric studies in general, ...). For “simultaneous” right-hand sides, the solvers of choice have been for years based on matrix factorizations as the factorization is performed once and simple and cheap block forward/backward substitutions are then performed. In order to effectively propose alternative to such solvers, we need to have efficient preconditioned Krylov subspace solvers. In that framework, block Krylov approaches, where the Krylov spaces associated with each right-hand sides are shared to enlarge the search space will be considered. They are not only attractive because of this numerical feature (larger search space), but also from an implementation point of view. Their block-structures exhibit nice features with respect to data locality and re-usability that comply with the memory constraint of multicore architectures. For right-hand sides available one after each other, various strategies that exploit the information available in the sequence of Krylov spaces (e.g. spectral information) will be considered that include for instance technique to perform incremental update of the preconditioner or to built augmented Krylov subspaces.

Flexible Krylov subspace methods with recycling techniques. In many situations, it has been observed that significant convergence improvements can be achieved in preconditioned Krylov subspace methods by enriching them with some spectral information. On the other hand effective preconditioning strategies are often designed where the preconditioner varies from one step to the next (e.g. in domain decomposition methods, when approximate solvers are considered for the interior problems, or more generally for block preconditioning technique where approximate block solution are used) so that a flexible Krylov solver is required. In that context, we intend to investigate how numerical techniques implementing subspace recycling and/or incremental preconditioning can be extended and adapted to cope with this situation of flexible preconditioning; that is, how can we numerically benefit from the preconditioning implementation flexibility.

Stopping criteria for iterative methods in the framework of backward error minimization algorithms. The stopping criterion is a central component of any iterative schemes since it is crucial to evaluate the numerical “quality” of the iterate when the solution is stopped. The backward analysis is a general framework that permits such a a posteriori analysis and stopping criteria of Krylov subspace method are commonly based on it. On the other hand the underlying heuristics that conduce to the selection of the iterate in the Krylov subspace at each step (Petrov-Galerkin approaches or residual norm minimization) are unrelated to this stopping criterion. We intend to investigate new iterate selection strategies guided by the objective to minimize the targeted backward error in order to fully couple the iterate selection and the targeted numerical quality of the final solution.

Krylov solver for complex symmetric non-Hermitian matrices. In material physics when the absorption spectrum of a molecule due to an exterior field is computed, we have to solve for each frequency a dense linear system where the matrix depends on the frequency. The sequence of matrices are complex symmetric non-Hermitian. While a direct approach can be used for small molecules, a Krylov subspace solver must be considered for larger molecules. Typically, Lanczos-type methods are used to solve these systems but the convergence is often slow. Based on our earlier experience on preconditioning techniques for dense complex symmetric non-Hermitian linear system in electromagnetism, we are interested in designing new preconditioners for this class of material physics applications. A first track will consist in building preconditioners on sparsified approximation of the matrix as well as computing incremental updates, eg. Sherman-Morrison type, of the preconditioner when the frequency varies. This action will be developed in the framework of the research activity described in Section 4.2.

Extension or modification of Krylov subspace algorithms for multicore architectures. Finally to match as much as possible to the computer architecture evolution and get as much as possible performance out of the computer, a particular attention will be paid to adapt, extend or develop numerical schemes that comply with the efficiency constraints associated with the available computers. Nowadays, multicore architectures seem to become widely used, where memory latency and bandwidth are the main bottlenecks; investigations on

communication avoiding techniques will be undertaken in the framework of preconditioned Krylov subspace solvers as a general guideline for all the items mentioned above.

Eigensolvers. Many eigensolvers also rely on Krylov subspace techniques. Naturally some links exist between the Krylov subspace linear solvers and the Krylov subspace eigensolvers. We plan to study the computation of eigenvalue problems with respect to the following three different axes:

- Exploiting the link between Krylov subspace methods for linear system solution and eigensolvers, we intend to develop advanced iterative linear methods based on Krylov subspace methods that use some spectral information to build part of a subspace to be recycled, either through space augmentation or through preconditioner update. This spectral information may correspond to a certain part of the spectrum of the original large matrix or to some approximations of the eigenvalues obtained by solving a reduced eigenproblem. This technique will also be investigated in the framework of block Krylov subspace methods.
- In the framework of an FP7 Marie project (MyPlanet), we intend to study parallel robust nonlinear quadratic eigensolvers. It is a crucial question in numerous technologies like the stability and vibration analysis in classical structural mechanics. Ranging from simple nonlinear stationary iterations to Newton's type approaches will be considered.
- In the context of the calculation of the ground state of an atomistic system, eigenvalue computation is a critical step; more accurate and more efficient parallel and scalable eigensolvers are required (see Section 4.2).

3.4. High performance Fast Multipole Method for N-body problems

Participants: Olivier Coulaud, Pierre Fortin, Luc Giraud, Jean Roman.

In most scientific computing applications considered nowadays as computational challenges (like biological and material 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. For instance, in the N-body problem arising from these application fields, we must compute all pairwise interactions among N objects (particles, lines, ...) at every timestep. Among these methods, the Fast Multipole Method (FMM) developed for gravitational potentials in astrophysics and for electrostatic (coulombic) potentials in molecular simulations solves this N-body problem for any given precision with $O(N)$ runtime complexity against $O(N^2)$ for the direct computation.

The potential field is decomposed in a near field part, directly computed, and a far field part approximated thanks to multipole and local expansions. In the former ScAlAppIix project, we introduced a matrix formulation of the FMM that exploits the cache hierarchy on a processor through the Basic Linear Algebra Subprograms (BLAS). Moreover, we developed a parallel adaptive version of the FMM algorithm for heterogeneous particle distributions, which is very efficient on parallel clusters of SMP nodes. Finally on such computers, we developed the first hybrid MPI-thread algorithm, which enables to reach better parallel efficiency and better memory scalability. We plan to work on the following points in HiEPACS .

3.4.1. Improvement of calculation efficiency

Nowadays, the high performance computing community is examining alternative architectures that address the limitations of modern cache-based designs. GPU (Graphics Processing Units) and the Cell processor have thus already been used in astrophysics and in molecular dynamics. The Fast Mutipole Method has also been implemented on GPU. We intend to examine the potential of using these forthcoming processors as a building block for high-end parallel computing in N-body calculations. More precisely, we want to take advantage of our specific underlying BLAS routines to obtain an efficient and easily portable FMM for these new architectures. Algorithmic issues such as dynamic load balancing among heterogeneous cores will also have to be solved in order to gather all the available computation power. This research action will be conducted on close connection with the activity described in Section 3.2.

3.4.2. *Non uniform distributions*

In many applications arising from material physics or astrophysics, the distribution of the data is highly non uniform and the data can grow between two time steps. As mentioned previously, we have proposed a hybrid MPI-thread algorithm to exploit the data locality within each node. We plan to further improve the load balancing for highly non uniform particle distributions with small computation grain thanks to dynamic load balancing at the thread level and thanks to a load balancing correction over several simulation time steps at the process level.

3.4.3. *Fast Multipole Method for dislocation operators*

The engine that we develop will be extended to new potentials arising from material physics such as those used in dislocation simulations. The interaction between dislocations is long ranged ($O(1/r)$) and anisotropic, leading to severe computational challenges for large-scale simulations. Several approaches based on the FMM or based on spatial decomposition in boxes are proposed to speed-up the computation. In dislocation codes, the calculation of the interaction forces between dislocations is still the most CPU time consuming. This computation has to be improved to obtain faster and more accurate simulations. Moreover, in such simulations, the number of dislocations grows while the phenomenon occurs and these dislocations are not uniformly distributed in the domain. This means that strategies to dynamically balance the computational load are crucial to achieve high performance.

3.4.4. *Fast Multipole Method for boundary element methods*

The boundary element method (BEM) is a well known solution of boundary value problems appearing in various fields of physics. With this approach, we only have to solve an integral equation on the boundary. This implies an interaction that decreases in space, but results in the solution of a dense linear system with $O(N^3)$ complexity. The FMM calculation that performs the matrix-vector product enables the use of Krylov subspace methods. Based on the parallel data distribution of the underlying octree implemented to perform the FMM, parallel preconditioners can be designed that exploit the local interaction matrices computed at the finest level of the octree. This research action will be conducted on close connection with the activity described in Section 3.3. Following our earlier experience, we plan to first consider approximate inverse preconditioners that can efficiently exploit these data structures.

3.5. Efficient algorithmics for code coupling in complex simulations

Participants: Olivier Coulaud, Aurélien Esnard, Damien Genet, Nicolas Richart, Jean Roman, Jérôme Soumagne.

Many important physical phenomena in material physics and climate modelling are inherently complex applications. They often use multiphysics or multiscale approaches, that couple different models and codes. There is typically one model per different scale or physics; and each model is implemented by a parallel code. For instance, to model crack propagation, one uses two scales: an atomistic model and a continuum model discretized by a finite element method. These phenomena are simulated by coupling different parallel codes such as molecular dynamic code and elasticity code.

The experience that we have acquired in the ScAlAppIix project through the activities in crack propagation simulations with LibMultiScale and in M-by-N computational steering (coupling simulation with parallel visualization tools) with EPSN shows us that if the model aspect was well studied, several problems in parallel or distributed algorithms are still open and not well studied. In the context of code coupling in HiePACS, we want to contribute more precisely to the following points.

3.5.1. *Efficient schemes for multiscale simulations*

As mentioned previously, many important physical phenomena, such as material deformation and failure (see Section 4.2), are inherently multiscale processes that cannot always be modeled via continuum model. Fully microscopic simulations of most domains of interest are not computationally feasible. Therefore, researchers must look at multiscale methods that couple micro models and macro models. Combining different scales

such as quantum-atomistic or atomistic, mesoscale and continuum, are still a challenge to obtain efficient and accurate schemes that efficiently and effectively exchange information between the different scales. We are currently involved in two national research projects (ANR), that focus on multiscale schemes. More precisely, the models that we start to study are the quantum to atomic coupling (QM/MM coupling) in the NOSSI ANR and the atomic to dislocation coupling in the OPTIDIS ANR (proposal for the 2010 COSINUS call of the French ANR).

3.5.2. Coupling of complex simulations based on the hypergraph model

The performance of the coupled codes depends on how well the data are distributed among the processors. Generally, the data distributions of each code are built independently from each other to obtain the best load-balancing. But once the codes are coupled, the naive use of these decompositions can lead to important imbalance in particular, when we have an overlap zone between the different models. Therefore, the modelling of the coupling itself is crucial to improve the performance and to ensure a good scalability of the coupled codes. The goal here is to find the best data distribution for the whole coupled code and not only for each standalone code. The main idea is to use an hypergraph model as the one provided by the ZOLTAN toolkit, and to take into account more information in the coupling than the classical one used by graph partitionner. Indeed, in the hypergraph model, the hyperedge cuts accurately measure communication volume, while, in the graph model, the edge cuts only approximate the communication volume. Moreover, recent works on hypergraph partitioning with fixed vertices have demonstrated their effectiveness for dynamic load balancing of adaptative simulations. As the load balancing problem is quite close to the redistribution one, we expect to provide new redistribution algorithm using similar strategies. For example, we should add in the communication cost the redistribution cost between codes (that depends on the volume of data exchanged); and we should add in the computation cost the interpolation cost, and so on. In addition, we expect the greater expressiveness of the hypergraph model help us to model each individual simulation code more accurately and thus enables us to improve their scalability thanks to a better partition quality. Another connected problem is the problem of resource allocation. This is particularly important for the global coupling efficiency and scalability, because each code involved in the coupling can be more or less computationally intensive, and there is a good trade-off to find between resources assigned to codes to avoid idle time. Typically, if we have a given number of processors and two coupled codes, how to split the processors among each code?

3.5.3. Steering and Interacting with complex coupled simulations

The computational steering is an effort to make the typical simulation work-flow (modelling, computing, analyzing) more efficient, by providing online visualization and interactive steering over the on-going computational processes. The online visualization appears very useful to monitor and to detect possible errors in long-running applications, and the interactive steering allows the researcher to alter simulation parameters on-the-fly and to immediately receive feedback on their effects. Thus, the scientist gains an additional insight in the simulation regarding to the cause-and-effect relationship.

In the ScAlApplix project, we have studied this problem in the case where both the simulation and the visualization can be parallel, what we call M-by-N computational steering, and we have developed a software environment called EPSN (see Section 5.3). More recently, we have proposed a model for the steering of complex coupled simulations and one important conclusion we have from these previous works is that the steering problem can be conveniently modeled as a coupling problem between one or more parallel simulation codes and one visualization code, that can be parallel as well. We propose in HiεPACS to revisit the steering problem as a coupling problem and we expect to reuse the new redistribution algorithms developed in the context of code coupling for the purpose of M-by-N steering.

In several applications, it is often very useful either to visualize the results of the ongoing simulation before writing it to disk, or to steer the simulation by modifying some parameters and visualize the impact of these modifications interactively. Nowadays, high performance computing simulations use many computing nodes, that perform I/O using the widely used HDF5 file format. One of the problems is now to use real-time visualization using high performance computing. In that respect we need to efficiently combine very large parallel simulation systems with parallel visualization systems. The originality of this approach is the

use of the HDF5 file format to write in a distributed shared memory (DSM); so that the data can be read from the upper part of the visualization pipeline. This leads to define a relevant steering model based on a DSM. It implies finding a way to write/read data efficiently in this DSM, and steer the simulation. This work is developed in collaboration with the Swiss National Supercomputing Centre (CSCS). As concerns the interaction aspect, we are interested in providing new mechanisms to interact with the simulation directly through the visualization. For instance in the ANR NOSSI, in order to speed up the computation we are interested in rotating a molecule in a cavity or in moving it from one cavity to another within the crystal lattice. To perform safely such interactions a model of the interaction in our steering framework is necessary to keep the data coherency in the simulation. Another point we plan to study is the monitoring and interaction with resources, in order to perform user-directed checkpoint/restart or user-directed load balancing at runtime.

4. Application Domains

4.1. Introduction

Currently, we have one major application which is material physics, and for which we contribute to all steps that go from modelling aspects to the design and the implementation of very efficient algorithms and codes for very large multi-scale simulations. Moreover, we apply our algorithmic research about linear algebra (see Section 3) in the context of several collaborations with industrial and academic partners. Our high performance libraries are or will be integrated in several complex codes and will be used and validated for very large simulations.

4.2. Material physics

Due to the increase of available computer power, new applications in nano science and physics appear such as study of properties of new materials (photovoltaic materials, bio- and environmental sensors, ...), failure in materials, nano-indentation. Chemists, physicists now commonly perform simulations in these fields. These computations simulate systems up to billion of atoms in materials, for large time scales up to several nanoseconds. The larger the simulation, the smaller the computational cost of the potential driving the phenomena, resulting in low precision results. So, if we need to increase the precision, there is two ways to decrease the computational cost. In the first approach, we improve classical methods and algorithms and in the second way, we will consider a multiscale approach.

Many applications in material physics need to couple several models like quantum mechanic and molecular mechanic models, or molecular and mesoscopic or continuum models. These couplings allow scientists to treat larger solids or molecules in their environment. Many of macroscopic phenomena in science depend on phenomena at smaller scales. Full simulations at the finest level are not computationally feasible in the whole material. Most of the time, the finest level is only necessary where the phenomenon 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.

We will focus on two applications; the first one concerns the computation of optical spectra of molecules or solids in their environment. In the second application, we will develop faster algorithms to obtain a better understanding of the metal plasticity, phenomenon governing by dislocation behavior. Moreover, we will focus on the improvement of the algorithms and the methods to build faster and more accurate simulations on modern massively parallel architectures.

4.2.1. Hybrid materials

There is current interest in hybrid pigments for cosmetics, phototherapy and paints. Hybrid materials, combining the properties of an inorganic host and the tailorable properties of organic guests, particularly dyes, are also of wide interest for environmental detection (oxygen sensors) and remediation (trapping and elimination of dyes in effluents, photosensitised production of reactive oxygen species for reduction of air and water borne contaminants). A thorough understanding of the factors determining the photo and chemical stability of hybrid pigments is thus mandated by health, environmental concerns and economic viability.

Many applications of hybrid materials in the field of optics exploit combinations of properties such as transparency, adhesion, barrier effect, corrosion, protection, easy tuning of the colour and refractive index, adjustable mechanical properties and decorative properties. It is remarkable that ancient pigments, such as Maya Blue and lacquers, fulfill a number of these properties. This is a key to the attractiveness of such materials. These materials are not simply physical mixtures, but should be thought of as either miscible organic and inorganic components, or as a heterogeneous system where at least one of the component exhibits a hierarchical order at the nanometer scale. The properties of such materials no longer derive from the sum of the individual contributions of both phases, since the organic/inorganic interface plays a major role. Either organic and inorganic components are embedded and only weak bonds (hydrogen, van der Waals, ionic bonds) give the structure its cohesion (class I) or covalent and ionic-covalent bonds govern the stability of the whole (class II).

These simulations are complex and costly and may involve several length scales, quantum effects, components of different kinds (mineral-organic, hydro-philic and -phobic parts). Computer simulation already contributes widely to the design of these materials, but current simulation packages do not provide several crucial functions, which would greatly enhance the scope and power of computer simulation in this field.

The computation of optical spectra of molecules and solids is the greatest use of the Time Dependent Density Functional Theory (TDDFT). We compute the ground state of the given system as the solution of the Kohn-Sham equations (DFT). Then, we compute the excited states of the quantum system under an external perturbation - electrical field of the environment - or thanks to the linear theory, we compute only the response function of the system. In fact, physicists are not only interested by the spectra for one conformation of the molecule, but by an average on its available configurations. To do that, they sample the trajectory of the system and then compute several hundred of optical spectra in one simulation. But, due to the size of interesting systems (several thousands of atoms) and even if we consider linear methods to solve the Kohn-Sham equations arising from the Density Functional Theory, we cannot compute all the system at this scale. In fact, such simulations are performed by coupling Quantum mechanics (QM) and Molecular mechanic (MM). A lot of works are done on the way to couple these two scales, but a lot of work remains in order to build efficient methods and efficient parallel couplings.

The most consuming time in such coupling is to compute optical spectra is the TDDFT. Unfortunately, examining optical excitations based on contemporary quantum mechanical methods can be especially challenging because accurate methods for structural energies, such as DFT, are often not well suited for excited state properties. This requires new methods designed for predicting excited states and new algorithms for implementing them. Several tracks will be investigated in the project:

- Typically physicists or chemists consider spectral functions to build a basis (orbital functions) and all the computations are performed in a spectral way. Due to our background, we want to develop new methods to solve the system in the real space by finite differences or by wavelets methods. The main expectation is to construct error estimates based on for instance the grid-size h parameter.
- For a given frequency in the optical spectra, we have to solve a symmetric non Hermitian system. With our knowledge on linear solvers, we think that we can improve the methods commonly used (Lanczos like) to solve the system (see Section 3.3).
- Improving the parallel coupling is crucial for large systems because the computational cost of the atomic and quantum models are really different. In parallel we have the following order of magnitude: one second or less per time step for the molecular dynamics, several minutes or more for

the DFT and the TDDFT. The challenge to find the best distribution in order to have the same CPU time per time step is really important to reach high performance. Another aspect in the coupling is the coupling with the visualization to obtain online visualization or steerable simulations. Such steerable simulations help the physicists to construct the system during the simulation process by moving one or a set of molecules. This kind of interaction is very challenging in terms of algorithmic and this is a good field for our software platform EPSN .

4.2.2. Material failures

Another domain of interest is the material aging for the nuclear industry. The materials are exposed to complex conditions due to the combination of thermo-mechanical loading, the effects of irradiation and the harsh operating environment. This operating regime makes experimentation extremely difficult and we must rely on multi-physics and multi-scale modelling for our understanding of how these materials behave in service. This fundamental understanding helps not only to ensure the longevity of existing nuclear reactors, but also to guide the development of new materials for 4th generation reactor programs and dedicated fusion reactors. For the study of crystalline materials, an important tool is dislocation dynamics (DD) modelling. This multiscale simulation method predicts the plastic response of a material from the underlying physics of dislocation motion. DD serves as a crucial link between the scale of molecular dynamics and macroscopic methods based on finite elements; it can be used to accurately describe the interactions of a small handful of dislocations, or equally well to investigate the global behavior of a massive collection of interacting defects.

To explore, i.e., to simulate these new areas, we need to develop and/or to improve significantly models, schemes and solvers used in the classical codes. In the project, we want to accelerate algorithms arising in those fields. We will focus on the following topics (in particular in the currently under definition OPTIDIS project in collaboration with CEA Saclay, CEA Ile-de-france and SIMaP Laboratory in Grenoble) in connection with research described at Sections 3.4 and 3.5.

- The interaction between dislocations is long ranged ($O(1/r)$) and anisotropic, leading to severe computational challenges for large-scale simulations. In dislocation codes, the computation of interaction forces between dislocations is still the most CPU time consuming and has to be improved to obtain faster and more accurate simulations.
- In such simulations, the number of dislocations grows while the phenomenon occurs and these dislocations are not uniformly distributed in the domain. This means that strategies to dynamically construct a good load balancing are crucial to achieve high performance.
- From a physical and a simulation point of view, it will be interesting to couple a molecular dynamics model (atomistic model) with a dislocation one (mesoscale model). In such three-dimensional coupling, the main difficulties are firstly to find and characterize a dislocation in the atomistic region, secondly to understand how we can transmit with consistency the information between the two micro and meso scales.

4.3. Application framework customers of high performance linear algebra solvers

We are currently collaborating with various research groups involved in geophysics, electromagnetics and structural mechanics. For all these application areas, the current bottleneck is the solution of huge sparse linear systems often involving multiple right-hand sides either available simultaneously or given in sequence. The robustness, efficiency and scalability of the numerical tools designed in Section 3.3 will be preliminary investigated in the parallel simulation codes of these partners.

More precisely, BRGM and TOTAL simulations require the solutions of huge linear systems with many right-hand sides given simultaneously. We notice that the collaborative work with TOTAL will also address the use of GPU for intensive numerical kernels in the Reverse Time Migration process for seismic imaging.

The CEA-CESTA simulation codes need the solution with simultaneous right-hand sides but also with right-hand sides given in sequence. The first situation arises in RCS calculations, but is generic in many parametric studies, while the second one comes from the nature of the solver that is based on a multiplicative Schwarz approach. The subproblems are solved several times in sequence. Many of the numerical approaches and possible outcoming software are well suited to tackle these challenging problems.

Research activities related to EDF and developed in the framework of the ANR SOLSTICE project have already stimulated interactions between members of the former ScAlApplix INRIA project team and members of the Parallel Algorithms team of CERFACS. These research activities have concerned direct and iterative solution methods for linear systems and eigenvalue computations. A major focus was on the efficient use of parallel sparse direct solution methods for large scale applications in structural mechanics in both in-core and out-of-core environments. These solution methods have been already integrated in the Code_Aster structural mechanics code developed at EDF. The use of hybrid solution methods will be investigated in structural mechanic applications and also in other different applications of interest for EDF such as neutronics or fluid mechanics.

On more academic sides, some ongoing collaborations with other INRIA EPs will be continued and others will be started. In collaboration with the NACHOS INRIA project team, we will continue to investigate the use of efficient linear solvers for the solution of the Maxwell equations in the time and frequency domains where discontinuous Galerkin discretizations are considered. Additional funding will be sought out in order to foster this research activity in connection with actions described in Section 3.3.

Jointly with the MAGIQUE3D INRIA project team, we intend to collaborate to design parallel efficient simulation codes for seismic wave propagation at the Earth scale where various huge linear systems have to be solved on large parallel platforms. The foreseen numerical techniques will be based on a mixed spectral finite element approach coupled with some boundary element techniques. The efficient solution of such problem will strongly rely on the activities described in Section 3.2 (e.g. complex load balancing problem) and in Section 3.3 (for the various parallel linear algebra kernels).

5. Software

5.1. Introduction

We describe in this section the software that we are developing. The first two (MaPHyS and EPSN) will be the main milestones of our project. The other software developments will be conducted in collaboration with academic partners or in collaboration with some industrial partners in the context of their private R&D or production activities. For all these software developments, we will use first the various (very) large parallel platforms available through CERFACS and GENCI in France (CCRT, CINES and IDRIS Computational Centers), and next the high-end parallel platforms that will be available via European and US initiatives or projects such that PRACE.

5.2. MaPHyS

MaPHyS (Massively Parallel Hybrid Solver) is a software package whose prototype was initially developed in the framework of the PhD thesis of Azzam Haidar (CERFACS) and further consolidated thanks to the ANR-CIS Solstice funding. This parallel linear solver couples direct and iterative approaches. The underlying idea is to apply to general unstructured linear systems domain decomposition ideas developed for the solution of linear systems arising from PDEs. The interface problem, associated with the so called Schur complement system, is solved using a block preconditioner with overlap between the blocks that is referred to as Algebraic Additive Schwarz. To cope with the possible lack of coarse grid mechanism that enables one to keep constant the number of iterations when the number of blocks is increased, the solver exploits two levels of parallelism (between the blocks and within the treatment of the blocks). This enables us to exploit a large number of processors with a moderate number of blocks which ensures a reasonable convergence behaviour.

The current prototype code will be further consolidated to end-up with a high performance software package to be made freely available to the scientific community. In that respect, an additional support has been obtained in the framework of the INRIA technologic development actions; 24 man-month engineer (Yohan Lee-Tin-Yien) have been allocated to this software activity. The roadmap for this software development is well defined; it should enable us to have interfaces with various graph partitioning packages at the end of the first year and interfaces with most of the parallel sparse direct solvers at the end of the second year. The MaPhyS package is very much a first outcome of the research activity described in Section 3.3. Finally, MaPhyS is a preconditioner that can be used to speed-up the convergence of any Krylov subspace method. We foresee to either embed in MaPhyS some Krylov solvers or to release them as standalone packages, in particular for the block variants that will be some outcome of the studies discussed in Section 3.3.

5.3. EPSN

EPSN (Environement for Computational Steering) is a software environment for the steering of parallel numerical simulations with visualization programs that can be parallel as well (see Figure 1). It is based on RedGRID, a software environment especially dedicated to the coupling of parallel codes, and more precisely to the redistribution of complex parallel data objects such as structured grids, particles and unstructured meshes.

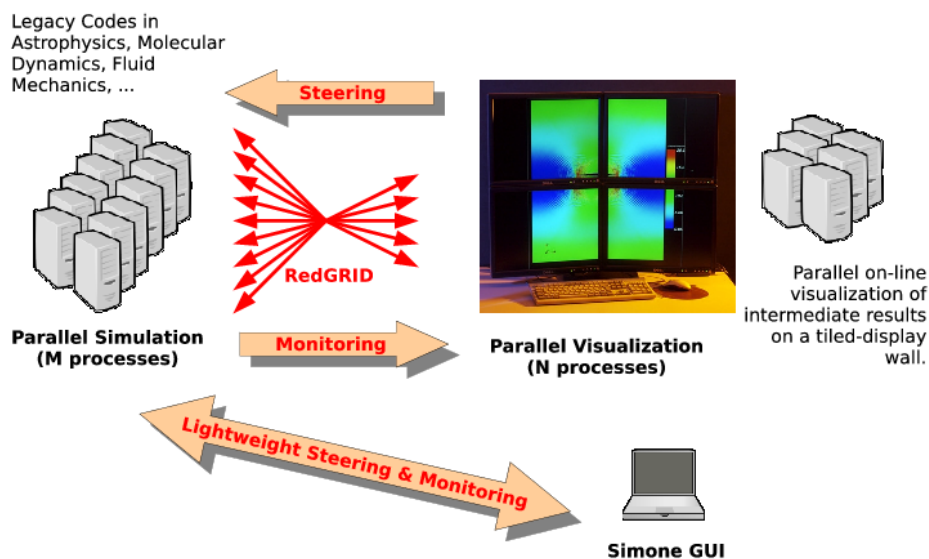


Figure 1. EPSN: software environment for $M \times N$ computational steering.

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 the integration of 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 us 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 us to reduce the steering overhead of the EPSN platform and allows us to process efficiently large dataset. It is also possible to exploit tiled-display wall with EPSN in order to reach high resolution image.

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, points or unstructured meshes. Both EPSN and RedGRID use a communication infrastructure based on CORBA which provides our platform with portability, interoperability and network transparency. EPSN has been supported by the ACI-GRID program (grant number PPL02-03), the ARC RedGRID, and more recently by the ANR program called MASSIM (grant number ANR-05-MMSA-0008-03). It is now involved in the ANR CIS NOSSI (2007). More informations are available on our web site: <http://www.labri.fr/projet/epsn>. This software is publicly available at INRIA Gforge (<http://gforge.inria.fr/projects/epsn>).

In the context of code coupling, it appears important to us to develop a simpler and more efficient way to couple codes than we have used in EPSN. Moreover, we also want to factorize different developments we have done in different projects such as MASSIM, NOSSI or within collaborations. That could be done through a new library in order to make easier the research, the development and the experimentation of the << new >> algorithms we are looking for, as described at section 3.5. The functionalities that we need are

- *a coupling layer*, that provides a simple API to easily couple several MPI-based codes, to deploy them, to interconnect them and optionnally to restart them for load balancing if it is useful;
- *a redistribution layer*, that acts as an algorithmic kernel, independant of the network technology. Based on the description of the distributed data involved in the coupling, the redistribution algorithm computes the data mapping between codes;
- *a communication layer*, that must efficiently perform the transfer according to the result provided by the redistribution layer;
- *a monitoring & steering layer*, that should incorporate previous works around M-by-N steering realized with EPSN.

This library will be based on the well-known MPI standard to obtain performance, and will fully exploit the new facilities provided by MPI-2. Indeed, the dynamic process management allowed by MPI-2 offers interesting possibility for the design of code coupling.

5.4. Other software

These software packages are or will be developed in collaboration with some academic partners (LIP6, LaBRI, CPMOH, IPREM, EPFL) or in collaboration with industrial partners (CEA, TOTAL, EDF) in the context of their private R&D or production activities.

- Fast Multipole with BLAS (FMB), developed in collaboration with P. Fortin (LIP6), is a high performance parallel implementation of the Fast Multipole Method for the Laplace equation. It is based on BLAS routines and on an hybrid MPI-Thread parallelization for both shared and distributed memory architectures (see Section 3.4).
- For the materials physics applications, a lot of development will be done in the context of ANR projects (NOSSI and proposal OPTIDIS, see Section 4.2) in collaboration with LaBRI, CPMOH, IPREM, EPFL and with CEA Saclay and Bruyère-le-Châtel.
- In the context of the PhD thesis of Mathieu Chanaud (collaboration with CEA/CESTA), we develop a new parallel platform based on a combination of a multigrid solver and a direct solver (the PaStiX solver developed in the previous ScAlAppIix project-team) to solve huge linear systems arising from Maxwell equations discretized with first-order Nédelec elements (see Section 3.3).
- Finally, we contribute to software developments for seismic analysis and imaging and for wave propagation in collaboration with TOTAL (use of GPU technology with CUDA in the context of the PhD thesis of Rached Abdelkhalek) and with BRGM (use of PaStiX and MaPhyS solvers in the context of the PhD of Fabrice Dupros in collaboration with Dimitri Komatitsch of MAGIQUE3D project team).

6. New Results

6.1. Algorithms and high-performance solvers

6.1.1. Hybrid direct/iterative solvers based on algebraic domain decomposition techniques

We have studied the parallel scalability of variants of an algebraic additive Schwarz preconditioner for the solution of large three dimensional convection diffusion problems in a non-overlapping domain decomposition framework. To alleviate the computational cost, both in terms of memory and floating-point complexity, we investigate variants based on a sparse approximation or on mixed 32- and 64-bit calculation. The robustness and the scalability of the preconditioners are investigated through extensive parallel experiments on up to two thousand processors. Their efficiency from a numerical and parallel performance view point are investigated. More detailed on this work can be found in [5].

Parallel numerical experiments with this solver were also performed in the framework of inverse problem in geophysics. Some analysis of the its behaviour with respect to other approaches for 3D simulations are reported in [6].

6.1.2. Hybrid solver based on a combination of a multigrid method and a direct method

In the context of a collaboration with the CEA/CESTA center, M. Chanaud continues a Ph.D. concerning a tight combination between multigrid methods and direct methods for the efficient solution of challenging 3D irregular finite element problems arising from the discretization of Maxwell or Helmholtz equations. A sequential prototype has been validated. A parallel solver dedicated to the ODYSSEE challenge (electromagnetism) of CEA/CESTA and the study of the numerical behaviour of this hybrid solver are ongoing activities.

6.1.3. Domain decomposition methods to solve neutron transport equations

This work, started with a collaboration between the EDF/SINETICS team and the former ScAlAppIix project, intended to design and develop techniques to optimize the efficiency of the codes used to simulate the physics of nuclear reactors. In the context of Bruno Lathuilère PhD (in collaboration with Pierre Ramet from BACCHUS), we have completed a study to parallelize a SPn simulation code by using a domain decomposition method applied for the solution of the neutron transport equations (Boltzmann equations). The defense of the thesis is planed at the beginning of February 2010.

6.1.4. High performance algorithms for wave propagation

A first work has been initiated during the ANR CIGC-05 NUMASIS project. The overall objective is the adaptation and the optimization of numerical methods in geophysics for large scale simulations on hierarchical and multicores architectures. Fabrice Dupros (BRGM) has started a PhD on these topics in February 2007 in the former ScAlApplix project. This work is also carried out in the framework of a collaboration with the INRIA MAGIQUE3D team (Dimitri Komatitsch) and BRGM. Several contributions can be underlined, for example the impact of the memory hierarchy for this class of simulations ([10], [14]). Large scale finite-elements computations for site effects in the French Riviera urban area have also been performed on the JADE GENCI/Cines platform using the PaStiX sparse parallel direct solver [3]. An ongoing topic is the evaluation of a spacetime decomposition for the time-domain finite-differences method (FDTD) and its application to the classical staggered-grid scheme [17]. The defense of this PhD is planned during the first semester of 2010.

A second work is currently carried on with TOTAL. The extraordinary challenge that the oil and gas industry must face for hydrocarbon exploration requires the development of leading edge technologies to recover an accurate representation of the subsurface. Seismic modeling and Reverse Time Migration (RTM) based on the full wave equation discretization, are tools of major importance since they give an accurate representation of complex wave propagation areas. Unfortunately, they are highly compute intensive. The recent development in GPU technologies with unified architecture and general-purpose languages coupled with the high and rapidly increasing performance throughput of these components made General Purpose Processing on Graphics Processing Units an attractive solution to speed up diverse applications. We have designed a fast parallel simulator that solves the acoustic wave equation on a GPU cluster ([1], [15]). Solving the acoustic wave equation in an oil exploration industrial context aims at speeding up seismic modeling and Reverse Time Migration. We consider a finite difference approach on a regular mesh, in both 2D and 3D cases. The acoustic wave equation is solved in a constant density or a variable density domain. All the computations are done in single precision, since double precision is not required in our context. We use nvidia CUDA to take advantage of the GPU computational power. We study different implementations and their impact on the application performance. We obtain a speed up of 16 for Reverse Time Migration and up to 43 for the modeling application over a sequential code running on general purpose CPU.

6.2. Efficient algorithmics for code coupling in complex simulations

In the context of load-balancing for complex code-coupling simulations, we have studied the particular case of a multiscale simulation, called LibMultiscale previously developed in the ScAlApplix project. More precisely, LibMultiScale simulates a crack propagation by using two parallel codes with different scales: a molecular dynamic code (atomistic model) and an elasticity code discretized by a finite element method (continuum model). The experience we have acquired with LibMultiScale shows us that the load-balancing of the whole coupled simulation is a difficult issue, that cannot be simply solved by partitioning each code independently from each other. Such a naive approach can lead to important imbalance in particular in the overlap zone where the two models coexist. Indeed, only a few processes were involved during the coupling phase, and the resulting communication pattern between the coupled codes were often unbalanced.

To overcome these difficulties, we have introduced a modelling of the whole coupled simulation based on the hypergraph model, that uses the notion of coupling hyperedge. The key idea is to provide a coupling-aware partitioning strategy based upon this model. One has evaluated different strategies for 2D cases of LibMultiScale. One of the most interesting strategy consists in partitioning the overlap area first (using all processes) and then, using the resulting fixed vertices, to partition the remaining model. As a result, both the computational phase and the coupling phase are well balanced. Moreover, the use of some special hyperedges has been used to reduce the number of coupling communications. These preliminary works has been realized during the Master internship of Mohamed Amine El Afrit.

6.3. Computational steering environment for distributed numerical simulations

6.3.1. 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), which affords to solve the coherency problem for such complex applications. The EPSN framework has been extended to handle this new kind of simulations. In the context of the ANR MASSIM and ANR NOSSI, 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. This work is done in the context of the PhD of Nicolas Richart and the defense is planned at the beginning of 2010.

6.3.2. Distributed Shared Memory approach for the steering of parallel simulations

As a different approach of the in-situ and steering framework of EPSN, we conceived and developed a light push-driven architecture for in-situ visualization. The architecture, part of ICARUS, is intended to address three principal objectives: Require little or no modification to the simulation code in order to allow a live visualization. Allow the simulation to be run on one parallel machine whilst the visualization is run on a separate (or the same) parallel machine. Provide good performance to ensure that massive simulations may be handled as easily as small test cases. The interface developed is built around the HDF5 file I/O library used commonly in HPC applications. The HDF5 API allows the derivation of custom virtual file drivers (VFDs) which may be instantiated at run-time on a per file basis to control how data is written to the file system. We have made use of this facility to create a specialized MPI based VFD which allows the simulation to write data in parallel to a file, which is actually redirected over the network to a visualization cluster which in turn stores the file in a Distributed Shared Memory (DSM) buffer - or in effect a virtual file system. The ParaView application acts as a server/host for this DSM and can read the file contents directly using the HDF5 API as if reading from disk. The transfer of data between simulation and visualization machines may be done using either an MPI based communicator shared between the applications, or using a socket based communication. The management of both ends of the network transfer is transparently handled by our DSM VFD layer, meaning that an application using HDF5 can make use of in-situ visualization without any code changes. It is only necessary to re-link the application against a modified version of the HDF library which contains our driver. This work has been made and is currently carrying on at CSCS - Swiss National Supercomputing Centre, under the co-supervision of Mr. John Biddiscombe, within the NextMuSE European project 7th FWP/ICT-2007.8.0 FET Open.

6.4. Material physics

6.4.1. Hybrid materials

The study of hybrid materials with a coupling method between molecular dynamics (MD) and quantum mechanism (QM) has begun in collaboration with IPREM (Pau) in the ANR CIS 2007 NOSSI. These simulations are complex and costly and may involve several length scales, quantum effects, components of different kinds (mineral-organic, hydro-philic and -phobic parts). Our goal is to compute dynamical properties of hybrid materials like optical spectra. The computation of optical spectra of molecules and solids is the most consuming time in such coupling. This requires new methods designed for predicting excited states and new algorithms for implementing them. Several tracks are investigated in the project and new results obtained as described bellow.

Optical spectra. Theory of electronic excitations contributes to our understanding of photovoltaic devices, dyes and photo-catalytic materials. Recent advances in polymer semiconductors pose new challenges for theoretical predictions because the size of the system and the absence of any symmetry in molecules of interest. Ab-initio approaches to the quantum theory of large molecules (hundreds of atoms) is often limited to density functional theory and its time-dependent counterpart (TDDFT) because these approaches allow for a

favorable complexity scaling in practical computations. An essential ingredient of TDDFT—the Kohn–Sham response function—has a simple expression in terms of molecular orbitals, although disregarding its inherent locality leads to a poor complexity scaling $O(N^4)$ with the number of atoms N . Moreover, the applications of Kohn–Sham response function go beyond TDDFT, therefore we have been further developing of a fast method for the Kohn–Sham response function. The fast Kohn–Sham response function allows for a favorable $O(N^2)$ complexity scaling ([4], [12]). Our implementation of the algorithm had been optimized and parallelized with the shared-memory approach ([12], [24]). This work allowed us to compute the response function and corresponding absorption spectra of fullerene C_{60} . The overall complexity for the absorption spectra of $O(N^2)$ has been achieved due to usage of an iterative method (bi-orthogonal Lanczos or GMRES) [13]. An article is in preparation. Future work on Kohn–Sham response function shall be concentrated on distributed memory parallelization because of high memory demand of the response function. The symmetry properties of the response function must be necessarily exploited in the MPI parallelized version of the program.

QM/MM algorithm. For structure studies or dynamical properties, we intend to couple QM model based on pseudo-potentials (SIESTA code) with dynamic molecular (DL-POLY code). Therefore we have developed first a new algorithm to avoid to count twice the quantum electric field in the molecular model. Secondly, we have introduced an algorithm to compute faster the electric field which polarizes the quantum atoms. Now, we are implementing our algorithm in Siesta and DL-POLY codes.

7. Contracts and Grants with Industry

7.1. Research and development contracts

CEA research and development contracts:

- Conception of an hybrid solver combining multigrid and direct methods (Mathieu Chanaud (PhD); David Goudin and Jean-Jacques Pesqué from CEA-CESTA; Luc Giraud, Jean Roman).

EDF research and development contract:

- Application of a domain decomposition method to the neutronic SPn equations (Bruno Lathuilière (PhD) and Pierre Ramet from BACCHUS INRIA project team; Jean Roman).

TOTAL research and development contract:

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

8. Other Grants and Activities

8.1. National initiatives

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

Participants: Olivier Coulaud, Aurélien Esnard, Damien Genet.

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 design 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>

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

Participants: Abdou Guermouche, 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 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/>

8.1.3. SOLSTICE: High performance solvers for frontier simulations

Participants: Abdou Guermouche, Luc Giraud, Jean Roman.

Grant: ANR-06-CIS

Dates: 2006 – 2010

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/>

8.2. International initiatives

8.2.1. Associated team PHyLeaS

Participants: Olivier Coulaud, Luc Giraud, Jean Roman.

Grant: INRIA

Dates: 2008-2009

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-sop.inria.fr/nachos/phyleas/>

9. Dissemination

9.1. Participation to the Scientific Community

Olivier Coulaud has been member of the INRIA COST GTAI (in charge of incentive actions).

Luc Giraud has been member of the scientific committee of the international conferences PARENG'09, PDSEC'09, SC'09 and Precond'09.

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 EuroMicroPDP'09 (IEEE), PARCO'09 and of the national conference Renpar'09. He is member of the "Strategic Comity for Intensive Computation" of the French Research Ministry and is member of the "Scientific Board" of the CEA-DAM.

9.2. Teaching

In complement of the normal teaching activity of the university members and of ENSEIRB-MATMECA members, Olivier Coulaud teaches at ENSEIRB-MATMECA and Luc Giraud teaches at ENSEEIHT and ISAE-ENSICA.

10. Bibliography

Major publications by the team in recent years

- [1] R. ABDELKHALEK, H. CALANDRA, O. COULAUD, J. ROMAN, G. LATU. *Fast Seismic Modeling and Reverse Time Migration on a GPU Cluster*, in "The 2009 High Performance Computing & Simulation - HPCS'09, Allemagne Leipzig", 2009.

- [2] O. COULAUD, P. FORTIN, J. ROMAN. *High-performance BLAS formulation of the Adaptive Fast Multipole Method*, in "Mathematical and Computer Modelling", vol. 51, n^o 3-4, 2009, p. 177-188.
- [3] F. DUPROS, F. DE MARTIN, E. FOERSTER, D. KOMATITSCH, J. ROMAN. *High-performance finite-element simulations of seismic wave propagation in three dimensional non linear in-elastic geological media*, in "Parallel Computing", 2010, to appear.
- [4] D. FOERSTER, P. KOVAL. *On the Kohn–Sham density response in a localized basis set*, in "The Journal of Chemical Physics", vol. 131, n^o 4, 2009, <http://hal.archives-ouvertes.fr/hal-00378987/en/>.
- [5] L. GIRAUD, A. HAIDAR. *Parallel algebraic hybrid solvers for large 3D convection-diffusion problems*, in "Numerical Algorithms", vol. 51, n^o 2, 2009, p. 151-177.

Year Publications

Articles in International Peer-Reviewed Journal

- [6] J. VIRIEUX, S. OPERTO, H. BEN-HADJ-ALI, R. BROSSIER, V. ETIENNE, F. SOURBIER, L. GIRAUD, A. HAIDAR. *Seismic wave modeling for seismic imaging*, in "The Leading Edge", vol. 25, n^o 8, 2009, p. 538-544.

Invited Conferences

- [7] A. HAIDAR, L. GIRAUD, H. BEN-HADJ-ALI, S. OPERTO, J. VIRIEUX. *Parallel Algebraic Hybrid Linear Solver for Frequency Domain Acoustic Wave Modeling*, in "SIAM Conference on Applied Linear Algebra SIAM Conference on Applied Linear Algebra, Monterey, California, USA", 2009, Minisymposium invited talk.
- [8] A. HAIDAR, L. GIRAUD. *Massively parallel linear solver for large 3D problems*, in "International Conference On Preconditioning Techniques For Large Sparse Matrix Problems In Scientific And Industrial Applications, Hong-Kong", 2009, Plenary invited talk.

International Peer-Reviewed Conference/Proceedings

- [9] 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.
- [10] F. DUPROS, C. POUSA, A. CARISSIMI, J.-F. MÉHAUT. *Parallel simulations of seismic wave propagation on NUMA architectures*, in "International Conference on Parallel Computing (PARCO'09), Lyon, France", 2009.
- [11] P. KOVAL, D. FOERSTER, O. COULAUD. *Fast construction of the Kohn–Sham response function for molecules*, in "Trends in nanotechnology - TNT 2009, Barcelona Espagne", 2009, <http://hal.inria.fr/inria-00437603/en/>.
- [12] P. KOVAL, D. FOERSTER. *A simple extension of the LCAO method to excited states*, in "Trends in nanotechnology - TNT 2009, Espagne Barcelona", 2009.

[13] P. KOVAL, O. COULAUD, D. FOERSTER. *Extension of LCAO to excited states: non interacting response and molecular spectra*, in "ETSF 2009: Ab-initio tools for the characterization of nanostructures, Evora Portugal", 2009, <http://hal.inria.fr/inria-00437608/PDF/prod-basis-etsf-evora-2009.pdf>.

[14] C. POUSA, M. CASTRO, F. DUPROS, A. CARISSIMI, L. G. FERNANDES, J.-F. MÉHAUT. *High-performance Applications on Hierarchical Shared Memory Multiprocessors*, in "Colloquium of computation: Brazil / INRIA, Cooperations, Advances and Challenges, SBC 2009, Bento-Goncalves, Brazil", 2009, p. 55-60.

Workshops without Proceedings

[15] R. ABDELKHALEK, H. CALANDRA, O. COULAUD, G. LATU, J. ROMAN. *FDTD Based Seismic Modeling and Reverse Time Migration on a GPU Cluster*, in "9th International Conference on Mathematical and Numerical Aspects of Waves Propagation - Waves 2009, France Pau", 2009.

[16] M. CHANAUD, D. GOUDIN, J.-J. PESQUÉ, J. ROMAN. *A dedicated method for solving large linear systems arising from finite element discretizations of high-frequency Maxwell problems*, in "Fourteenth Copper Mountain Conference on Multigrid Methods", 2009.

[17] F. DUPROS, A. DUCELLIER, H. AOCHI, D. KOMATITSCH, J. ROMAN. *Three-dimensional parallel finite-difference modeling of seismic wave propagation on multicore architectures: Application to Nice, France, urban*, in "9th International Conference on Mathematical and Numerical Aspects of Waves Propagation - Waves 2009, France Pau", 2009.

[18] A. HAIDAR, L. GIRAUD, H. BEN-HADJ-ALI, S. OPERTO, J. VIRIEUX. *On the use of a parallel algebraic hybrid linear solver for frequency domain acoustic wave modeling*, in "9th International Conference on Mathematical and Numerical Aspects of Waves Propagation - Waves 2009, France Pau", 2009.

Other Publications

[19] E. AGULLO, A. GUERMOUCHE, J.-Y. L'EXCELLENT. *Reducing the I/O Volume in Sparse Out-of-core Multifrontal Methods*, 2010, SIAM Journal on Scientific Computing, to appear.

[20] P. AMESTOY, I. S. DUFF, A. GUERMOUCHE, T. SLAVOVA. *Analysis of the Solution Phase of a Parallel Multifrontal Approach*, 2010, Parallel Computing, to appear.

[21] M. CHANAUD, D. GOUDIN, J.-J. PESQUÉ, J. ROMAN. *Parallel solver coupling multigrid and direct methods for Maxwell equations*, 2010, SIAM Conference on Parallel Processing for Scientific Computing (PP10), February 24-26, Seattle.

[22] L. GIRAUD, A. HAIDAR, S. PRALET. *Using multiple levels of parallelism to enhance the performance of domain decomposition solvers*, 2010, Parallel Computing, to appear.

[23] L. GIRAUD, A. HAIDAR, Y. SAAD. *Sparse approximations of the Schur complement for parallel algebraic hybrid linear solvers in 3D*, 2010, Numerical Mathematics: Theory, Methods and Applications, to appear US

[24] P. KOVAL, D. FOERSTER, O. COULAUD. *Fast construction of the Kohn–Sham response function for molecules*, Wiley, 2010, <http://hal.inria.fr/inria-00457652/en/>, Physica Status Solidi B, to appear.