First Brazil-France workshop on
High performance computing and scientific data management driven by highly demanding applications

INRIA Sophia Antipolis – Méditerranée research center
July 25 – 27, 2012

Scientific context

This workshop is organized in the context of the HOSCAR collaborative project between CNPq and INRIA. The general objective of the HOSCAR project is to setup a multidisciplinary Brazil-France collaborative effort for taking full benefits of future high-performance massively parallel architectures. The targets are the very large-scale datasets and numerical simulations relevant to a selected set of 5 applications in natural sciences: (i) Resource prospection, (ii) Reservoir simulation, (iii) Ecological modeling, (iv) Astronomy data management, and (v) Simulation data management. The project involves computer scientists and numerical mathematicians divided in 3 fundamental research groups: (i) Numerical schemes for PDE models (Group 1), (ii) Scientific data management (Group 2), and (iii) High-performance software systems (Group 3). The project also proposes the integration of the scientific results produced by the project within a common, user-friendly computational platform deployed over the partners' HPC facilities and tailored to the aforementioned applications.

Objectives

This first meeting between the Brazilian and French partners of the HOSCAR project will concern the research and development activities planned in Group 1: Numerical schemes for PDE models. The primary objective of the workshop is to discuss about the research and development activities of the researchers (applied mathematicians and computer scientists) participating in this group that will be conducted in the HOSCAR project. A secondary objective is to prepare the official kickoff of the project and in particular, the organization of the larger workshop that will take place on September 10-14, 2012 in Petropolis.

The Embassy of France in Brazil, INRIA, funds this workshop for what concern the participation of the French partners. The Brazilian partners also benefit from partial funding from the Embassy of France in Brazil.

Address and practical details

The workshop will take place at the INRIA Sophia Antipolis – Méditerranée research center. The research center is about 20 km away from the Nice Côte d'Azur International Airport,

http://en.nice.aeroport.fr

Details on how to reach the research center can be found using the following link,

http://www.inria.fr/en/centre/sophia/overview/how-to-reach-the-centre

Day 1: HPC numerical solvers and tools:

- scalable linear system solvers,
- scalable mesh generation,
- parallel graph remapping and parallel graph repartitioning.

Day 2: Numerical schemes for PDE models with applications to resource prospection: we will discuss here about numerical schemes for related wave propagation problems.

Day 3: Numerical schemes for PDE models with applications to reservoir simulation: we will discuss here about numerical schemes for related fluid flow problems.

Participants

Hélène Barucq, researcher, INRIA Bordeaux – Sud-Ouest
Lionel Boillot, PhD student, INRIA Bordeaux – Sud-Ouest
Konstantin Brenner, postdoctoral student, CEMEF, Sophia Antipolis
Juliette Chabassier, postdoctoral student, INRIA Bordeaux – Sud-Ouest
Alvaro Coutinho, researcher, COPPE/UFRJ
Florian Danard, intern, University of Nice–Sophia Antipolis
Cindy Guichard, postdoctoral student, University of Nice–Sophia Antipolis and INRIA Sophia Antipolis–Méditerranée
Clément Durochat, PhD student, INRIA Sophia Antipolis – Méditerranée
Luc Giraud, researcher, University of Bordeaux and INRIA Bordeaux – Sud-Ouest
Jérôme Jaffré, researcher, INRIA Paris - Rocquencourt
Cédric Lachat, PhD student, INRIA Bordeaux – Sud-Ouest et INRIA Sophia Antipolis – Méditerranée
Stéphane Lanteri, researcher, INRIA Sophia Antipolis – Méditerranée
Alexandre Madureira, researcher, LNCC
Roland Masson, professor, University of Nice–Sophia Antipolis and INRIA Sophia Antipolis–Méditerranée
François Pellegrini, professor, University of Bordeaux and INRIA Bordeaux – Sud-Ouest
Fabien Peyrusse, PhD student, INRIA Sophia Antipolis – Méditerranée
Pierre Ramet, associate professor, University of Bordeaux and INRIA Bordeaux – Sud-Ouest
Jean Roberts, researcher, INRIA Paris - Rocquencourt
Jean Roman, INRIA Bordeaux – Sud-Ouest
Claire Scheid, associate professor, University of Nice–Sophia Antipolis and INRIA Sophia Antipolis–Méditerranée
Souhila Sabit, INRIA Rennes – Bretagne Atlantique
Frédéric Valentin, researcher, LNCC
Mawussi Zounon, INRIA Bordeaux – Sud-Ouest

Day 1: HPC numerical solvers and tools
Wednesday, July 25

10h00 – 10h15
Welcome talk, introduction of the scientific collaborative context
Stéphane Lanteri

10h15 – 11h00
Alvaro Coutinho
Recent advances in high performance algorithms for computational science and engineering

In this talk we review some recent advances in high performance algorithms for relevant computational science and engineering applications. Of particular interest here are advanced nonlinear equation solvers, mesh generation and adaptive refinement and coarsening. We will discuss their implementation in high performance clusters and topics such as efficient data structures, partitioning strategies and related aspects. Multiphysics applications, as found in geoscience and offshore engineering will be used as demonstration problems. As an example of big-data issues, we will briefly review the use of parallel scientific workflow management tools as enabling technologies to support uncertainty quantification.

11h00 – 11h45
Pierre Ramet
Sparse direct solvers with accelerators over DAG runtimes

The current trend in the high performance computing shows a dramatic increase in the number of cores on the shared memory compute nodes. Algorithms, especially those related to linear algebra, need to be adapted to these new computer architectures in order to be efficient. PaStiX is a sparse parallel direct solver, which incorporates a dynamic scheduler for strongly hierarchical modern architectures. In this work, we study the replacement of this internal highly integrated scheduling strategy by two generic runtime frameworks: DAGuE and StarPU. Those runtimes will give the opportunity to execute the factorization tasks graph on emerging computers equipped with accelerators. As for previous work done in dense linear algebra, we will present the kernels used for GPU computations inspired by the MAGMA library and the DAG algorithm used with those two runtimes. A comparative study of the performances of the supernodal solver with the three different schedulers is performed on manycore architectures and the improvements obtained with accelerators will be presented with the StarPU runtime. These results demonstrate that these DAG runtimes provide uniform programming interfaces to obtain high performance on different architectures on irregular problems as sparse direct factorizations.

11h45 – 12h30
Cédric Lachat and François Pellegrini
A short tour of PaMPA 0.1

This talk will present the structure and operations of PaMPA (Parallel Mesh Partitioning and Adaptation), a middleware library dedicated to the management of unstructured meshes distributed across the processors of a parallel machine. Its purpose is to relieve solver writers from the tedious and error prone task of writing again and again service routines for mesh handling, data communication and exchange, remeshing, and data redistribution.

12h45 – 14h30 Lunch

14h30 – 15h15
Luc Giraud
Parallel algebraic domain decomposition linear solvers

The solution of linear systems is often the most computational consuming kernel in large complex numerical simulations. In this talk, we will describe a parallel algebraic hierarchical linear solver for sparse linear systems. The numerical scheme based on a partition of the adjacency graph of a sparse matrix, that leads to the solution of a Schur complement system, will be presented as well as the related preconditioning technique. Parallel numerical experiments of the hybrid direct/iterative technique will be described on 3D examples from both academic and industrial relevance.

15h15 – 16h00
Mawussi Zounon
Towards resilient algorithms in numerical linear algebra

The advent of exascale machines will require the use of parallel resources at an unprecedented scale, leading to a high rate of hardware faults. High Performance Computing applications that aim at exploiting all these resources will thus need to be resilient, i.e., being able to compute a correct output in presence of faults. Contrary to checkpointing techniques or Algorithm Based Fault Tolerant (ABFT) mechanisms, strategies based on interpolation for recovering lost data do not require extra work or memory when no fault occurs. We apply this latter strategy to Krylov iterative solvers, which are often the most computational intensive kernels in HPC simulation codes. Our main contribution is the proposition and discussion of several variants compared to previous works. For that, we propose a new variant for recovering data, we study the occurrence of multiple faults, we consider the GMRES, CG and BICGSTAB solvers, and we inject faults according to an advanced model of fault distribution. We assess the impact of the recovery method, the fault rate and the number of processors on resilience. Rather than implementing a particular actual parallel code, we assess all our strategies based on sequential Matlab implementations that simulate parallel executions.

16h00 – 16h45
François Pellegrini and Sébastien Fourestier
Current challenges for parallel graph (re)partitioning and (re)mapping

Graph partitioning is an ubiquitous technique which has applications in many fields of computer science and engineering. Its use in HPC includes the computation of partitions for domain decomposition methods and the ordering of sparse matrices for linear system solving. Scotch and its offspring PT-Scotch are software tools dedicated to the computation of high quality graph partitions. Scotch is a robust sequential tool that has been developed for almost 20 years. Its parallel offspring PT-Scotch, which aims at providing the same features in parallel, can partition graphs with sizes up to several billion vertices, distributed over thousands of processors. The advent of massively parallel, NUMA machines, represents a new challenge for software designers, in order for partitioning tools to scale up to hundred thousands of processing elements. The purpose of this talk is to present the key issues which we are considering, within the Scotch project, for the development of the next generation of scalable parallel graph (re)partitioning and (re)mapping algorithms.

16h45 – 17h30 Discussion

20h30 Diner at a restaurant in Antibes or nearby

Day 2: Numerical schemes for PDE models with applications to resource prospection and radioactive waste management
Thursday, July 26

10h00 – 10h45
Jocelyne Erhel, Edouard Canot and **Souhila Sabit**
Numerical methods of reactive transport: global approach DAE

Modeling reactive transport of contaminants in porous media is a complex time-dependent problem, due to combining the difficulties of modeling transport and chemistry, especially the coupling between them. In this work, we are interested to solve this type of coupling. Several methods have been developed for the resolution for solving this type problem. We choose to solve this problem by a global approach, which considers all the equations as a whole system of differential algebraic equations (DAE), which come from the spatial-only discretization of the equations (method of lines). This approach uses implicit schemes, which imply solving many large linear systems with the Jacobian matrix. The differential algebraic system (DAE) is solved by the solver IDA Sundials. Our new technique is implemented in the GRT3D software; we have observed that the CPU time increases very fast with the size of the system. Our aim is thus to reduce this computation time. Profiling tools have shown that an important part of this computation is due to the linear solving related to the Jacobian matrix. We focus our effort on improving this part, by exploiting the 3x3 block-structure of the Jacobian matrix, via a Gaussian block elimination technique. Our simulations are performed on academic test cases, which involve few chemical components (4 to 5) for both 1D and 2D geometries, giving a number of unknowns up to 72000. First results have shown that our technique is very promising, because the CPU time is reduced by approximately 40%. After this part, we eliminated the tracer in our test cases. In GRT3D-SL software, we calculated the concentrations directly without using the Logarithms and with this software, we have reduced the CPU time to 50 %.

10h45 – 11h30
Clément Durochat, Stéphane Lanteri and Claire Scheid
High order non-conforming multi-element discontinuous Galerkin method for time-domain electromagnetics

Nowadays, a variety of methods exist for the solution of the time domain Maxwell equations ranging from the well-established FDTD (Finite Difference Time-Domain) methods to FETD (Finite Element) methods and DGTD (Discontinuous Galerkin) methods. The recent years have witnessed an increased interest in DGTD methods, which have been developed on quadrangular (2D case) or hexahedral (3D case) meshes, as well as on triangular (2D case) or tetrahedral (3D case) meshes. One of the main features of DGTD methods is their enhanced flexibility with regards to the type of meshes they can deal with. Indeed, DGTD methods can easily handle irregular possibly non-conforming meshes, constituted of elements of various types and shapes. In this talk, we are concerned with the study of a DGTD method formulated on non-conforming and hybrid quadrangular-triangular (2D case) or hexahedral-tetrahedral (3D case) meshes for the solution of the time-domain Maxwell equations. Within each mesh element, the electromagnetic field components are approximated by a high order nodal polynomial, using a centered approximation for the surface integrals. Time integration of the associated semi-discrete equations is achieved by a second or fourth order Leap-Frog scheme. We prove the L2 stability (3D) of this method by establishing the conservation of a discrete analog of the electromagnetic energy and a sufficient CFL-like stability condition is exhibited. The theoretical convergence of the scheme is also studied; this leads to a-priori error estimate that takes into account the hybrid nature of the mesh.

11h30 – 12h15
Fabien Peyrusse, Nathalie Glinsky and Stéphane Lanteri
Discontinuous Galerkin method for wave propagation in viscoelastic media

We present a Discontinuous Galerkin method to solve the viscoelastic first-order velocity-stress system. Compared to the usual elastic system, it contains additional equations for the anelastic functions accouting for the strain history of the material. Numerical tests on 2D unstructured meshes with realistic topography are performed to demonstrate the accuracy of the scheme.

12h15 – 12h45 Discussion

12h45 – 14h30 Lunch

14h30 – 15h15
Jérôme Jaffré and Jean Robert
Nonlinear complementarity constraints for hydrogen-water flow with hydrogen dissolution in porous media

The migration of hydrogen produced by the corrosion of nuclear waste packages in an underground storage including the dissolution of hydrogen can be modeled by a set of nonlinear partial differential equations with nonlinear complementarity constraints. A modern and efficient solution strategy, the Newton-min method, is applied to this geoscience problem. We investigate its applicability and efficiency from both the theoretical and the numerical points of view for solving linear and nonlinear complementarity problems. Numerical and theoretical results will be presented to show the ability of our solver.

15h15 – 15h35
Lionel Boillot
Absorbing boundary condition for anisotropic elastodynamic media

Reverse Time Migration (RTM) technique produces images thanks to the propagation of waves. We focus on Tilted Transverse Isotropy (TTI) media which come into the category of anisotropic media. The finite element solution of this problem requires the design of efficient boundary conditions (BC) that are able to attenuate possible artificial reflected waves generated by the boundaries of the computational domain. In the case of elastic waves, Perfectly Matched Layers (PML) are widely used. However, instabilities may appear in the layers, either due to the numerical discretization or to the continuous PML problem as for the TTI case. RTM produces images thanks to the cross-correlation of the source propagation. This process can consider artificial reflected waves as noise which is not taken into account by the imaging condition. That is why we propose to use high-order Absorbing Boundary Conditions (ABC) and to study how they perform in the RTM framework. Moreover, they are easily included in a high-order Discontinuous Galerkin formulation. Numerical experiments of RTM will be performed to see how the ABCs impact on the accuracy of the images.

15h35 – 16h20
Juliette Chabassier
Stability and dispersion analysis of improved time discretization for prestressed Timoshenko systems. Application to the stiff piano string

Piano strings can be modeled as simply supported Timoshenko prestressed beams. This model accounts for inharmonicity of the tranversal displacement, via a coupling with a shear angle resulting in the propagation of flexural and shear waves with highly contrasted speeds. Our concern is to develop a new implicit time discretization, which will be associated with finite element methods in space, in order to reduce numerical dispersion while allowing the use of a large time step. After showing the continuous system and its two branches of eigenfrequencies, the classical θ-schemes will be presented. We will give some insights of innovative proofs of stability by energy techniques that provide uniform results with respect to the time step. A dispersion analysis will recall that $\theta=1/12$ reduces numerical dispersion, but the associated stability condition will prove very severe for our application case. We will propose a new θ-scheme based on different θ-approximations for the flexural and shear terms of the equations, which will allow to reduce numerical dispersion while relaxing the stability condition. Stability proofs can be extended for this new scheme. Theoretical results will be illustrated with numerical experiments corresponding to the simulation of a realistic piano string.

16h20 – 17h30 Discussion

Day 3: Numerical schemes for PDE models with applications to reservoir simulation
Friday, July 27

10h00 – 10h45
Rodolfo Araya, Christopher Harder, Diego Paredes and **Frédéric Valentin**
Multiscale hybrid-mixed methods for porous media flows

This work proposes a new family of finite element methods for porous media flows, named Multiscale Hybrid-Mixed (MHM) methods. The MHM method is a consequence of a hybridization procedure, and emerges as a method that naturally incorporates multiple scales while provides solutions with high-order precision for the primal and dual (or flux) variables. The computation of local problems is embedded in the upscaling procedure, which are completely independent and thus may be naturally obtained using parallel computation facilities. Also interesting is that the flux variable preserves the local conservation property using a simple post-processing of the primal variable. The general framework is illustrated for the Darcy equation, and further extended to other operators taking part in the modeling of porous media flows (advection-diffusion and elasticity equations, for instance). The analysis results in a priori estimates showing optimal convergence in natural norms and providing a face-based a posteriori estimator. Regarding the latter, we prove that reliability and efficiency hold. Numerical results verify the optimal convergence properties as well as a capacity to accurately incorporate heterogeneity and high-contrast coefficients, showing in particular the great performance of the new a posteriori error estimator in driving mesh adaptativity. We conclude that the MHM method, along with its associated a posteriori estimator, is naturally shaped to be used in parallel computing environments and appears to be a highly competitive option to handle realistic multiscale boundary value problems with precision on coarse meshes.

10h45 – 11h30
Manuel Barreda, Ana Carolina Carius and **Alexandre L. Madureira**
RFB methods: a consistent framework to derive multiscale finite element methods

We discuss some finite element techniques that are suitable to deal with problems of multiscale type, and are prone to parallel implementations. In particular, we discuss Residual Free Bubble (RFB) methods, and as an application, we look into some equations that result from dimension reduction techniques (the art of deriving 2D equations from 3D problems in slender domains). We also propose a RFB type method for a nonlinear, multiscale problem. We derive the method, and present some possible alternatives, depending on linearization choices. We show some error estimates under the assumption that the coefficients are highly oscillatory. The analysis is quite delicate due to the presence of the small-scale parameter. The method itself applies in quite general situations, but the error estimates are restricted to the cases of periodic coefficients.

11h30 – 12h15
Robert Eymard, **Cindy Guichard**, Raphaèle Herbin and Roland Masson
Vertex centered finite volume scheme for compositional multiphase flows in porous media

This talk deals with the discretization on general 3D meshes of multiphase compositional Darcy flows in porous media. The model accounts for the coupling of the mass balance of each component with the pore volume conservation and the thermodynamical equilibrium, and dynamically manages phase appearance and disappearance. The spatial discretization of the multiphase compositional Darcy flows is based on the Vertex Approximate Gradient scheme (VAG). It leads to an unconditionally coercive scheme for arbitrary meshes and permeability tensors. The stencil of this vertex-centered scheme typically comprises 27 points on topologically Cartesian meshes, and the number of unknowns on tetrahedral meshes is considerably reduced, compared with usual cell-centered approaches. The efficiency of this approach is exhibited on several examples, including the near well injection of miscible CO_2 in a saline aquifer taking into account the vaporization of H_2O in the gas phase as well as the precipitation of salt.

12h15 – 12h45 Discussion

12h45 – 14h30 Lunch

14h30 – 15h15
Robert Eymard, Cindy Guichard and **Roland Masson**
Finite volume discretization of two-phase porous media flows with discontinuous capillary pressures on general meshes

This talk deals with the extension of the vertex approximate gradient scheme (VAG) to take into account discontinuous capillary pressures for two-phase Darcy flows. This is an important issue to modelize capillary barriers arising for example in CO2 storage or basin modelling. Our method is based on the recent work of Clement Cances, Konstantin Brenner and Danielle Hilhorst that introduces a convergent cell centered two-point flux discretization of such problem on admissible meshes. The proposed extension of the VAG scheme is based on a specific choice of the primary unknowns at the vertices that allows for discontinuous saturations at the interface between different rock types and respect the continuity of the phase pressures. Numerical tests exhibit the efficiency of this approach on general meshes.

15h15 – 16h00
Clément Cancès, **Konstantin Brenner** and Danielle Hilhorst
A convergent finite volume scheme for two-phase flows in porous media with discontinuous capillary pressure field

We propose a finite volume scheme for an incompressible and immiscible two-phase (e.g. oil and water) flow in heterogeneous porous media. More precisely we assume that the flow takes place in the domain made of two different rocks, so that the capillary pressure is discontinuous across the interface between the rocks. As a consequence the solution itself is discontinuous across the interface with nonlinear transmission conditions. We discretize the problem by means of a numerical scheme which reduces to a standard finite volume scheme in each sub-domain and prove the convergence of a sequence of approximate solutions towards a weak solution of the continuous problem. The numerical tests show that the scheme can reproduce the oil-trapping phenomenon.

16h00 – 17h00 Concluding discussion, planning of future actions, etc.

17h00 End of the meeting