C2S@Exa
Computer and Computational Sciences at Exascale
An Inria Project Lab dedicated to high performance computing driven by applications

Thematic pole 2 - Numerical schemes for PDEs

Stéphane Lanteri
Inria Sophia Antipolis - Méditerranée, France

Midterm evaluation meeting
Inria Place d'Italie - January 13, 2015

Problem characteristics and algorithmic challenges

- Future exascale supercomputers will enable the numerical treatment of more challenging problems involving, on one hand, discretized models with higher spatial resolution and, on the other hand, more complex physical models possibly encompassing multiple space and time scales
- However, the simulation of such complex physical phenomena will require very accurate and efficient numerical schemes, that will ideally be able to automatically switch between arbitrary high order accuracy in regions where the problem solution is smooth, and low order accuracy combined to local adaptivity of the discretization mesh in less regular regions
- From the algorithmic point of view, the implementation of the proposed numerical schemes will have to be optimized for maximizing both the single core computational performance and the scalability in view of exploiting massive parallelism

Use cases PDE models and schemes

- Nuclear waste management (ANDRA)
  TRACES simulation software

  Darcy's equation (hydraulic flow conditions)
  Advection-diffusion-reaction (radionuclide transport)

  Spatial discretization of 3d problems relies on a discontinuous finite element method for the advection part and mixed hybrid finite element method for the other parts, formulated on conforming, structured or unstructured grid with hexahedral elements

  Time integration relies on explicit or implicit schems and in the latter case, a Newton method is used for the linearization of the impicit systems leading to the solution of sparse linear systems of equations

Use cases PDE models and schemes

- Nuclear energy production (CEA-IRFM)
  GYSELA simulation software

  Global non-linear electrostatic code which solves the gyrokinetic equations (Vlasov) in a five dimensional phase space with a semi-Lagrangian method

  Structured mesh grid which is kept fixed in time in the phase space (Eulerian method)

  The Vlasov equation is integrated along the trajectories (Lagrangian method) using the invariance of the distribution function along the characteristics

  The Vlasov gyrokinetic equation is coupled to a Poisson-like equation which gives the 3D electrostatic field generated by the distribution function of the particles

Objectives and activities in C2S@Exa

Use cases PDE models and schemes

- Nuclear energy production (CEA-IRFM)
  JOREK simulation software

  JOREK is used for magneto-hydrodynamic (MHD) modeling of plasma dynamic in Tokomak geometries

  JOREK is dedicated to Edge Localized Modes (ELMs) and disruptions

  Numerical approximation is based on finite elements where 3D basic functions are tensor products of poloidal (2d) by totoidal (1d)

  Basically, JOREK uses curved bicubic isoparametric elements in 2d and a spectral decomposition (sine, cosine) into the toroidal axis

  Time integration makes use of an implicit schemes leading to sparse linear systems of equations with a special block structure

Objectives and activities in C2S@Exa

Research directions

1. Improvement of numerical schemes in TRACES, GYSELA and JOREK
   - Accuracy, robustness, and computational efficiency
   - TRACES software

     PhD thesis of Nabil Birgle in the Pomdapi project-team, in collaboration with ANDRA

     Robust and efficient finite element solver for subsurface flows
2. Study of innovative schemes
   - Mainly (so far) for hyperbolic systems of PDEs
   - High order accuracy in space and time
   - Adpated to heterogeneous parallel systems
   - Focus on Discontinuous Galerkin schemes
   - Scalable Vlasov-Maxwell solver for plasma dynamics

     Tonus project-team in relation with CEA-IRFM
   - Scalable Debye-Maxwell solver for bioelectromagnetics

     Nachos project-team in the context of the DEEP-ER FP7 project

Participants

- Nachos project-team
  Tristan Cabel (fixed-term engineer, Simon ADT, Inria/C2S@Exa)
  Alexandra Christophe (postdoc, Inria/C2S@Exa)
  Rachid El Khaoulani (fixed-term engineer, ANDRA grant)
  Stéphane Lanteri (researcher)
  Raphaël Léger (fixed-term engineer, DEEP-ER FP7 project)
  Ludovic Moya (fixed-term engineer, TECSER project)
- Pomdapi project-team
  Nabil Birgle (PhD, Inria/C2S@Exa)
  Jérôme Jaffré (researcher)
  Michel Kern (researcher)
- Sage project-team
  Edouard Canot (researcher)
  Jocelyne Erhel (researcher)
- Tonus project-team
  Philippe Helluy (professor, University of Strasbourg)
  Sever Hirstoaga (researcher)
  Thomas Strub (PhD, CIFRE/AxesSim)

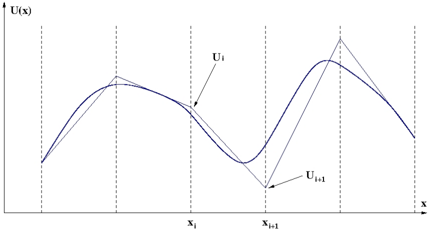

Continuous P1 interpolation

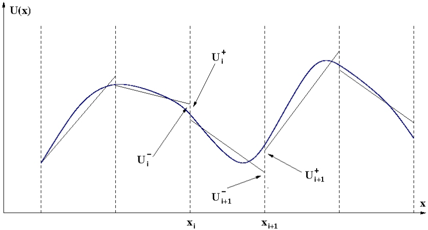

Discontinuous P1 interpolation

- Naturally adapted to heterogeneous media and discontinuous solutions
- Can easily deal with unstructured, possibly non-conforming meshes (h-adaptivity)
- High order with compact stencils and non-conforming approximations (p-adaptivity)
- Usually rely on polynomial interpolation but can also accomodate alternative functions (e.g plane waves)
- Yield block diagonal mass matrices when coupled to explicit time integration schemes
- Amenable to efficient parallelization
- But leads to larger problems compared to continuous finite element methods

Why using DG for time-domain electromagnetics ?

- Heterogeneity is ideally treated at the element level
  - Discontinuities occur at material (i.e element) interfaces
  - Mesh generation process is simplified
- Wavelength varies with ϵ and μ
  - For a given mesh density, approximation order can be adapted at the element level in order to fit to the local wavelength
- Yield block diagonal mass matrices when coupled to explicit time stepping schemes

Application context: computational electromagnetics DEEP-ER FP7 project

- Development of a flexible and efficient finite element simulation tool adapted to hybrid MIMD-SIMD parallel systems for the study of 3D ElectroMagnetic (EM) wave propagation problems in complex domains and heterogeneous media
- Application to the numerical modeling of human exposure to EM fields
- Numerical ingredients
  - Unstructured meshes (tetrahedra in 3D)
  - Discontinuous Galerkin Time-Domain method with polynomial interpolation (DGTD-$\mathbb{P}_p$ method)
- Hybrid MIMD/SIMD parallelization strategy
  - Coarse grain : mesh partitioning+MPI
  - Fine grain : loop parallelization with OpenMP

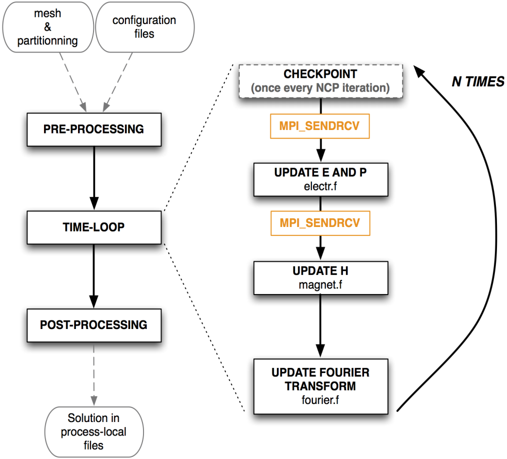

Structure of the DGTD code for solving Debye Maxwell system

Application context: computational electromagnetics
SC'10 - Hybrid CPU-GPU computing

- HPC resource made available by GENCI (allocation 2010-t2010065004)
- Hybrid CPU-GPU Bull cluster of the CCRT
- 1068 Intel CPU nodes with two quad-core Intel Xeon X5570 Nehalem processors operating at 2.93 GHz each
- 48 Teslas S1070 GPU systems with four GT200 GPUs and two PCI Express-2 buses each
- Network is a non-blocking, symmetric, full duplex Voltaire InfiniBand double data rate organized as a fat tree
- Hybrid MPI-CUDA programming model
- Simulations are performed in single precision arithmetic

Application context: computational electromagnetics
SC'10 - Hybrid CPU-GPU computing

- Mesh: # elements = 5,536,852
- Total # dof is 132,884,448 (DGTD-$\mathbb{P}_1$ method) and 332,211,120 (DGTD-$\mathbb{P}_2$ method)
- Time on 64 CPU cores for the DGTD-$\mathbb{P}_1$ method: 7 h 10 mn

# GPU	DGTD-$\mathbb{P}_1$			DGTD-$\mathbb{P}_2$		
	Time	GFlops	Speedup	Time	GFlops	Speedup
64	12 mn	2762	-	59 mn	4525	-
128	7 mn	4643	1.7	30 mn	8865	1.95

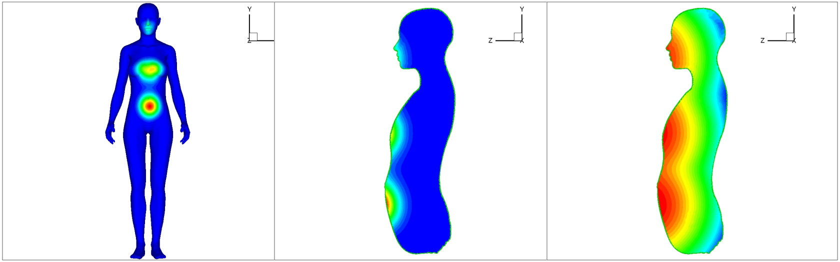

Application context: computational electromagnetics DEEP-ER FP7 project

- Hybrid MPI-OpenMP programming model
- Parallel I/O SIONlib library
- OmpSs runtime

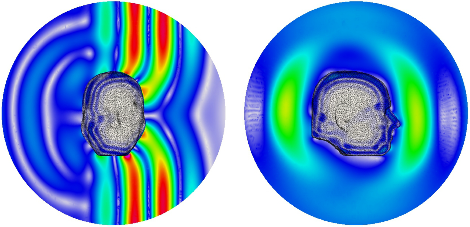

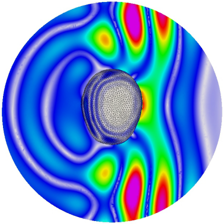

Application context: computational electromagnetics DEEP-ER FP7 project

In the DEEP project, an innovative architecture for heterogeneous HPC systems has been developed based on the combination of a standard HPC Cluster and a tightly connected HPC Booster built of manycore processors

DEEP-ER will use second-generation Intel Xeon Phi manycore CPUs that boot without the help of an attached Intel Xeon processor for the Booster part

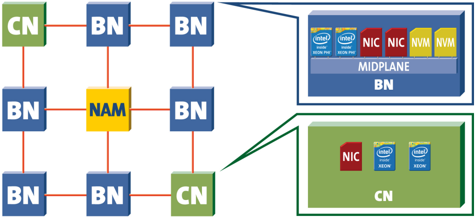

Discontinuous Galerkin (DG) methods

DGTD for Debye-Maxwell system

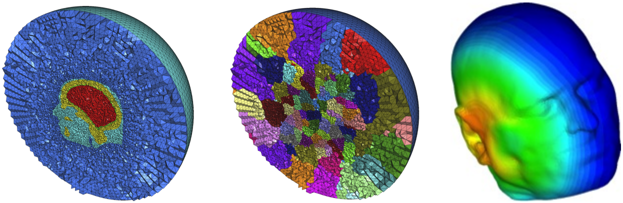

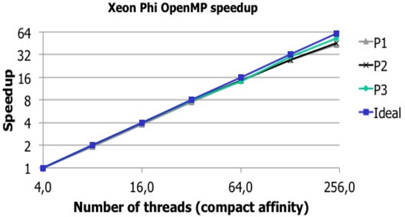

Simulation of the interaction of an electromagnetic wave with biological tissues for numerical dosimetry studies using a high order DG-based Debye-Maxwell solver (top figures). Parallel speedup of the OpenMP parallelization on a KNC (Knights Corner) Xeon Phi chip (strong scaling for different interpolation orders in the DGTD method, top figure).

Application context: computational plasma dynamics

- CLAC (Conservation Laws Approximation on manyCores) is a C++ library developed by Inria Nancy - Grand-Est (Tonus project-team) and AxesSim (SME in Strasbourg)
- Discontinus Galerkin methods for systems of conservations laws
- Abstract physical model: the user provides the numerical flux and the source term (Maxwell, Vlasov-Maxwell, MHD, Euler, Navier-Stokes, etc.)
- Actually used for electromagnetics/plasms dynamics simulations
- Hybrid MPI-OpenCL programming model
- Subdomain decomposition: each domain is associated to a MPI process, and each MPI process is associated to an OpenCL device (CPU or GPU)
- Zone decomposition: each subdomain is split into volume zones and interface zones
- A zone possess identical elements (same order, same geometry, same physical model)
- A computation kernel is compiled for each zone (for avoiding branch tests)
- Fine grain parallelization based on a task dependency graph

Task-based programming model

- Well adapted to (massive) fine grain parallelism
- A set of inter-dependent tasks can be modelled as a Directed Acyclic Graph (DAG)
- The DAG is traversed breadth-first by a task scheduler assigning the tasks to the cores
- Software infrastructure for such task-based parallelism: QUARK, Cilk and StarPU
- Programming models: OpenCL, OpenM 3.0 and Intel's TBB

CLAC task graph
By courtesy of P. Helluy, Tonus project-team

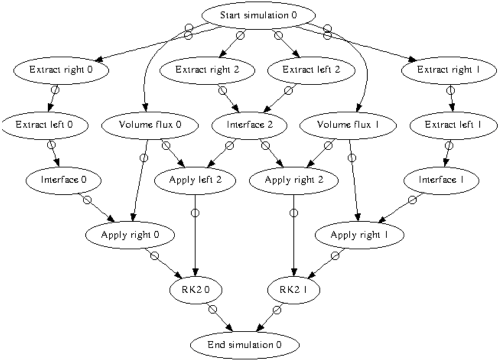

Activities under this thematic pole follow two paths:

- Specific activities for improving the *performances* of existing numerical schemes in the TRACES, GYSELA and JOREK simulation software
  Pomdapi, Sage project-teams, and to some extent, Tonus project-team

- Activities for developing *high performance* numerical schemes and *scalable solvers* incorporated within in-house simulation software (so far)
  Physical settings: electromagnetic waves (and potentially other wave models) and material physics
  Need for appropriate application contexts (use cases)
  Hiepacs, Nachos and Tonus project-teams