Development of a high performance manycore-enabled finite element type solver for the numerical modeling of nanoscale light/matter interaction

Master internship project

Host institutions and supervisors
Stéphane Lanteri, stephane.lanteri@inria.fr
Inria Sophia Antipolis – Méditerranée, Nachos project-team
Jean-François Méhaut, jean-francois.mehaut@imag.fr
Inria Grenoble – Rhône Alpes, Corse project-team

1. Scientific and technological context

Nanostructuring of materials has opened up a number of new possibilities for manipulating and enhancing light-matter interactions, thereby improving fundamental device properties. Low-dimensional semiconductors, like quantum dots, enable one to catch the electrons and control the electronic properties of a material, while photonic crystal structures allow synthesizing the electromagnetic properties. These technologies may, e.g., be employed to make smaller and better lasers, sources that generate only one photon at a time for applications in quantum information technology, or miniature sensors with high sensitivity. The incorporation of metallic structures into the medium adds further possibilities for manipulating the propagation of electromagnetic waves. In particular, this allows subwavelength localization of the electromagnetic field and, by subwavelength structuring of the material, novel effects like negative refraction, e.g. enabling super lenses, may be realized. Nanophotonics is the field of science and technology aimed at establishing and using the peculiar properties of light and light-matter interaction in various nanostructures. Nanophotonics includes all the phenomena that are used in optical sciences for the development of optical devices. Because of its numerous scientific and technological applications (e.g. in relation to telecommunication, energy production and biomedicine), nanophotonics represents an active field of research increasingly relying on numerical modeling beside experimental studies. Computational nanophotonics is the discipline aiming at studying nanoscale light/matter interaction using computer simulations.

2. Mathematical and numerical modeling

Numerical modeling of light/matter interaction with metallic nanostructures requires to solve the system of time-domain Maxwell equations coupled with an appropriate material law for metals at optical frequencies. During the last twenty years, numerical methods formulated on unstructured meshes have drawn a lot of attention in computational electromagnetics with the aim of dealing with irregularly shaped structures and heterogeneous media. In particular, the discontinuous Galerkin time-domain (DGTD) method has progressively emerged as a viable alternative to well established finite-difference time-domain (FDTD) and finite-element time-domain (FETD) methods for the numerical simulation of electromagnetic wave propagation problems in the time-domain. In this work, we will consider such a DGTD method that has been recently designed at Inria in the Nachos project-team for the simulation of nanoscale light/matter interaction problems [1].

The DGTD method can be considered as a finite element method where the continuity constraint at an element interface is released. While it keeps almost all the advantages of the finite element method (large spectrum of applications, complex geometries, etc.), the DGTD method has other nice properties, which explain the renewed interest it gains in various domains in scientific computing:

- ✓ It is naturally adapted to a high order approximation of the unknown field. Moreover, one may increase the degree of the approximation in the whole mesh as easily as for spectral methods but, with a DGTD method, this can also be done locally i.e. at the mesh cell level. In most cases, the approximation relies on a polynomial interpolation method but the method also offers the flexibility of applying local approximation strategies that best fit to the intrinsic features of the modeled physical phenomena.
- ✓ When the discretization in space is coupled to an explicit time integration method, the DG method leads to a block diagonal mass matrix independently of the form of the local approximation (e.g the type of polynomial interpolation). This is a striking difference with classical, continuous FETD formulations. Moreover, the mass matrix is diagonal if an orthogonal basis is chosen.
- ✓ It easily handles complex meshes. The grid may be a classical conforming finite element mesh, a non-conforming one or even a hybrid mesh made of various elements (tetrahedra, prisms, hexahedra, etc.). The DGTD method has been proven to work well with highly locally refined meshes. This property makes the DGTD method more suitable to the design of a *hp*-adaptive solution strategy (i.e. where the characteristic mesh size *h* and the interpolation degree *p* changes locally wherever it is needed).
- ✓ It is flexible with regards to the choice of the time stepping scheme. One may combine the discontinuous Galerkin spatial discretization with any global or local explicit time integration scheme, or even implicit, provided the resulting scheme is stable.
- ✓ It is naturally adapted to parallel computing. As long as an explicit time integration scheme is used, the DGTD method is easily parallelized. Moreover, the compact nature of method is in favor of high computation to communication ratio especially when the interpolation order is increased.

As in a classical finite element framework, a discontinuous Galerkin formulation relies on a weak form of the continuous problem at hand. However, due to the discontinuity of the global approximation, this variational formulation has to be defined at the element level. Then, a degree of freedom in the design of a discontinuous Galerkin scheme stems from the approximation of the boundary integral term resulting from the application of an integration by parts to the element-wise variational form. In the spirit of finite volume methods, the approximation of this boundary integral term calls for a numerical flux function, which can be based on either a centered scheme or an upwind scheme, or a blend of these two schemes.

Because of its inherently local nature, the DGTD method is ideally suited to hybrid MIMD/SIMD parallelism. A coarse grain MIMD parallelization is obtained by partitioning the computational mesh and adopting a message-passing programming model implemented with the MPI standard. Achieving an efficient fine grain SIMD parallelization is more tricky but can lead to high floating point rates when systems

based of GPU cards are considered [2].

3. Optimization of the auto-tuning of computing kernels

The Corse project-team at Inria Grenoble – Rhône Alpes develops the BOAST environment that can be used to develop portable compute intensive kernels that are optimized for a specific architecture such as manycore CPUs or GPUs. BOAST [3] provides scientific application developers with a framework to develop and test application-computing kernels. The developer starts from an application kernel (either designed or implemented) and write it in a dedicated language. This language provides enough flexibility for the kernel to be meta-programmed with several orthogonal optimizations. Then the developer selects a kernel configuration and a target language. Those parameters define the output source code that will be generated by BOAST. The resulting code source is then built according to the user specified compiler and options. If input data are available then the kernel can be benchmarked and tested for non-regression. Based on the results, other optimizations can be selected or new optimizations can be added to the BOAST sources. The process can be repeated until a good candidate is found on the target platform. The resulting kernel is then added to the application. Several steps in the workflow can be automated: optimization and compiler flags exploration, non-regression testing, benchmarking, etc. This automation can be scripted by the user or by interfacing other dedicated tools.

4. Objectives of the study

The coarse grain MIMD parallelization of the DGDT nanophotonics solver developed in [1] is well mastered. Beside, its fine grain SIMD parallelization has note been studied in details. In this context, the main objective of this study is to specify and implement a strategy for achieving a portable, scalable and computationally efficient SIMD parallelization of this solver. The main hardware target is a GPU accelerator although a manycore system based on the Intel MIC architecture will also be of interest to this study. The internship project will consist of the following steps:

1. Familiarization with the DGTD nanophotonics solver [1], identification of the main compute intensive kernels that will be considered for the SIMD parallelization. Bibliography study for a state-of-the-art in SIMD (in particular GPU) parallelization of DGTD solvers for electromagnetics and nanophotonics (reference [2] is only one example of existing works) – 1 month;
2. Specification of SIMD parallelization strategies of the identified compute intensive kernels targeting a GPU architecture and the Intel MIC architecture – 1 month;
3. Familiarization with the BOAST environment and its DSL. Development of the DSL modules corresponding to the SIMD parallelization strategies of the identified compute intensive kernels and preliminary validations on a model test problem – 2 months;
4. Optimization of DSL modules, final validations and detailed performance evaluation on a realistic nanophotonics problem – 2 months.

The internship will take place both in Grenoble and Sophia Antipolis with a time schedule to be determined. The duration of the internship will be 6 months and the monthly net salary will be around 1200 €.

References

[1] J. Viquerat, *Simulation of electromagnetic waves propagation in nano-optics with a high-order discontinuous Galerkin time-domain method*, PhD thesis, University of Nice-Sophia Antipolis, December 2015. https://tel.archives-ouvertes.fr/tel-01272010v1

[2] R. Diehl, *Analysis of metallic nanostructures by a discontinuous Galerkin time-domain Maxwell solver on graphics processing units*, PhD thesis, KIT, December 2012. http://d-nb.info/1029764786/34

[3] J. Cronsioe, B. Videau and V. Marangozova-Martin, *BOAST: Bringing optimization through automatic source-to-source transformations*, In Embedded Multicore SoCs (MCSoC), 2013 IEEE 7th International Symposium on, pages 129–134. IEEE, 2013. http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6657917