

Efficient Parallel Computing for Laser-Gas Quantum Interaction and Propagation

E. Lorin

University of Ontario Institute of Technology
Faculty of Science
2000, Simcoe North Street, L1H 7K4, Oshawa, Canada
Emmanuel.Lorin@uoit.ca

A. D. Bandrauk

Laboratoire de chimie théorique Faculté des Sciences
Université de Sherbrooke, Qué, J1K 2R1, Canada
Andre.Bandrauk@USherbrooke.ca

Abstract

We present in this short paper some domain decomposition and high performance numerical techniques for simulating laser-gas interaction and propagation using a multi-scale Maxwell-Schrödinger model. Ionization and high order harmonic generation are in particular taken into account in this micro-macro model, leading to additional simulation difficulties. We propose some benchmarks to validate the presented techniques.

1. Introduction

Simulation of intense and ultrashort laser-matter interaction necessitates a very precise micro-macro modeling. In [4], [5] we have introduced a multi-scale model coupling the macroscopic Maxwell equations modeling the propagation of an electro-magnetic field in a gas modeled by many Time Dependent Schrödinger Equations (TDSE's). The complexity of this model especially in 3-D, requires the use of efficient techniques. In this paper, we present some numerical and parallel methods for simulating efficiently ultrashort (less than 10^{-14} second), intense (more than 10^{13} W · cm⁻²) and high frequency (less than 800 nm) laser pulses propagating in dense (more than 10^{17} mol · cm⁻³) gaseous media. At this point, numerical schemes used are still relatively standard and consists of the use of a Yee-like scheme for the Maxwell equations and Crank-Nicolson scheme for the TDSE's and multi-grid techniques. As discussed later in the paper the TDSE approximation is the most costly part of the simulation due the fact that an accurate description of

the gas requires the computation of hundreds or thousands TDSE's.

In Sections 2, 3 we present shortly the Maxwell-Schrödinger model and its numerical approximation. We will consider the case of non Born-Oppenheimer approximations (moving nuclei) even if most of the computations are done under the Born-Oppenheimer approximation allowing to reduce the complexity of the problem. We focus in particular on the boundary conditions for laser-molecule TDSE's (still in Section 3) allowing a crucial reduction of the algorithmic complexity of the numerical scheme approximating these equations. Then a domain decomposition approach is presented and its natural parallelization in Section 4. A numerical experiment is then proposed to validate the chosen method.

2 Physical problem and modeling

The model we consider is a coupling between the microscopic laser-molecule Schrödinger equations and macroscopic Maxwell's equations. It allows to take into account high order harmonic generation, ionization and is then more precise than classical nonlinear Schrödinger's equations. Some informations related to the model may be found in [5] and some applications in [6] or [7]. As it is not the purpose of this paper we simply present the system of equations (in atomic unit). First, we introduce some important notations. For the Maxwell equations, we will denote by $\Omega \subset \mathbb{R}^3$ the spatial domain with a regular boundary denoted by Γ and by $\mathbf{r} = (x, y, z)$ the space variable in Ω . At the molecule scale we will denote by $(\mathbf{r}' = (x', y', z'), R') \in \mathbb{R}^3 \times \mathbb{R}_+^*$ the space variable (for electrons and ions). The molecular

density is supposed to be constant in time and is given by a function n . The equations we consider are the following:

$$\left\{ \begin{array}{l} \partial_t \mathbf{B}(\mathbf{r}, t) = -c \nabla \times \mathbf{E}(\mathbf{r}, t), \\ \partial_t \mathbf{E}(\mathbf{r}, t) = c \nabla \times \mathbf{B}(\mathbf{r}, t) - 4\pi \partial_t \mathbf{P}(\mathbf{r}, t), \\ \nabla \cdot \mathbf{B}(\mathbf{r}, t) = 0, \\ \nabla \cdot (\mathbf{E}(\mathbf{r}, t) + 4\pi \mathbf{P}(\mathbf{r}, t)) = 0, \\ \mathbf{P}(\mathbf{r}, t) = n(\mathbf{r}) \sum_{i=1}^{\ell} \mathbf{P}_i(\mathbf{r}, t) = n(\mathbf{r}) \sum_{i=1}^{\ell} \chi_{\Omega_i}(\mathbf{r}) \\ \int_{\mathbb{R}^4} \psi_i(R', \mathbf{r}', t) \mathbf{r}' \psi_i^*(R', \mathbf{r}', t) d\mathbf{r}' dR', \\ i \partial_t \psi_i(R', \mathbf{r}', t) = -\frac{\Delta_{\mathbf{r}'}}{2} \psi_i(R', \mathbf{r}', t) \\ -\frac{\Delta_{R'}}{m_p} \psi_i(R', \mathbf{r}', t) \\ + \left(V_i(R') + V_c(R', \mathbf{r}') \right. \\ \left. + \mathbf{r}' \cdot \mathbf{E}_{r_i}(t) \right) \psi_i(R', \mathbf{r}', t). \end{array} \right.$$

In the previous equation V_c denotes the Coulomb potential and V_i the nuclei potential. We impose Dirichlet conditions on Γ :

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{B}(\mathbf{r}, t) = \mathbf{0}, \quad \forall t \geq 0, \quad \forall r \in \Gamma.$$

Ω_i denotes the domain associated to ψ_i wavefunction of the i^{th} TDSE, and \mathbf{P}_i the polarization associated to this domain (see fig. 1). Functions χ_{Ω_i} are defined by $\chi \otimes \mathbf{1}_{\Omega_i}$ where $\chi \in \mathcal{C}_0^\infty$ is a Plateau function and $\mathbf{1}_{\Omega_i}$ the characteristic function of Ω_i . Naturally we have $\cup_{i=1}^{\ell} \Omega_i = \Omega$ and we denote $\bar{\psi} = (\psi_1, \dots, \psi_\ell)^T$. Results of existence and uniqueness of weak solutions for this system will be soon presented in [3].

3 Numerical aspects

3.1 Numerical methods

A modified Yee scheme is used for solving the Maxwell equations. This is an order 2 corrector-predictor finite difference scheme [12] where the electric and magnetic fields are computed on two dual temporal and spatial staggered grids. It is stable under a CFL conditions relating time and space steps. As it is an explicit scheme its numerical complexity is relatively small.

A Crank-Nicolson scheme with three-point stencil is used to approximate the TDSE's. At least on an infinite domain this is an unconditionally stable and order 2 scheme in space and time. Physical and numerical considerations lead to choose properly the space and time steps for the Maxwell and TDSE solvers.

The coupling between these two schemes is an order 2 splitting that allows to conserve a global order 2 in space and time as follows:

$$\mathbf{P}^n(\mathbf{r}') \xrightarrow{\text{ME}} \mathbf{E}^{n+1}(\mathbf{r}') \xrightarrow{\text{TDSE}} \psi_{\mathbf{r}'}^{n+1}(\mathbf{r}) \longrightarrow \mathbf{P}^{n+1}(\mathbf{r}')$$

Classical linear algebra methods are used to solve sparse linear systems GMRES (see [10]) and eigenvalue problems (Arnoldi using the Library PRIMME [9, 11]). Compressed Row or Column Storage are naturally used.

The principle of the scheme is summarized on fig. 1 It consists physically and then numerically of subdividing

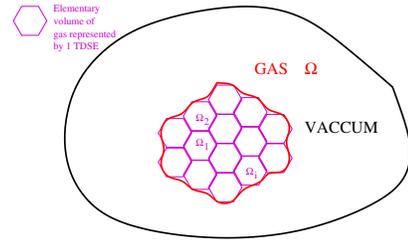


Figure 1. Physical modeling

in small volumes of gas $(\Omega_i)_{i=1, \dots, \ell}$ the sample we study $\Omega = \cup_{i=1}^{\ell} \Omega_i$. Each small volume of gas is represented by one single TDSE. From a computing point of view this subdivision can be reproduced easily. At every time step, we solve independently one TDSE per subdomain of gas Ω_i . We then deduce polarization (supposed to be constant inside each subdomain Ω_i). Once this is done the updated polarization can be used by the Maxwell equation solver.

3.2 Algorithmic complexity

The computation of each TDSE requires $O(N^{3/2})$ operations every time step, where N is the size of the matrix then also the number of degrees of freedom in one TDSE computational domain. This is due to the use of a preconditioned linear solver for sparse linear systems (GMRES). Denoting ℓ the number of TDSE's that are solved in the system, every time step $O(\ell N^{3/2})$ operations are then necessary. Polarization computation requires $O(N)$ operations with a very small prefactor and is then negligible compared to the TDSE computations. At the same time, as the Yee scheme is an explicit scheme, the computation of Maxwell's equations requires $O(M)$ operations where M is the number of degrees of freedom in the Maxwell grid. In practice there is no relation between M and N and usually $N^{3/2} \ell$ is much larger than M . For large gaseous media (M large), ℓ has to be chosen large in order to describe precisely the laser-gas interactions. The larger ℓ is chosen the more precise the gas description will be.

3.3 Artificial Boundary Conditions

In order to maintain a good accuracy of the numerical schemes and of the physical model the most natural way to reduce the algorithmic complexity is to reduce as much as possible the TDSE computational domain. To this end, we introduce transparent and artificial boundary conditions. In theory N has to be very large as the TDSE computational domain is supposed to represent \mathbb{R}^3 ($\mathbb{R}^3 \times \mathbb{R}_+$ beyond Born-Oppenheimer approximation). However at infinity the wavefunction modulus decreases to zero, so that N is finite (but very large for large laser pulse intensities). We shortly recall the principle of the method. Details can be found in particular in [8]. The very general idea consists of the following. Denoting by $S(x, t, D)$ the Schrödinger operator and d the physical dimension, we consider

$$\begin{cases} S(x, t, D)u = 0, & (x, t) \in \mathbb{R}^d \times \mathbb{R}_+, \\ u(x, 0) = u_0(x), & x \in \mathbb{R}^d. \end{cases}$$

The idea of transparent (or artificial for approximate ones) boundary conditions consists of finding κ (as small as possible) subset of \mathbb{R}^d containing the support of u_0 and a pseudo-differential operator $\mathcal{B}(x, t, D)$, such that the solution of

$$\begin{cases} S(x, t, D)v = 0, & (x, t) \in \kappa \times \mathbb{R}_+, \\ \mathcal{B}(x, t, D)v = 0, & (x, t) \in \partial\kappa \times \mathbb{R}_+, \\ v(x, 0) = u_0(x), & x \in \mathbb{R}^d \end{cases}$$

verifies for all time $t \in \mathbb{R}_+$, $u|_{\kappa}(\cdot, t) = v(\cdot, t)$.

In practice κ represents one TDSE computational domain. An adapted choice of \mathcal{B} allows a drastic reduction of the size of the computational domain κ , that is of the number of degrees of freedom N and this, for all the TDSE's used in the gas modeling. Taking Dirichlet or Neumann boundary conditions lead as is well-known to spurious reflecting waves inside the computational domain. Although Dirichlet-Neumann boundary conditions seem relatively adapted we have chosen Volkov-like boundary conditions. Based on the solution of laser-molecule TDSE without electronic potential and coupled *via* splitting with the Coulomb potential using Filon's integration scheme for highly oscillating functions [1], [2]. This approach gives very promising results and are presented in [8].

4 Parallelism

The physical configuration and model complexity leads to consider efficient parallel techniques.

4.1 Parallel Approach

The parallel approach we use can be summarized as follows. To simplify the presentation we suppose that

there is one processor (2 in practice on mammoth, <http://ccs.usherbrooke.ca>) per memory node and we work with N_p processors. From time t^n to t^{n+1} the process is the following.

1. From time t^n to $t^{n+1/2}$ we solve the ℓ TDSE's on m processors with then ℓ/m TDSE's per processors. Note that each TDSE is solved sequentially. We then deduce the corresponding polarization in each subdomain.
2. Polarization is then sent to the nodes in charge of the Maxwell equation computation (in fact only the nodes related to gas regions of the Maxwell domain, as polarization is zero in vacuum).
3. Maxwell's equations are solved by domain decomposition $D = \cup_i^s D_i$ where the sample of gas $\Omega \subset D$. The complementary $D - \Omega$ is vacuum. From time $t^{n+1/2}$ to t^{n+1} we solve the Maxwell equations on processors p_1, \dots, p_s with 1 subdomain per processor. We then have $N_p = s + m$. The updated electric field is then sent to the nodes in charge of the TDSE computation.
4. Data storage follows the same principle. Maxwell's equation data are stored on the nodes associated to processors p_1, \dots, p_s and TDSE data and the complementary ones.

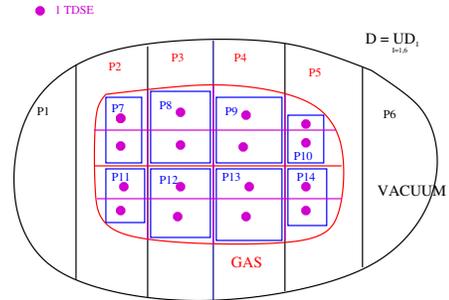


Figure 2. Parallelism principle - 1

As an example on figs. 2, 3, we take $N_p = 14$, $\ell = 16$, $s = 6$ $m = 8$. From time t^n to t^{n+1} the TDSE's ($\ell = 16$) are solved in parallel by processors p_7 to p_{14} ($m = 8$). Then the updated polarization is sent to processors p_1 to p_6 ($s = 6$). Maxwell's equations are solved in parallel by these processors and the updated electric field is sent to processors p_7 to p_{14} . Such a parallelism allows in theory to reduce the algorithmic complexity as follows. Starting sequentially (1 processor) from $O(\ell N^{3/2}) + O(M)$ we can expect at most to obtain using $N_p = p + m$ processors, a final complexity equal to $O(\ell N^{3/2}/m) + O(M/s)$. Communications between processors only involve the

polarization and the electric field (vectors) that require relatively small data storage and then leads in theory to a very good speed-up.

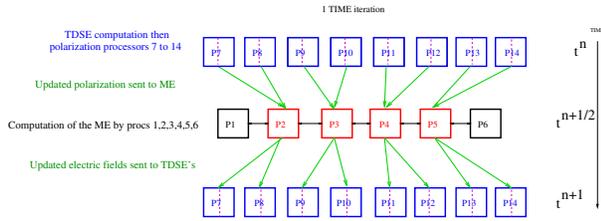


Figure 3. Parallelism principle - 2

Remark. As said above, TDSE computational time is much larger than Maxwell’s equation computation time. Consequently most of the time the s processors solving Maxwell’s equations do not work. This issue is due to the fact that because of the storage of big data (related to Maxwell’s equation discretization) on the nodes associated to these processors, it is not efficient to release (and then reallocate) the memory of these nodes and to use the corresponding processors for the TDSE computation (that also necessitates very big data storage). This corresponds to a lack of efficiency, but limited to the fact that in practice $p \gg s$.

4.2 Results

We now present a benchmark to illustrate how efficient is this parallelization. The numerical Object-Oriented (C++) code is installed on mammouth, super-computer of the RQCHP (<http://ccs.usherbrooke.ca>). Typical simulations require up to 512 Intel Xeon EM64T, 3.6 GHz, (RAM 8 GB) processors during several days. This corresponds to more than 10000 1-D TDSE’s (but some important transversal effects are lost) or several hundreds 3-D TDSE’s coupled with 3-D Maxwell’s equations. The general principle of the presented results is then: on each processor we compute sequentially n TDSE’s. The TDSE’s are coupled at each time step with the Maxwell equations solved by domain decomposition on m processors. In the following experiments we vary the number of TDSE’s and processors for solving these Schrödinger equations. The computations we present here are purely 3-D (TDSE and Maxwell) and in practice we solve $n = 1$ TDSE per processor with 2 processors per memory node. We represent the CPU time for k TDSE’s solved on k processors with ($k = 8, 16, 32, 54, 120$) coupled with the Maxwell equations solved on $m = 5$ processors, following the principle described in figs 2 and 3. An ideal speed-up would correspond to a constant value of the computational time when

increasing the number of TDSE’s and processors.

The TDSE’s are solved on a $40 \times 40 \times 401$ node grids and the Maxwell equations are solved on a grid of $300 \times 300 \times 300$ nodes. MPI is used as a library of message passing. CPU times are given after 50 Maxwell time steps corresponding to 5×50 Schrödinger time steps (each Maxwell time step is subdivided in q Schrödinger time steps with q varies between 5 and 20 depending on the physical data). As expected computational time is almost constant when the number of TDSE’s and processors increases. The reason is as follows. Compared to the Maxwell equation computation, the TDSE computation is much more time consuming. As these TDSE computations are done totally independently at each iteration, and as communications between nodes involves very small data (polarization) a very good speed-up has been reached. On the other hand, if we would consider much larger Maxwell domains but with small TDSE computation domains (for instance for low intense laser pulses but large propagation distances) the efficiency would certainly be reduced. Indeed, the domain decomposition technique for solving the Maxwell equations although relatively efficient, is not as efficient as the TDSE computation parallel technique.

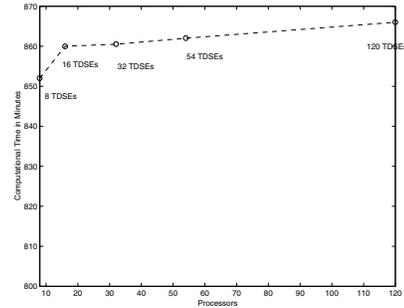


Figure 4. CPU time when solving 3-D Maxwell-Schrödinger’s equations. We solve 1 TDSE per processor (up to 120 TDSE’s and processors) and use 5 processors to solve Maxwell’s equations.

5 Conclusion

We have presented a multi-scale Maxwell-Schrödinger model for intense laser-gas interactions and some efficient numerical and parallel techniques based on domain decompositions, and perfectly adapted to the physical modeling. The numerical code allows a precise description of laser-pulse propagation in dense gaseous media taking into account high order harmonic generation, ionization under or

beyond the Born-Oppenheimer approximation what can not do classical nonlinear Schrödinger or Maxwell models.

References

- [1] A. Iserles. On the numerical quadrature of highly-oscillating integrals. I. Fourier transforms. *IMA J. Numer. Anal.*, 24(3):365–391, 2004.
- [2] A. Iserles. On the numerical quadrature of highly-oscillating integrals. II. Irregular oscillators. *IMA J. Numer. Anal.*, 25(1):25–44, 2005.
- [3] E. Lorin. Global existence and uniqueness of solution for a micro-macro Maxwell-Schrödinger model. In preparation.
- [4] E. Lorin, S. Chelkowski, and A. Bandrauk. A Maxwell-Schrödinger model for non-perturbative laser-molecule interaction and some methods of numerical computation. *Proceeding CRM, American Mathematics Society*, vol 41, 2007.
- [5] E. Lorin, S. Chelkowski, and A. Bandrauk. A numerical Maxwell-Schrödinger model for laser-matter interaction and propagation. *Comput. Phys. Comm.*, 177(12):908–932, 2007.
- [6] E. Lorin, S. Chelkowski, and A. Bandrauk. Propagation effects on attosecond pulse generation. *Proceedings SPIE*, 6733, 2007.
- [7] E. Lorin, S. Chelkowski, and A. Bandrauk. Attosecond pulse generation from aligned molecules - dynamics and propagation in H_2^+ . *New J. Phys.*, 10(025033), 2008.
- [8] E. Lorin, S. Chelkowski, and A. Bandrauk. Mathematical modeling of boundary conditions for laser-molecule time dependent Schrödinger equations and some aspects of their numerical computation - One-dimensional case. *Numer. Methods Partial Differential Equations*, Published online, 2008.
- [9] J. R. McCombs and A. Stathopoulos. Iterative validation of eigensolvers: a scheme for improving the reliability of Hermitian eigenvalue solvers. *SIAM J. Sci. Comput.*, 28(6):2337–2338 (electronic), 2006.
- [10] Y. Saad and W. Schultz. GMRES: a generalized minimal algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Static. Comput.*, 7(3):856–869, 1986.
- [11] A. Stathopoulos. Nearly optimal preconditioned methods for Hermitian eigenproblems under limited memory. I. Seeking one eigenvalue. *SIAM J. Sci. Comput.*, 29(2):481–514 (electronic), 2007.
- [12] K. Yee. Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Transaction on Antennas and Propagation*, AP-16:302–307, 1966.