Numerical investigations on nonlinear nonparaxial beam propagation using graphics processing units

Pedro Chamorro-Posada, Julio Sánchez-Curto and Graham S. McDonald

***Abstract*—We study the performance of a nonparaxial beam propagation method accelerated using massively parallel computation in graphic processing units. The implementation is tested in two different NVIDIA hardware architectures, Tesla and Fermi, and the results are compared with a CPU-based parallel implementation using Open MPI.**

***Index Terms*—Nonparaxial nonlinear beams, Helmholtz solitons, massivelyl parallel processing.**

I. Introduction

EFFICIENT and accurate nonparaxial numerical solutions of the propagation of optical beams in nonlinear media [1] have served as the cornerstone of extensive analytical and computational research on Helmholtz solitons [2-9]. This numerical scheme has also been assessed against the full solution of Maxwell equations in the time domain using the Transmission Line Matrix (TLM) method [10]. Both the time-domain [10] and complex envelope numerical methods [11] have been parallelized and tested in various types of architectures.

Here, we present the results of numerical investigations on nonparaxial beam propagation in nonlinear optical media using graphics processing units (GPUs). The numerical method described in [1] makes extensive usage of the FFT algorithm and is particularly well suited for its implementation in GPUs. We make a performance evaluation of the GPU implementation in an NVIDIA Tesla C1060 and a (Fermi architecture) NVIDIA GTX 480 GPU. The results are compared with those obtained from an implementation in a shared-memory CPU parallel architecture.

Manuscript received September 30, 2011. This work was supported by the Spanish MICINN under Grant TEC2010-21303-C04-04. We also acknowledge support from NVIDIA.

P. Chamorro-Posada is with the Departamento de Teoría de la Señal y Comunicaciones e Ingeniería Telemática at Universidad de Valladolid, ETSI Telecomunicación, Campus Miguel Delibes, Paseo Belén 15, E-47011 Valladolid, Spain (corresponding autor, phone: +34 983 185545; fax: +34 983 423667; e-mail: pedcha@tel.uva.es).

J. Sánchez-Curto is with the Departamento de Teoría de la Señal y Comunicaciones e Ingeniería Telemática at Universidad de Valladolid, ETSI Telecomunicación, Campus Miguel Delibes, Paseo Belén 15, E-47011 Valladolid, Spain (e-mail: julsan@tel.uva.es).

G.S. McDonald is with the Scholl of Computing, Science and Engineering at the University of Salford, Joule Physics Laboratory, Newton Building, Salford M5 4WT, UK (e-mail: g.s.mcdonald@salford.ac.uk).

II. Nonparaxial nonlinear beam propagation

The propagation of the field envelope $u(\xi,\eta,\zeta)$ of a scalar continuous-wave (CW) optical beam in a nonlinear dielectric medium is accounted for by the equation [1]

$$\kappa\frac{\partial^2 u}{\partial\zeta^2}+i\frac{\partial u}{\partial\zeta}+\frac{1}{2}\nabla^2_{\xi\eta}u+f(|u|;\xi,\eta,\zeta)u=0\,, \qquad (1)$$

where f is a function of the beam intensity and the spatial coordinates that defines the medium nonlinearity and any non-homogeneity that may be present associated either to its linear or nonlinear dielectric properties. The transverse (x,y) and longitudinal (z) coordinates are normalized according to a reference Gausian beam of waist w_0 and diffraction length $L_D=kw_0^2/2$ as

$$\xi=\frac{\sqrt{2}x}{w_0}\,,\ \eta=\frac{\sqrt{2}y}{w_0}\ \text{ and }\ \zeta=\frac{z}{L_D}\,. \qquad (2)$$

$\kappa=1/(kw_0^2)$ is a parameter associated to the degree of nonparaxialiaty of the beam. Equation (1) is fully equivalent to the corresponding nonlinear Helmholtz equation [2]. If the term $\kappa\partial_{\zeta\zeta}u$ is neglected in Eq. (1) one recovers the paraxial nonlinear wave equation. This term can be neglected in the nonlinear evolution equation if three conditions are fulfilled [2]: (i) That the beam is not too narrow (in comparison with the optical wavelength), so the angular spectrum is limited to the paraxial region, (ii) that the beam peak intensity is not too large and (iii) that the beam propagates very close to the optical axis. Under typical nonlinear propagation conditions, (i) and (ii) are intimately linked. So, there are two distinct situations that may lead to a strong nonparxiality in the propagating beam: strong focusing of ultra-intense beams, which result in a large broadening of the angular spectrum, or the propagation of broad beams (in terms of the wavelength) with large angles relative to the optical axis. The validity of Eq. (1) is questionable in the high-intensity nonparaxial scenario, where polarization [1] or other high-order effects are expected to come into play and, possibly, dominate the

nonparaxial correction itself. Nevertheless, the widely accepted conditions that assure the accuracy of the scalar treatment of the optical field are the same for the paraxial equation and Eq. (1) when nonparaxiality has an angular origin.

When restricted to the (1+1)*D* case, Eq. (1) describes the evolution TE-polarized soliton beams under nonparaxial conditions of the angular type, which has also been termed as Helmholtz nonparaxiality. This model equation has been shown to hold exact bright [2] and dark [3] soliton solutions, multi-component [4] and also solitons in non-Kerr nonlinear media [5,6]. Furthermore, even though the evolution equation is non-integrable [7], several analytical properties of its solutions: transformation invariance [2] and conserved quantities [2,5] have been derived. These Helmoltz soliton generalizations of paraxial optical solitons permit to deal with relevant situations with an intrinsic strong angular character, such as the interaction of spatially multiplexed optical beams [7] or the refraction of bright and dark optical solitons at nonlinear interfaces [8,9].

The study of the propagation and stability properties of Helmholtz solitons has relied heavily on intensive numerical simulations [2-9]. In turn, the existence of exact solutions and the knowledge of their analytical properties has played a key role in the development of the required accurate numerical tools [1]. The nonparaxial algorithm employed in the numerical investigations is based on the solution of the difference-differential equation [1]

$$u_{n+1}(\xi,\eta)=\frac{1}{2\kappa+i\Delta\zeta}\left[\left(4\kappa-\Delta\zeta^2\hat{O}-2\Delta\zeta^2 f\left(\left|u_n(\xi,\eta)\right|;\xi,\eta,n\Delta\zeta\right)\right)u_n(\xi,\eta)-\left(2\kappa-i\Delta\zeta\right)u_{n-1}(\xi,\eta)\right], \quad (3)$$

where operator $\hat{O}$ has two components: the transverse diffraction term $\nabla^2_{\xi\eta}$ and a filter that eliminates the evanescent backward propagating components of the beam angular spectrum that would result in numerical instability [1]. Operator $\hat{O}$ is evaluated in the spectral domain.

III. Results and Discussion

There are three basic ingredients in the propagation step described by (3) that must be taken into account: (i) Computation of direct and inverse Fourier transforms using the FFT algorithm, (ii) point-wise matrix multiplications both in the spatial and transformed domains and (iii) complex additions. Even though more elaborate approaches are under progress, here we focus on radix-2 FFTs and well-established state-of-the-art libraries for their parallel computation: the MPI parallel version of the FFTW library [12] (version 3.3) for CPU computations and the NVIDIA CUFFT library [13] for GPU computations.

For the CPU parallel implementation of the whole algorithm (3), we have used C and Open MPI as message passing library. NVIDIA CUDA C extensions have been used for the GPU calculations. In this latter case, all the calculations involved are computed using CUDA kernels and take place within the GPU device.

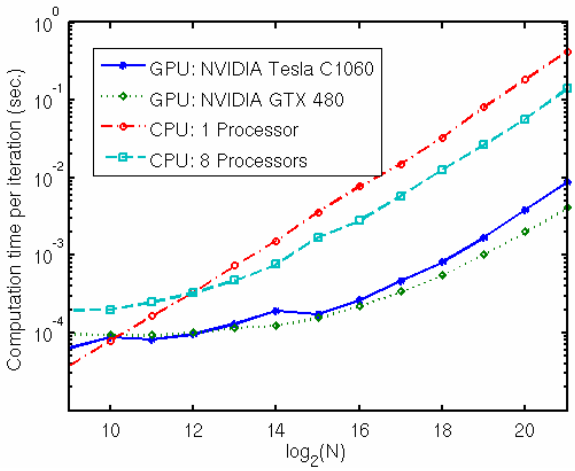

Fig. 1. Logarithmic plot of the floating pointing single precision computation time for one step of the numerical scheme calculated on a run of ten thousand iterations.

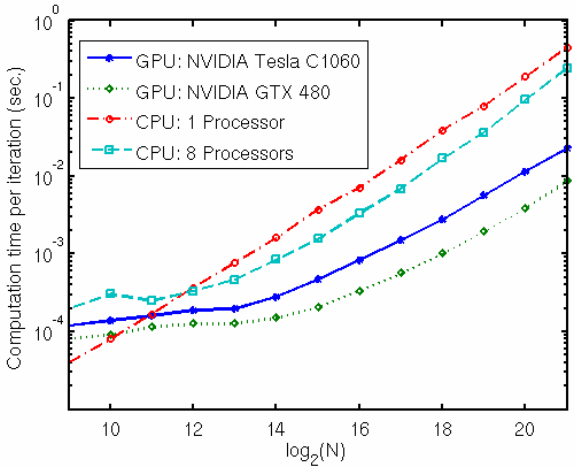

Fig. 2. Logarithmic plot of the floating point double precision computation time for one step of the numerical scheme calculated on a run of ten thousand iterations.

Two different devices have been used for the assessment of the GPU acceleration. A NVIDIA Tesla C1060 (1.3 compute capability) hosted on a Intel i7 at 2.67 GHz with 8 GB of memory and Ubuntu 9.04 (Jaunty) operating system and a Fermi architecture NVIDIA GeForce GTX 480 (2.0 compute capability) hosted on a Intel Core2 Quad CPU Q8400 with 4 GB of memory at 2.66 GHz and Ubuntu 10.10 (Maverick) operating system. The allocation of the GPU resources has been performed using the CUDA occupancy calculator provided by NVIDIA [13]. Both the sequential and parallel CPU calculations are performed on the same i7 CPU where the Tesla C1060 is hosted.

In order to assess the computational efficiency of the nonparaxial propagation scheme in the various architectures, the (real) total execution time for a fixed number of ten thousand steps is calculated. From this, the computation time per iteration is estimated.

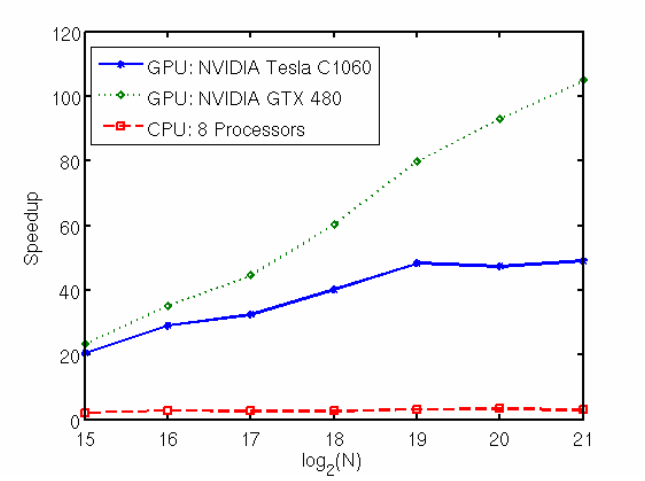

Fig. 3. Measured speedup relative to the one-processor execution for single precision floating point calculations as a function of the number of data points.

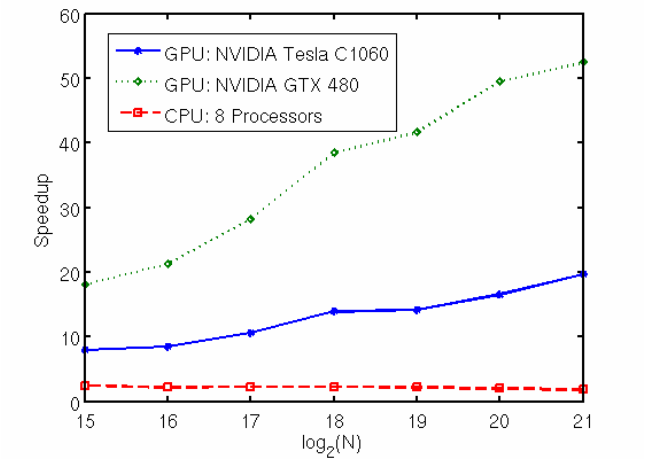

Fig. 4. Measured speedup relative to the one-processor execution for double precision floating point calculation as a function of the number of data points.

Figures 1 and 2 display the total computation time per iteration step for single and double precision, respectively, as a function of the number of data points. Whereas the sequential computational time grows exponentially starting from small numbers of data points, the parallel executions always show a small growth in the computing time for small problem sizes. This effect is due to the initialization time of the GPU devices and MPI. In all cases, the initialization overhead for the MPI is larger than for the GPU devices. For single precision calculations, this overhead is very similar for the Fermi and Tesla architectures, whereas in the case of double precision calculations, the Tesla architecture displays larger values of this parameter. For $N \geq 2^{15}$, the growth of the computation time with the problem size is approximately exponential. In this regime, parallel CPU calculations with eight processors are more efficient that sequential calculations, but the performance is largely improved by using GPU acceleration hardware. Quite remarkably, the improvement increases as the problem size grows.

It is expected that Tesla and Fermi GPUs perform in a similar manner for single precision calculations and about a fourfold improvement in the Fermi hardware for double precision [13]. Nevertheless, as shown in Figure 1, the NVIDIA GTX 480 displays better performance that the Tesla C1060 even for single precision calculations as the problem size grows. The details of the speedup obtained relative to the sequential calculations for $N \geq 2^{15}$ are shown in Figures 3 and 4.

IV. Conclusion

We have carried on a survey on the acceleration of a widely used nonparaxial nonlinear beam propagation algorithm obtained using GPU hardware relative to CPU hardware for sequential and parallel execution. The results have permitted to compare two distinct GPU architectures and have demonstrated an impressive computational performance leap gained from the use of massively parallel NVIDIA GPUs.

References

[1] P. Chamorro-Posada, G.S. McDonald and G.H.C. New, "Nonparaxial beam propagation methods," *Opt. Commun.*, vol. 192, pp. 1-12, 2001.

[2] P. Chamorro-Posada, G.S. McDonald and G.H.C. New, "Nonparaxial solitons," *J. Mod. Opt.*, vol. 45, pp.1111-1121, Jun. 1998.

[3] P. Chamorro-Posada, G.S. McDonald, "Helmholtz Dark Solitons," *Opt. Lett.*, vol. 15, pp. 827-827, May 2003

[4] J.M. Christian, G.S. McDonald and P. Chamorro-Posada, "Helmholtz-Manakov solitons," *Phys. Rev. E*, vol. 74, pp. 066612 1-5, Dec. 2006.

[5] J.M. Christian, G.S. McDonald and P. Chamorro-Posada, "Bistable Helmholtz solitons in cubic-quintic materials," *Phys. Rev. A*, vol. 76, pp. 033833 1-9, Sept. 2007.

[6] J.M. Christian, G.S. McDonald and P. Chamorro-Posada, "Bistable Helmholtz bright solitons in saturable materials," *J. Opt. Soc. Am. B*, vol. 26, pp. 2323-2330, Dec. 2009.

[7] P. Chamorro-Posada and G.S. McDonald, "Spatial Kerr soliton collisions at arbitrary angles," *Phys. Rev. E*, vol. 74, pp. 36609 1-4, Sept. 2006.

[8] J. Sánchez-Curto, P. Chamorro-Posada and G.S. McDonald, "Helmholtz solitons at nonlinear interfaces," Opt. Lett. , vol. 32, pp. 1126-1128, May 2007.

[9] J. Sánchez-Curto, P. Chamorro-Posada and G.S. McDonald, "Black and grey Helmholtz Kerr soliton refraction," *Phys. Rev. A*, vol. 83, pp. 013828 1-9, Jan. 2011.

[10] P. Chamorro-Posada and G.S. McDonald, "From Maxwell equations to Helmholtz solitons," in Nonlinear Guided Waves and their Applications (Topical Meeting) in CD-ROM (The Optical Society of America, Washington, DC, 2005), presentation WD3.

[11] J. Sánchez-Curto and P. Chamorro-Posada, "On a faster parallel implementation of the Split-Step Fourier method," *Parallel Computing*, vol. 34, pp. 539-549, Sep. 2008.

[12] M. Frigo and S. G. Johnson, "The Design and Implementation of FFTW3," *Proc. IEEE*, vol. 93, pp. 216–231, Feb. 2005.

[13] D.B. Kirk and W.W. Hwu, *Programming Massively Parallel Processor. A Hands-on approach.* Burlingon, MA: Elsevier, 2010.