Enhanced Functionality for Hardware-Based FDTD Accelerators

Petersen F. Curt [1], James P. Durbano [1], Michael R. Bodnar [2], Shouyuan Shi [2], and Mark S. Mirotznik [3]

[1] Accelerated Computing Division
EM Photonics, Inc., Newark, DE 19711, USA
pcurt@emphotonics.com, durbano@emphotonics.com

[2] Department of Electrical and Computer Engineering
University of Delaware, Newark, DE 19716, USA
bodnar@ee.udel.edu, sshi@ee.udel.edu

[3] Department of Electrical Engineering and Computer Science
The Catholic University, Washington, DC, 20064, USA
mirotznik@cua.edu

Abstract ⎯ In this paper we present several architectural enhancements to our previously published hardware-based FDTD acceleration platform. This includes the addition of several new sources, including H-polarized point sources, voltage and current sources, Gaussian beams, and user-defined sources, such as waveguide mode profiles. We also discuss the recent support for extending objects into the absorbing boundary, which minimizes non-physical back reflections. With the addition of these features, the FDTD acceleration hardware has become a more robust and powerful tool, enabling the rapid simulation of a wider breadth of applications, including antennas, waveguides, and optics.

Keywords ⎯ Finite Difference Methods, Simulation, Hardware Acceleration, Electromagnetic Analysis, FPGA.

I. INTRODUCTION

Although the need for advanced electromagnetic analysis in a variety of applications is readily apparent, the long runtimes associated with these simulations frequently limit what can be realistically modeled. Consequently, designers are often forced to artificially limit the scope of their simulations in order to analyze problems within a reasonable time frame. Fortunately, with the advent of hardware-based FDTD solvers, many of these limitations have been overcome, resulting in more designers relying on acceleration hardware for solutions to their computational needs. In [1, 2], we presented such a platform that demonstrated considerable improvements over software-based solutions in both speed and maximum problem size (Fig. 1). Although a significant achievement, this platform was limited in the types of problems it could solve. Specifically, the acceleration hardware was only capable of analyzing problems that incorporated either E-polarized point sources, uniform plane waves, or their temporally modulated variations. While this did allow a variety of simulations to be performed, a host of applications remained that could not be modeled, as they required more advanced source types. In this paper, we present our most recent architectural developments, which have focused on the incorporation of new source types. Notable platform additions include support for H-polarized point sources, voltage/current sources, and the introduction of a connecting boundary, which enables support for both Gaussian beams and also user-defined sources, such as waveguide mode profiles. We also discuss the recent support for extending objects into the absorbing boundary, which is vital to minimize non-physical back reflections associated with a variety of problems, including waveguides. Despite numerous publications on FDTD acceleration hardware, a fully 3D, hardware-based FDTD accelerator that encompasses such features has never been described until now [1-6].

Fig. 1. The FDTD hardware accelerator. This FPGA-based board provides the platform for the acceleration architecture. It includes the largest FPGA on the market and supports up to 16 GB DDR SDRAM. This platform has demonstrated performance comparable to 100-node PC clusters.

II. NEW SOURCE CONDITIONS

The original FDTD acceleration platform only supported two source types: single electric point sources and uniform plane waves. While the underlying hardware architecture was capable of quickly performing large simulations, the lack of support for more advanced source conditions prevented its application to numerous problems. For example, although the uniform plane wave source could be applied to many scattering problems, a spatially windowed plane wave (previously unsupported) was necessary to model "infinite" structures. Similarly, although a single point source (E polarized) was sufficient to model simple radiation patterns, the hardware accelerator did not support impedance-matched current sources and, thus, could not model many antenna structures. To expand the capabilities of the hardware accelerator, several new source conditions have been added to the architecture. In this section, we discuss four areas in which the hardware architecture has been enhanced, namely support for arbitrary magnitude and phase specification, magnetic field excitation, point source extensions, and the incorporation of a connecting boundary. These additions have directly enabled support for a wide range of sources, including magnetically polarized point sources, Gaussian beams, and guided mode profiles.

A. Arbitrary Magnitude and Phase Specification

Although the previous hardware design supported plane wave sources, their flexibility was limited. For example, it was not possible to specify an arbitrary magnitude (or phase) at the various points along the wavefront and, thus, the solver was limited to uniform plane waves. Although a temporal envelope function could be applied, spatially modulated waveforms, such as Gaussian beams, could not be implemented. Originally, uniform plane waves were supported because they were relatively easy to implement given that the magnitude along the phase front was constant.

However, it was clear that more advanced waveforms would be required to support a wider array of problems. Thus, the hardware solver was extended such that both the magnitude and phase of each point could be represented as a function of position.

Controlling the parameters of individual points was achieved by modifying the field update equations to incorporate both magnitude and phase terms in the source computation engine. Although this provided accurate answers, it also had the unfortunate consequence of doubling the memory requirements of the simulator and severely limiting performance. To remedy this, the incident field expressions were reformulated to use two new terms, dependant on the magnitude and phase, which have a limited range of [-1, 1]. To understand where these terms came from, assume that the incident electric and magnetic fields can be written as

$$E^{inc}(i, j, k, n) = Source(i, j, n - 0.5)$$
$$H^{inc}(i, j, k, n) = Source(i, j, n) \quad (1)$$

where:

$$Source(i, j, n) =$$
$$A(i, j) f_{envelope}(n) \left[\sin(\omega n \Delta t + \phi(i, j)) \right] . \quad (2)$$

We can then expand the sine term in equation (2)

$$A(i, j) f_{envelope}(n) \begin{bmatrix} \sin(\omega n \Delta t) \cos(\phi(i, j)) + \\ \cos(\omega n \Delta t) \sin(\phi(i, j)) \end{bmatrix} \quad (3)$$

and rearrange to obtain

$$A(i, j) \cos(\phi(i, j)) \left[f_{envelope}(n) \sin(\omega n \Delta t) \right] +$$
$$A(i, j) \sin(\phi(i, j)) \left[f_{envelope}(n) \cos(\omega n \Delta t) \right] . \quad (4)$$

We can then rewrite equation (4) as

$$term1(i, j) f_{\cos}(n) + term2(i, j) f_{\sin}(n) , \quad (5)$$

where

$$f_{\sin}(n) = f_{envelope}(n) \sin(\omega n \Delta t) ,$$
$$f_{\cos}(n) = f_{envelope}(n) \cos(\omega n \Delta t),$$
$$term1(i, j) = A(i, j) \cos(\phi(i, j)), \quad (6)$$
$$term2(i, j) = A(i, j) \sin(\phi(i, j)) .$$

By limiting the range to [-1, 1], it was possible to compress the terms and store them in the same amount of memory as the previous architecture. Specifically, in the previous design, a 32-bit floating-point value was used to store the source information. In this updated design, the same memory footprint (32 bits) was used, but was subdivided into two 16-bit numbers (one for magnitude, one for phase). By limiting the range of these values, there was no need for an exponent field (as would be required by a floating-point number) and the 16 bits could be used to accurately represent a decimal value with minimal precision loss as compared to a true floating-point representation.

B. Magnetic Field Excitation

In the previous acceleration architecture, sources were introduced by adding the incident electric field to the appropriate mesh points, as determined by a lookup table. However, the hardware lookup table and associated computational logic did not support the direct excitation of magnetic field components. In the new architecture, this lookup table was expanded, as was the control and computational datapaths, to provide support for the introduction of magnetic incident fields (Table 1). This extension directly enabled support for two additional source conditions: a magnetic point source (H polarized) and a connecting boundary, consisting of both electric and magnetic source fields. These new source conditions, as well as other point source additions, are now described.

Table 1. Hardware Lookup Table. This table details the four coefficients (A-D) that must be stored for each material used in the simulation. Entries are included for both electric and magnetic fields. The newest lookup table entries, which support magnetic sources, are shaded.

	Mat #	A	B	C	D
E	0	Ae(0)	Be(0)	Ce(0)	De(0)
	1	Ae(1)	Be(1)	Ce(1)	De(1)
	…	Ae(…)	Be(…)	Ce(…)	De(…)
H	0	Ah(0)	Bh(0)	Ch(0)	Dh(0)
	1	Ah(1)	Bh(1)	Ch(1)	Dh(1)
	…	Ah(…)	Bh(…)	Ch(…)	Dh(…)

C. Point Source Extensions

With respect to point sources, recall that the previous acceleration architecture only provided support for a single E-polarized point source. In order to expand the capabilities of hardware platform, two primary point source extensions were incorporated, specifically, magnetically polarized point sources, as well as direct support for voltage and current sources.

When modeling simple radiators, a magnetically polarized point source is of little value, as a corresponding E-polarized source could be used (assuming the structure's impedance is known). However, the importance of such a source becomes readily apparent when simulating devices of higher complexity, where the electric and magnetic components of the source cannot necessarily be related by real impedance. For example, the characteristic impedance of active devices, such as power supplies, amplifiers, and transistors, may change over the duration of simulation. In these cases, it would be impossible to model an equivalent source without using mathematically complex approximations, which would complicate the overall design and reduce the accelerator's performance. To this end, the acceleration architecture was modified to generate H-polarized point sources.

Support for both E- and H-polarized sources was achieved by designing more generalized source computation engines that can be quickly switched to produce either electric or magnetic incident fields as needed. At the end of each half timestep, central control instructs the source computation engines to reconfigure themselves to perform the appropriate field update. This switch command is pipelined with the existing data and therefore does not impede the accelerator's performance.

After providing support for magnetic point sources, the hardware solver was then extended to provide direct support for voltage and current sources. By specifying the source impedance, in addition to the time-domain waveform, Thévenin and Norton equivalents could be constructed that made use of the underlying support of the E- and H-polarized point sources [4, 7]. This capability is useful for simulating the operation of circuits, microwave amplifiers, and custom antennas designed for a matched input impedance.

D. Connecting Boundary

A connecting boundary is the cornerstone of the FDTD total-field/scattered-field (TF/SF) formulation and is used to introduce a variety of incident source conditions, as well as easily detect scattering from an object [4]. These scattered results can then be analyzed

using post-processing techniques, such as Fourier transforms, near-to-far field transformations, and also radar cross-section calculations. Although the previous processing architecture supported a scattered-field formulation, which made determining the scattered fields quite easy, this formulation could not support extensions for more advanced sources, such as spatially modulated plane waves. Thus, it became necessary to implement a connecting boundary within the hardware solver.

The implementation of the hardware's connecting boundary was achieved by modifying the preprocessing performed on the host PC. Before a hardware simulation begins, the host analyzes the input parameters and loads the hardware accelerator's on-board memory with the appropriate runtime data. In this manner, the host can construct an effective connecting boundary by calculating the particular magnitude and phase required at every point along the boundary such that the incident wave propagates forward, and the associated backward traveling wave is cancelled. Specifically, two 32-bit floating-values are calculated along the connecting boundary before the simulation begins. These additional computations add minimal overhead and do not affect performance as the bulk of the computational time occurs while iteratively solving the solution space over many thousands of time steps.

By combining these updates with the recent addition of arbitrary magnitude and phase specification, the connecting boundary enabled support for a variety of additional source conditions, including Gaussian beams. Not only does this extend the capabilities of the accelerated solver, it can also lead to a reduction of problem size. For instance, applications such as optical filters, switches, and mode converters require the guiding of light from a waveguide onto another device. If the underlying platform cannot model the waveguide output directly (e.g., guided mode profile), it becomes necessary to simulate both the input waveguide along with the ultimate device under test. This leads to much larger mesh sizes, which results in longer simulation times, as well as increased numerical dispersion. However, if the guided mode profile can be entered directly, modeling the source waveguide along with the device under test is no longer necessary. Thus, by using the hardware's connecting boundary capability, the input waveguide problem can be solved separately, using the result to apply a known guided source profile directly at the device.

In this section we presented four areas in which the FDTD acceleration platform was extended, including arbitrary magnitude and phase specification, magnetic field excitation, point source extensions, and a connecting boundary. The incorporation of these capabilities into the hardware solver enabled a much wider array of incident source conditions, such as voltage and current sources, Gaussian beams, and guided mode profiles, and thus directly provides support for a much broader application base, including advanced antenna and waveguide simulations. The next step in this process was to enhance the capabilities of the absorbing boundary conditions in order to further broaden the domain of the accelerated solver.

III. ADVANCED BOUNDARY CONDITION SUPPORT

When performing FDTD simulations, an appropriate absorbing boundary condition is necessary to prevent the non-physical reflection of outward-going waves back into the observation region. To this end, the original acceleration architecture incorporated Perfectly Matched Layer (PML) absorbing boundaries [8]. These boundaries "match" the outermost boundary layers to the background material of the computational space and then attenuate the outgoing fields. Unfortunately, if a material other than the background material were placed next to the boundary, such as the structure being simulated, the boundary nodes would no longer "match" the computational region and back reflections would occur. Additionally, if the device being simulated is part of a larger system, it may be necessary to extend a piece of the device into the boundary to give the appearance that it is part of a complete system. To resolve these issues, the hardware accelerator was modified to enable the extension of objects in the boundary layers. This capability was achieved by allowing mesh material information to be stored with both computation and boundary nodes. Previously, mesh material information was only stored for nodes inside the computation region. This was acceptable because nodes in the absorbing boundary were always assumed to be free space. Additional storage has been added in the new architecture to maintain information about material properties in the absorbing boundaries. Now, the acceleration engine can use this information to calculate the appropriate PML coefficients before they are used in the boundary updates.

The new architecture allows any dielectric structure to be extended into the boundary to create the appearance of an infinite space. Although an outward going wave may still be reflected after reaching the last absorbing layer, enough layers can be provided such that the magnitude of the reflection is inconsequential, as is the case in any software implementation. Specific

applications that benefit from this implementation include optical waveguides, periodic structures, and substrates with etched defects.

When coupled with the new source configuration, the extension of materials into the absorbing boundary enables much more efficient simulations for a variety of applications. In the next section, we present two such applications, the design of a coupled resonator optical waveguide (CROW) and a structure exhibiting left-handed behavior (LHM), and demonstrate the accuracy of the hardware accelerator.

IV. RESULTS

In order verify the accuracy of these platform additions, a CROW structure was modeled and compared against published results [9]. A CROW structure is a common building block in many optical communication systems because of its ability to store/delay optical signals. Unfortunately, the exact behavior of CROW structures is difficult to determine analytically and, thus requires computational methods for their analysis. After extending the capabilities of the hardware platform, as described above, the accelerated solver could now be applied to such devices.

The following example uses a CROW as a delay element. The input and output waveguides run parallel to one another, but are separated by two ring resonators. This particular structure is made of silicon and placed in a free space environment. A 16 layer PML boundary surrounds the computational space. Because this structure is a component in a larger system, the input and output waveguides must be extended into the absorbing boundary. The input wave is introduced at one end of the left waveguide using the new connecting boundary source condition. The problem was simulated using a Gaussian modulated input pulse of bandwidth 0.2% for 200,000 timesteps, which is equivalent to about 20 ps. Fig. 2 shows the transient results at several points in the simulation.

The hardware accelerator was able to solve this 5.71 million node problem in 7.55 hours, corresponding to roughly 42 Mnps of sustained computational throughput (see [1, 2] for a detailed discussion of this performance measurement). Fig. 2 clearly shows how the input pulse is placed on the output after a delay of several picoseconds, which is in agreement with the two-dimensional simulation results presented in [9].

Next, we simulated a twenty-ring resonator structure, with a computational region size of 49.61 million FDTD nodes. This simulation was performed for 450,000 time steps and required 147.6 hours of computation time. Previously, such a simulation would have been impossible in a standard desktop computing environment. The CROW achieved a delay of 15.78 ps delay (Fig. 3). In comparing these results against those of the two-ring structure, note that the resonant frequency is the same, but the overall transmission efficiency has decreased, which is due to the increased propagation losses associated with additional rings, including bending loss, roughness loss, and coupling loss. Furthermore, these simulations have revealed previously unknown scalability relationships that will simplify future CROW structure analyses [10]. A more thorough treatment of these results, including simulation data for five- and ten-ring structures, can be found in [10].
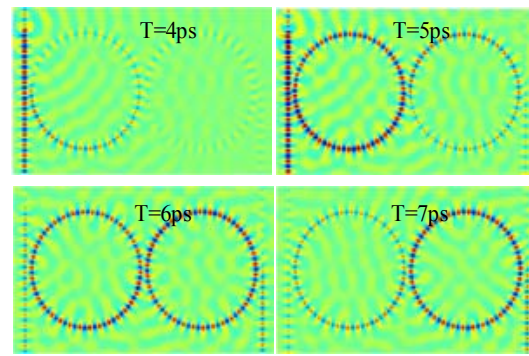

Fig. 2. Propagation through the 2-ring CROW structure. Here we see snapshots of the incident pulse as it travels through the ring resonators and, ultimately, to the output waveguide. The delay from input to output directly corresponds to previously published results.
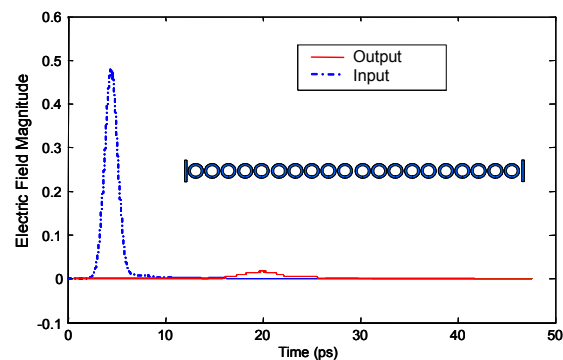

Fig. 3. Propagation through the 20-ring CROW structure. Here we see the input and output pulses associated with the 20-ring CROW. Note that the output pulse is a delayed (and attenuated) version of the input pulse.

Next, we modeled an LHM structure, which was composed of a periodic array of unit cells (Fig. 4).

Each unit cell consisted of an SRR structure patterned on one side of a dielectric substrate (ε_r = 3.4) (Fig. 5) and a metallic wire (0.5 mm × 1.0 mm × 20 mm) patterned on the other side of the substrate. The thickness of the SRR, the wire, and the substrate were chosen to be 0.5 mm. This unit cell, which measured 20 mm × 20 mm × 20 mm, was then replicated twenty times in the YZ plane, ten times in the XY plane, and three times in XZ plane to form an array of 600 SRR+wire pairs (lattice constant a = 30 mm). This LHM structure was placed inside a two-dimensional waveguide formed by parallel, perfectly conducting plates, each 40.15 cm wide and 24 cm long, in order to provide vertical confinement and more accurately reflect a guided mode source.

Ultimately, the computational domain for this structure measured 24 cm × 40.15 cm × 5 cm, or 481 × 803 × 101 cells (~ 40 million nodes), not including the PML absorbing boundary region. For an FDTD discretization of 0.5 mm, the corresponding timestep is 9.63 x 10^{-4} ns. The 20,000 timestep simulation required less than 2 GB of memory and approximately 5.5 hours of computation time on the Celerity™ accelerator card [2]. However, because the acceleration platform contains 16 GB of RAM, an LHM structure consisting of up to 4,700 unit cells could be simulated.
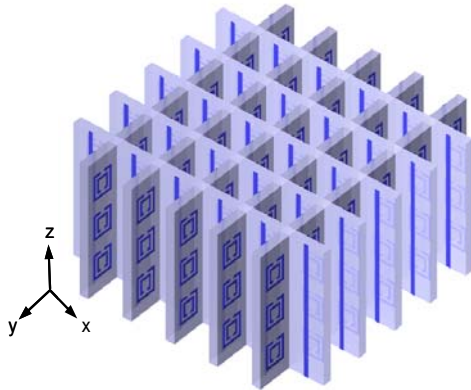

Fig. 4. LHM structure. This is a model (scaled down) of the LHM structure simulated on the hardware platform. From this model, it is easy to see both the SRRs and the wires.

To determine frequencies at which negative permittivity and permeability might exist, the transmission spectra of the LHM structure is measured to identify stopband frequency regions. Once these frequency regions are identified, a continuous wave is used to examine the steady-state behavior of the LHM structure. For this particular LHM structure, a broadband, z-polarized windowed plane wave, propagating along the x-direction, was used. A point detector was placed at the far end of the LHM structure in order to measure the

transmission spectra through the periodic array of SRRs. Specifically, the point detector recorded the time-varying electric and magnetic field amplitudes, which were then normalized to the source. The frequency response was then obtained by performing a Fast Fourier Transform (FFT) of the normalized data. Two simulations were performed: one with the SRRs alone (without wires) and one with both SRRs and wires. From these results, we see that the SRR structure, at resonance, has a stopband 38 dB down at 2.75 GHz (Fig. 6). For both wires and SRRs together, a small passband exists between 3.75 GHz and 4.25 GHz, near the resonance of the rings.
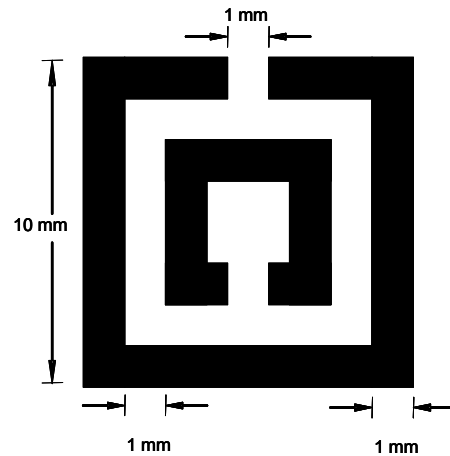

Fig. 5. A split-ring resonator. Here we see the SRR used in the LHM model. The length and width of each SRR was 10 mm, and the azimuthal inter-ring gaps were 1 mm.
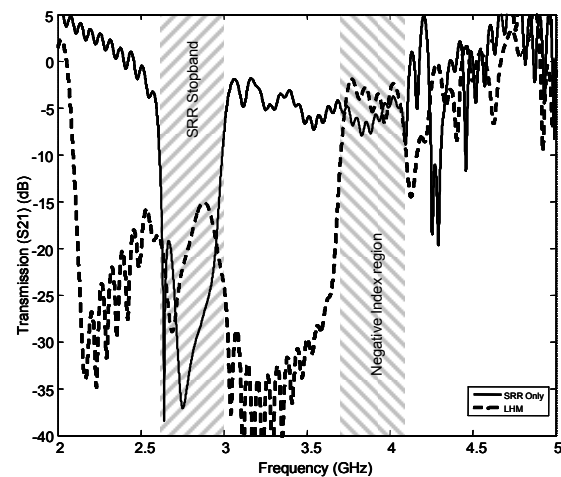

Fig. 6. Transmission Spectra of SRRs and SRRs + Wires (LHM). From these results, we see that the SRR structure, at resonance has a stopband 38 dB down at 2.75 GHz, with regions of positive and negative permeability on both sides. For both wires and SRRs together, a small passband exists between 3.75 GHz and 4.25 GHz, near the resonance of the rings.

To numerically demonstrate that the composite structure formed from the combination of SRRs and wires posses a negative index of refraction, we constructed a 26° prism and embedded the LHM structure within it. By measuring the direction of the power leaving the prism, it is possible to calculate the index of refraction using Snell's Law.

The prism/LHM structure was simulated using a windowed plane wave source. The frequency of the incident beam was varied according to the transmission spectra results previously obtained (Fig. 6). Specifically, simulations were performed at frequencies of 2.88 GHz, 3.0 GHz, 3.8 GHz, and 3.9 GHz. The first frequency, 2.88 GHz, was chosen because it lies within the resonance band of the SRR structure. Such a frequency finds the structure to be highly attenuating, and a positive index of refraction was seen. Next, we shifted the frequency away from the resonance band of the SRRs to the edge of the stop band at 3.0 GHz. At this frequency, the overall LHM structure is highly dispersive, but it still possesses a positive index of refraction. The next frequency tested was 3.8 GHz, for which the structure is less dispersive and we note that the index of refraction has been slightly changed (n = -1.4356). Specifically, the wave front is now directed away from the surface of the prism, indicating a negative refraction index (Fig. 7). Finally, the structure was excited using a 3.9 GHz source, for which the structure is less dispersive but with a positive index of refraction.
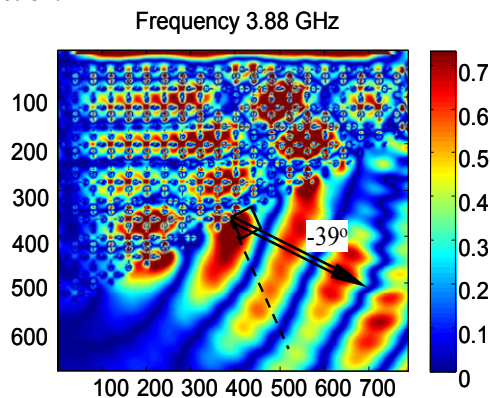

Fig. 7. Demonstrated Negative Refractive Index. Here we see the wave exiting the LHM/Prism structure. Notice that the wave front leaves the prism to the left of the normal, indicating negative refraction.

V. CONCLUSION

In this paper we presented several architectural additions to our previously published hardware-based FDTD acceleration platform, including H-polarized point sources, voltage and current sources, Gaussian beams, spatially windowed plane waves, and user-defined sources, such as waveguide mode profiles. We also discussed the recent support for extending objects into the absorbing boundary, which minimizes non-physical back reflections. With the addition of these features, the FDTD acceleration hardware has become a more robust and powerful tool, enabling the rapid simulation of a wide breadth of applications, including antennas, waveguide structures, and optics. Ultimately, these features continue to demonstrate the capabilities of hardware-based acceleration tools for computational electromagnetics and prove, once again, the viability of such platforms for both academic and industrial applications.

REFERENCES

[1] J. P. Durbano, F. E. Ortiz, J. R. Humphrey, D. W. Prather, and M. S. Mirotznik, "Hardware implementation of a three-dimensional Finite-Difference Time-Domain algorithm," *IEEE Antennas and Wireless Propagation Letters*, vol. 2, pp. 54-57, 2003.

[2] J. P. Durbano, J. R. Humphrey, F. E. Ortiz, P. F. Curt, D. W. Prather, and M. S. Mirotznik, "Hardware Acceleration of the 3D Finite-Difference Time-Domain Method," presented at IEEE AP-S International Symposium on Antennas and Propagation, 2004.

[3] W. Chen, P. Kosmas, M. Leeser, and C. Rappaport, "An FPGA Implementation of the Two-Dimensional Finite-Difference Time-Domain (FDTD) Algorithm," presented at The Twelfth ACM International Symposium on Field-Programmable Gate Arrays (FPGA), 2004.

[4] A. Taflove and S. C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, 3rd ed. Norwood: Artech House, 2005.

[5] L. Verducci, P. Placidi, G. Matrella, L. Roselli, F. Alimenti, P. Ciampolini, and A. Scorzoni, "A feasibility study about a custom hardware implementation of the FDTD algorithm," presented at The 27th General Assembly of the URSI, Maastricht, Netherlands, 2002.

[6] R. N. Schneider, M. M. Okoniewski, and L. E. Turner, "A Software-Coupled 2D FDTD Hardware Accelerator," presented at IEEE AP-S International Symposium on Antennas and Propagation, Monterey, CA, 2004.

[7] M. M. Mano, *Digtial Design*, 3rd ed. Upper Sadle River: Prentice Hall, 2001.

[8] J.-P. Berenger, "Three-Dimensional Perfectly Matched Layer for the Absorption of Electromagnetic Waves," *Journal of Computational Physics*, vol. 127, pp. 363-379, 1996.

[9] J. K. S. Poon, J. Scheuer, S. Mookherjea, G. T. Paloczi, Y. Y. Huang, and A. Yariv, "Matrix analysis of microring coupled-resonator optical waveguides," *Optics Express*, vol. 12, pp. 90-103, 2004.

[10] J. P. Durbano, A. S. Sharkawy, S. Shi, F. E. Ortiz, and P. F. Curt, "Accelerated Design Tools for Nanophotonic Devices and Applications," in *Handbook of Nanoscience, Engineering, and Technology*, D. Brenner, Ed., 2nd ed: CRC Press, 2006.

Petersen F. Curt is currently a Senior Engineer in the Accelerated Computing Division of EM Photonics, Inc. He received his B.Cp.E and M.E.E degrees from the University of Delaware in 2003, and 2005. Mr. Curt's research interests i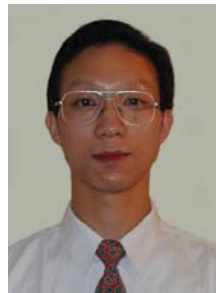nclude hardware acceleration of computationally-intense scientific algorithms and the study of reconfigurable logic in embedded systems.

James P. Durbano, Ph.D., is currently the Chief Hardware Architect of EM Photonics, Inc., Newark, DE and an Assistant Professor in the Computer and Information Sciences Department at the University of Delaware. He received the B.Cp.E., M.E.E., and Ph.D. degrees from the University of Delaware in 2000, 2002, and 2004, respectively. Dr. Durbano's research interests involve hardware-based acceleration of scientific computing using a variety of platforms, including field-programmable gate arrays (FPGAs), digital signal processors (DSPs), and graphics processing units (GPUs). Much of his work in this area has been focused on accelerating the Finite-Difference Time-Domain method.

Michael R. Bodnar is currently a M.S. student in the Electrical and Computer Engineering Department at the University of Delaware. He received a B.Cp.E. from the University of Delaware in 2005. Mr. Bodnar's research interests include hardware acceleration of computationally-intense scientific algorithms.

Shouyuan Shi, Ph.D., is currently an Associate Professor in the Department of Electrical and Computer Engineering, University of Delaware, Newark. He received the B.S., M.S., and Ph.D. degrees from Xidian University, Xi'an, China in 1991, 1994 and 1997 respectively, all in electrical engineering. His research interests include electromagnetic numerical modeling, electromagnetic imaging, antenna design, diffractive optical elements, RF microphotonic modulators, microdisk lasers, left-handed materials, and photonic crystals and their applications.

Mark S. Mirotznik, Ph.D, is currently an associate professor of Electrical Engineering and Computer Science at The Catholic University of America. He received a BSEE degree from Bradley University, Peoria Illinois, in 1988, MS degrees in Electrical Engineering and Biomedical Engineering from The University of Pennsylvania in 1991 and a Ph.D. in Biomedical Engineering from The University of Pennsylvania in 1992. Dr. Mirotznik's research interests include computational electromagnetics, biomedical instrumentation and micro/nano scale optical elements.