Progress in Parallel Implementation of the Multilevel Plane Wave Time Domain Algorithm

Yang Liu*[1], Hakan Bağcı[2], and Eric Michielssen[1]
(1) Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA
(2) Division of Computer, Electrical, and Mathematical Sciences and Engineering, King Abdullah University of Science and Technology, Thuwal 23955-6900, KSA

The computational complexity and memory requirements of classical schemes for evaluating transient electromagnetic fields produced by N_s dipoles active for N_t time steps scale as $O(N_t N_s^2)$ and $O(N_s^2)$, respectively. The multilevel plane wave time domain (PWTD) algorithm [A.A. Ergin et al., Antennas and Propagation Magazine, IEEE, vol. 41, pp. 39-52, 1999], viz. the extension of the frequency domain fast multipole method (FMM) to the time domain, reduces the above costs to $O(N_t N_s \log^2 N_s)$ and $O(N_s^\alpha)$ with $\alpha = 1.5$ for surface current distributions and $\alpha = 4/3$ for volumetric ones. Its favorable computational and memory costs notwithstanding, serial implementations of the PWTD scheme unfortunately remain somewhat limited in scope and ill-suited to tackle complex real-world scattering problems, and parallel implementations are called for.

In [Y. Liu et al., in *URSI Digest*, 2012], we presented a parallel implementation of the multilevel PWTD algorithm on a distributed-memory hybrid CPU-GPU cluster. The proposed implementation allocates the computation and memory loads among compute-nodes using a hierarchical spatial, angular, and temporal partitioning strategy that extends the parallel FMM scheme in [J. Fostier et al., Electr. Lett., vol.44, pp.1111-1113, 2008; O. Ergul et al., Trans. Ant. Prop., vol.57, no.6, pp.1740-1750, 2009] to PWTD. This parallelization strategy was proven scalable and exhibits favorable load balance and computation-to-communication ratios. Here we propose several techniques to further improve the CPU, memory, and communication efficiency of the CPU-only version of this parallelization scheme. (i) The scheme's memory requirements are reduced by exploiting temporal sparsity exhibited by typical high-frequency plane wave pulses via adaptive wavelet packet based compression. (ii) The PWTD geometric tree generation is parallelized; specifically, compute nodes generate partial PWTD trees from knowledge of subdomains they control and communication completes partition of the PWTD tree. (iii) Hybrid MPI/OpenMP schemes are applied to avoid intra-node MPI communication and memory duplication. The improved parallel scheme is used to evaluate transient electromagnetic fields generated by constellations of $N_s = 10{,}000{,}000$ surface-bound and $N_s = 64{,}000{,}000$ volumetrically distributed dipoles, a vast improvement over past parallel PWTD implementations.