Multi-GPU-based Acceleration of the Explicit Time Domain Volume Integral Equation Solver Using MPI-OpenACC

Saber Feki, Ahmed Al-Jarro*, and Hakan Bağcı
King Abdullah University of Science and Technology (KAUST)
Thuwal 23955-6900, Kingdom of Saudi Arabia
http://maxwell.kaust.edu.sa

An explicit marching-on-in-time (MOT)-based time-domain volume integral equation (TDVIE) solver has recently been developed for characterizing transient electromagnetic wave interactions on arbitrarily shaped dielectric bodies (A. Al-Jarro et al., *IEEE Trans. Antennas Propag.*, vol. 60, no. 11, 2012). The solver discretizes the spatio-temporal convolutions of the source fields with the background medium's Green function using nodal discretization in space and linear interpolation in time. The Green tensor, which involves second order spatial and temporal derivatives, is computed using finite differences on the temporal and spatial grid. A predictor-corrector algorithm is used to maintain the stability of the MOT scheme. The simplicity of the discretization scheme permits the computation of the discretized spatio-temporal convolutions on the fly during time marching; no "interaction" matrices are pre-computed or stored resulting in a memory efficient scheme. As a result, most often the applicability of this solver to the characterization of wave interactions on electrically large structures is limited by the computation time but not the memory.

To render the explicit MOT-TDVIE solver useful in real-life scenarios, where several millions of discretization elements are typical; a parallelization scheme, which makes use of unstructured mesh partitioning, has been proposed (A. Al-Jarro, and H. Bağcı, *ACES*, 2013). The partitioning strategy ensures the even distribution of the workload needed for computing the discretized spatio-temporal convolutions. Additionally, even distribution of the spatial finite difference computations is achieved through minimization of the number of boundary elements that reside in between the partitions. To reduce the communication costs, global/collective MPI operations are eliminated using a one-way pipeline communication strategy, so called "rotating tiles" paradigm. The performance of the parallelized MOT-TDVIE solver is enhanced through hybridization with OpenMP for the intra-process data exchange.

In this work, we build upon this hybrid MPI/OpenMP framework and extend it for execution on a cluster of Graphics Processing Units (GPUs). The GPU-accelerated MOT-TDVIE solver exploits the OpenACC application program interface (API) (http://www.openacc-standard.org), which annotates the MPI code with compiler directives; thus, targeting the GPUs in a fashion similar to how OpenMP targets shared memory architectures on CPUs. This high-level approach permits porting of the code originally parallelized for CPUs onto the multi-GPU environment with minimal re-coding effort. Performance results, which demonstrate the computational efficiency of the MOT-TDVIE solver's MPI/OpenACC implementation via comparison to that of its counterpart MPI/OpenMP implementation, will be presented at the conference.