GPU Implementation of Parallelized Microwave Tomography Algorithm

Michael Holman, and Sima Noghanian
Department of Electrical Engineering, University of North Dakota

Microwave tomography (MWT) has good potential to be used for medical imaging, however, most of MWT algorithms rely on local optimization methods and need regularization to find the a solution to inverse scattering problems. By using global optimization method for optimization, the non-deterministic nature of the genetic algorithm allows the inverse solver to avoid local minima without the use of regularization methods such as Tikhonov regularization. By not relying on regularization assumptions, high contrast areas of the imaging target can be resolved, whereas regularizations assume smooth dielectric contrast gradients. Resolving areas of high permittivity contrast is necessary to detect small tumors, less than a millimeter in length, as required for effective treatment. Our goal is to implement a fast MWT algorithm based on Finite Difference Time Domain (FDTD) forward solver and global optimization methods. In this regards, we propose the use of graphics processing unit (GPU) for FDTD computation. We have developed a FDTD program using NVidia's CUDA C language. The GPU implemented FDTD simulation was tested to yield 100-fold speed increase from standard Central Processing Unit (CPU) FDTD simulations.

Each cell in the FDTD problem space was executed by a single thread. Given that each block on the GPU can only sustain 512 threads, the simulation was computed on multiple streaming multiprocessors. Because of this, the threads were not able to synchronize at the end of computation in order to wait until all threads completed their tasks before moving on to the next time step. In order to get around this, we developed the algorithm such that it computed a single time step using a single kernel for the electric and magnetic field updates. The field-update kernels were initialized by the CPU for every time step. The field values were stored on the GPU for the entire simulation, in order to avoid memory bottlenecks. The field constants, a function of the permittivity and conductivity of the cell region, were transferred to the GPU before the simulation.

The forward solver was tested on University of North Dakota's Shale High Performance Computing cluster. Four workstations of the Shale cluster are equipped with Tesla S1070 GPUs. A CPU-computed version of the GPU FDTD program was also developed to test runtime speed. Field values at one cell in the simulation were sampled in order to test whether the GPU implementation introduced any calculation errors. While the CPU program took 433 seconds to finish a single FDTD simulation, the GPU parallelized FDTD completed simulation in a mere four seconds. Comparisons of the computed field values throughout the simulation showed that the GPU FDTD produced identical results to those of the CPU algorithm. With this speed, a MWT GA inverse solver will take a matter of hours to converge on the correct image, as opposed to a matter of weeks.