GPU Acceleration using CUDA for Computational Electromagnetics

Constantine Sideris[(1)]

(1) University of Southern California, Los Angeles, CA 90089, USA (csideris@usc.edu)

Although initially designed for graphics rendering applications, Graphics Processing Units (GPUs), due to their massively parallel many-core nature, have seen widespread use in accelerating many computationally expensive algorithms in recent years. In particular, GPUs have revolutionized the field of machine learning, enabling training deep neural network models which are multiple orders of magnitude larger than previously possible. Even consumer-grade GPUs, which have become very affordable and can be found in personal workstations (e.g., the NVIDIA GeForce RTX 4090), can achieve upwards of 80 TeraFLOPS computing performance.

GPU hardware has also been used in recent years to accelerate electromagnetic solvers, especially ones based on finite difference methods. Many commercial solvers include GPU acceleration support, and several open-source solutions exist as well. In this talk, I will begin by introducing the typical architecture of a modern GPU and discuss how GPUs can be used to effectively accelerate the solution of computational electromagnetics algorithms. I will start with examples based on finite difference methods and subsequently will cover our recent work in hardware acceleration of high-order integral equation methods [1]. Specific examples will be presented, as well as timing comparisons between CPU and GPU, of 3D Maxwell simulation of complicated nanophotonic devices reaching up to millions of unknowns.

The authors gratefully acknowledge support by the Air Force Office of Scientific Research (FA9550-20-1-0087) and the National Science Foundation (CCF-2047433) and the National Science Foundation (CCF-CCF-2047433).

[1] Garza, E. and Sideris, C. (2022). "Fast Inverse Design of 3D Nanophotonic Devices Using Boundary Integral Methods." ACS Photonics. doi.org/10.1021/acsphotonics.2c01072

978-1-7335096-7-1© 2024 ACES