IEEE HPEC 2016 NVIDIA Tutorial Abstract

Title: **GPU Computing – CUDA, Graph Analytics and Deep Learning**
Speakers: Larry Brown Ph.D. & Joe Eaton Ph.D.

This tutorial will have three main sections. The first section will provide an overview of GPU computing, the NVIDIA hardware roadmap and software ecosystem. We will discuss key factors driving recent developments in data analytics software, and the key trends in high performance computing and how they have impacted the evolution of modern computing architectures. There will be an overview on how to use the GPU including a brief intro to CUDA programming.

The second section will focus on the new nvGRAPH library. We will work through some common applications using the library, starting from how to format data sets and load graphs into the library's data structures. Then we will work through examples using the PageRank, SSSP and Widest Path algorithms. Subgraph extraction will be used to focus on a smaller graph of interest from a large graph of data. Future work on nvGRAPH will provide a graphBLAS compatible interface, and we will discuss the connection between the core features of nvGRAPH and the core of the graphBLAS.

The final section of the tutorial will be on the topic of Deep Learning. We will describe the impact of GPUs in the area of Deep Learning, and the ramifications for system development. We will describe a number of use cases of GPU accelerated deep learning, and review the tools and techniques for getting started with deep learning development. There will be a demo of the open-source DIGITS software by NVIDIA, which provides a web browser interface for applying deep learning to image processing.

Speaker Biographies

Larry Brown is a Sr. Solution Architect with NVIDIA, where he helps customers design and deploy GPU accelerated workflows in high performance computing and data analytics. Larry has over 15 years of experience designing, implementing and supporting a variety of advanced software and hardware systems for defense and national security applications. He has designed electro-optical systems for head-mounted displays and training simulators, developed GIS applications for multi-touch displays, and adapted computer vision code in UGVs for the GPU. He has a Ph.D. from the Johns Hopkins University in the area of Vision Science and a graduate certificate in Software Engineering from the University of Colorado.

Joe Eaton leads NVIDIA's efforts to build a world class graph analytics platform, to accelerate enterprise and national security applications. Joe managed the AmgX, cuSPARSE, and cuSOLVER CUDA libraries before shifting his attention to graphs and big data and creating nvGRAPH. Prior to NVGRAPH development, Joe has worked on RDF and SPARQL graph databases, reservoir simulation for the oil and gas industry, data management and predictive modeling. Joe has a BSc and MS in Mechanical Engineering, and a Ph.D. in Computational and Applied Mathematics from the University of Texas at Austin, TICAM program.