

Seminar



HPC/AI Convergence, Workload Matters

Abstract. In the era of convergence of HPC with AI, capturing and understanding your workload is critical to getting the most from the hardware you have or the hardware you will acquire. Couple that with the post Moore's Law world we live in now make the balance of CPU and GPU more important than ever. The talk will explore capturing Workload data and important aspects of GPU nodes for acceleration at the shared infrastructure level. The discussion will conclude with available performance features including Tensor Cores and the differences there in between Turing and Volta GPU's. I also briefly cover Automated Mixed Precision at the applications layer.

Biography. Mr. Mark Hill joined NVIDIA in September 2017 as a Sr. Solutions Architect supporting Academic and Research organizations on the east coast. In this role, he works to accelerate research and data center workloads with GPU's. Mark draws on his broad experience in R&D, hardware, systems software, networking, and IT infrastructure operations, within the Manufacturing, Defense, and Life Sciences industries.

Prior to joining NVIDIA, Mark was the Chief Technologist for Hyperscale on the east coast on HPE's technology team. His focus in the 6+ years at HPE was High Performance Computing, Big Data and HPC Storage Solutions, Mark lead the design of solutions and architecture for HPE's higher education and commercial research customers. He also developed an HPC Storage Architecture from a business and technology approach that was intended to scale from 1PB to 100's of PB and ultimately 1000PB. Prior to joining HPE, Mark spent over 23 years at Corning Incorporated where he started in the Mathematics and Statistical Analysis research group, supporting all disciplines in a world renowned RD&E organization rapidly growing to become Corning's HPC Architect responsible for Scientific Computing, Manufacturing, and Enterprise Systems Operations, Strategy, and Architecture globally. Mark was the principle in the re-engineering design of Corning's Scientific Computing Services to increase the speed of innovation and to resolve scientific and engineering problems encountered by Corning's Global technical community.

Before Corning, Mark worked 5 years in software and systems engineering capacities contracting to GE, Kodak, and Polaroid. Developing solutions in the Defense and Manufacturing sectors. Mark is presently transitioning the pursuit of his MS/PhD in Computer Science at Binghamton University to NOVA Southeastern University. Mark holds an AS in Engineering and a BS in Mathematics from Elmira College.



Mr. Mark Hill
Senior Solutions Architect
NVIDIA Corporation

+1-561-227-8805
mhill@nvidia.com
<https://www.nvidia.com/en-us/>

Friday February 7, 2020
Time: 2:00 - 2:50pm
Room: Bate 1009

Contact: Dr. Kamran Sartipi
Dept. of Computer Science, ECU
www.cs.ecu.edu/sartipi/CSseminar/

Seminar



February 7, 2020

<http://www.cs.ecu.edu/sartipi/CSseminar/images/Nvidia-07Feb2020.MOV>

Contact: Dr. Kamran Sartipi
Dept. of Computer Science, ECU
www.cs.ecu.edu/sartipi/CSseminar/



HPC AND AI CONVERGENCE: WORKLOAD MATTERS

Mark Hill

Sr. Solutions Architect, Higher Education and Research

January 2019



AI - A NEW INSTRUMENT FOR SCIENCE

HPC

- > +40 years of algorithms based on first principles theory.
- > Proven statistical models for accurate results

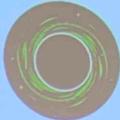
AI

- > Neural networks that learn patterns from large data sets.
- > Previously unmanageable data sets

Dramatically Improves Accuracy and Time-to-Solution



Commercially viable fusion energy



Understanding cosmological dark energy and matter



Clinically viable precision medicine



Improvement and validation of the Standard Model of Physics



Climate/weather forecasts with ultra-high fidelity





