Northeastern University

Advisors: Madeline Leger, Xiang Gong, Professor David Kaeli

Students: Jason Booth, Spencer Hance, Maddy Leger, Zachary Marcus, Carter McCardwell, Selean Ridley, Tristan Sweeney

Meet the Team

Our diverse team is made up of 7 students in Northeastern's Electrical & Computer Engineering Department. All of the students are part of NUCAR, a computer architecture research group led by Professor David Kaeli, and all have research interests in HPC, big data, and parallelization. We worked with students from Auburn University to prepare for SCC.

Jason Booth
Sophomore, Computer Engineering/Science

Spencer Hance
Third Year, Computer Engineering

Maddy Leger
Third Year, Computer Engineering, Business Minor

Zachary Marcus
Third Year, Computer Engineering/Science

Carter McCardwell
Third Year, Computer Engineering

Selean Ridley
Sophomore, Computer Engineering

Tristan Sweeney
Third Year, Computer Engineering/Science

Our System

Our cluster is a two node system, featuring two dense nodes able to tackle large workloads. Each node features:

- Two Intel Xeon E5 – 2630 v4 processors
- One AMD Firepro 9150 (5.07 TF @ 22.1 GF/W: 16GB GDDR5)
- One AMD R9 Nano (8.19 TF @ 47.92 GF/W: 4GB HBM)
- 128GB RAM

The nodes are connected over Infiniband to maintain a low communication overhead. The chosen hardware achieves both high performance and high energy efficiency. The decision to take Intel CPUs with AMD GPUs was to allow AMD's ROCm software stack its largest performance benefits, which can best be seen in heterogeneous compute solutions.

Applications & Strategies

Our Strategy: One team member is the "system administrator", and each remaining team member focuses on one specific app. This way, we can effectively focus on each app and how it interacts with our unique hardware setup. All members work on prepping for the mystery app.

LINPACK

The open source HPL-GPU leverages heterogeneous architectures and a high performance DGEMM for AMD GPUs. Featuring dynamic workload distribution and power modification, it should make excellent use of the high speed interconnect and 128GB of RAM per node.

HPCG

HPCG is a compute-intensive application that will utilize our Xeon processors. Although we will not be able to employ our GPUs, we can still provide a competitive score. Additionally, our two-node setup allows us less overhead with irregular memory accesses.

ParConnect

ParConnect is an I/O heavy application. We will be running Password Cracking and ParConnect concurrently to maximize the CPU/GPU utilization, since so much of ParConnect is I/O operations.

Password Cracking

Password cracking is attempting to reverse one-way cryptographic hashes. The two algorithms, MD5 and Bcrypt, are incredibly computationally intensive for both GPUs and CPUs and our custom software will ensure the best speed possible on our hardware.

Paraview

Paraview will be able to take advantage of our CPUs for data processing, and our powerful GPUs for rendering. Using filters in the appropriate order to cull the data will reduce memory overhead and computation time.

Mystery App

Our strategy for the mystery app is to not only take full advantage of the strength of our system, fueled by AMD and Intel processors, but also the versatility of our experienced members, coming back for a repeat performance.

Why We'll Win

We are utilizing AMD's Radeon Open Compute Platform, which:

- Is optimized for HPC and facilitating whole system computation
- Includes support for compiling conforming parallel programming standards into AMD's native ISA
- Can perform source-to-source transformations of CUDA code to a portable C++ via HIP

Our hardware selection, with the compute power provided by the Xeon processors and 256GB total RAM, the high speed interconnect, and the powerful performance from the two very distinct AMD GPU architectures, would be competitive using any software stack.The Radeon Open Compute Platform, meant to leverage exactly that hardware, further amplifies and enhances that advantage. Our compute platform will conform to the needs of each of the applications due to it's power and versatility of hardware.

Acknowledgements

We would like to acknowledge our faculty advisor Professor David Kaeli, our graduate advisor Xiang Gong, Kurt Keville from MIT, Dr. Anthony Skjellum from Auburn, and Greg Stoner, Guy Ludden, Perhaad Mistry and Yash Ukidave from AMD, for the support they've given us. We also thank AMD, Supermicro and Mellanox for their generous hardware donations, and we thank the employees of those companies who helped set up our system.