

Experimental Heterogeneous Computing with OpenCL

Andrew Kongs
The University of Tulsa
Tulsa, OK
kongs@utulsa.edu

Kyle Cook
The University of Tulsa
Tulsa, OK
kylecook@utulsa.edu

Michael Frohlich
The University of Tulsa
Tulsa, OK
michael-frohlich@utulsa.edu

Peter Hawrylak
The University of Tulsa
Tulsa, OK
peter-hawrylak@utulsa.edu

John Hale
The University of Tulsa
Tulsa, OK
john-hale@utulsa.edu

ABSTRACT

This abstract describes a new kind of Heterogenous computing system where CPUs, MICS and FPGAs work together to contribute each of their strengths to a system. An instance of this concept has been fielded by combining Intel CPUs, Altera FPGAs, Xeon Phi MiC cards, high-speed Interconnects, and fast flash-based storage. This platform is dedicated to HPC security analytics, but also serves as a proving ground for this new breed of heterogeneous compute node clusters.

1. ARCHITECTURE AND TOOLCHAIN

The 12 node heterogeneous compute cluster (named “Hammer”) is built using Intel’s Haswell-E Xeon processors in dual socket configuration with 64GB of RAM (Figure 1). High-speed interconnect is provided by 56GB/s FDR Infiniband to each node. Each node contains a pair of 3100 Series Xeon Phi Co-processors and an Altera Stratix V A7 Accelerator card, the Nallatach 385, with 8GB of RAM per card. In total, the machine provides 192 Xeon CPU cores, 1368 Phi-cores, and approximately 7.4 million FPGA logic elements.




Figure 1: Node Configuration

The key element in the Hammer toolchain is OpenCL. The idealized engagement of the machine is to use OpenCL across Xeon CPUs, Xeon Phi Cards, and the FPGAs. OpenCL

code can be deployed to whichever computing platform is most appropriate. The OpenCL kernels remain unchanged –only the host application would need to be rewritten. The intent is to determine the development efficiency and performance characteristics of this programming paradigm on highly heterogeneous clusters. Moreover, we seek to evaluate options for distributing the work across FPGA cards in the cluster. Currently we are exploring hybrid MPI+OpenCL code to coordinate multiple cards (Figure 2).

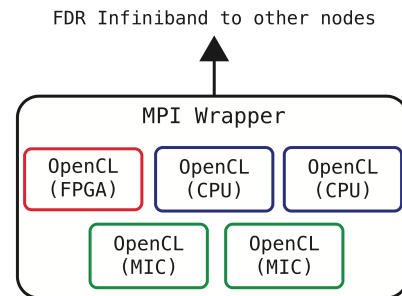


Figure 2: OpenCL Device Communication

2. APPLICATIONS AND FUTURE WORK

Hammer is dedicated to security analytics and related applications. It is being used to explore FPGA solutions for MD5 cracking and parallel generation of hybrid attack graphs. Experimentation with its toolchain continues, as does the expansion of security analytics applications on it. A second cluster, featuring GPUs and FPGAs in each of 16 nodes, is also under construction. GPUs represent an interesting contrast to the MiC co-processing capabilities, potentially exposing additional opportunities or accessible problem spaces in hybridized heterogeneous computing.

ACKNOWLEDGEMENTS

We gratefully acknowledge support from the Army Research Office, DURIP-ARO contract W911NF-15-1-0509. This material is based on work supported by the National Science Foundation under Grant No. 1524940. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.