

Parallel Programmability on Heterogeneous Architectures

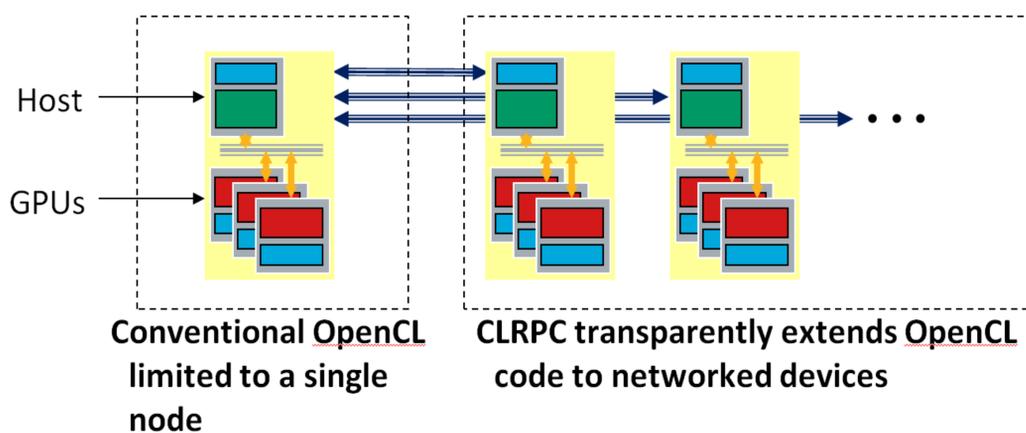


S&T Campaign: Computational Sciences
Advanced Computing Architectures

Song J Park, (410) 278-5444
song.j.park.civ@mail.mil

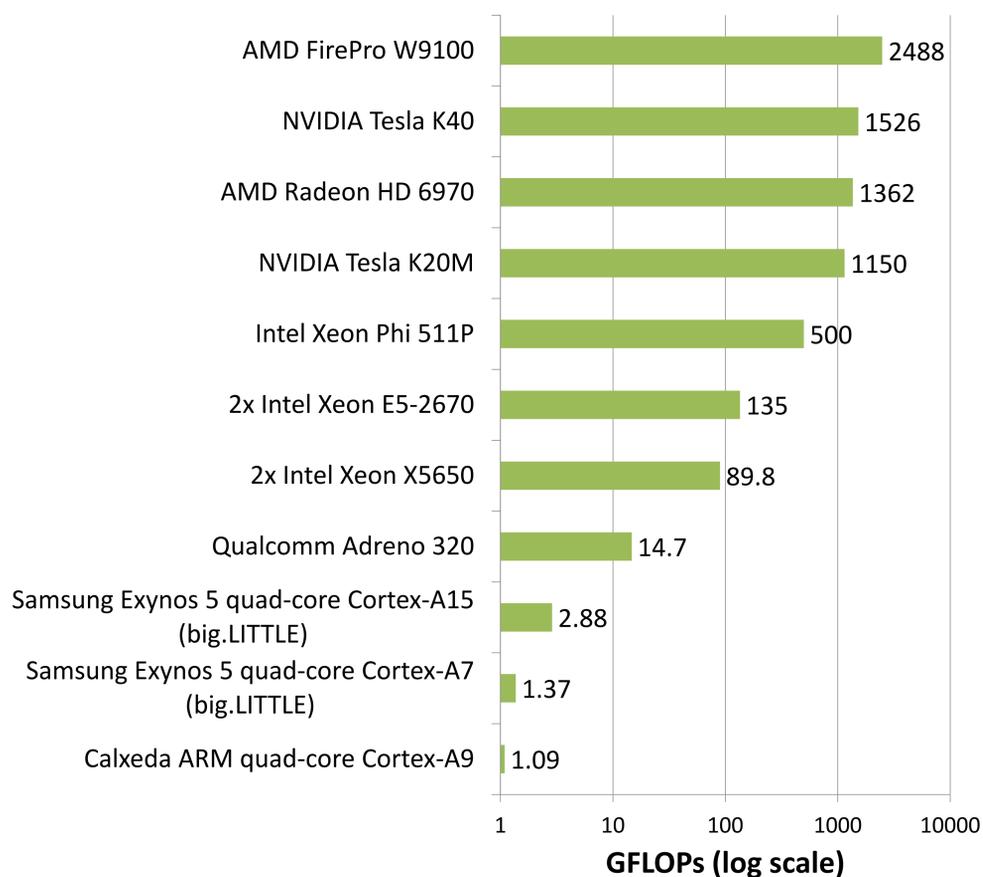
Research Objective

- Explore programming methodologies for the next generation hardware to achieve performance portability in current, emerging, and tomorrow's computational resources.



Challenges

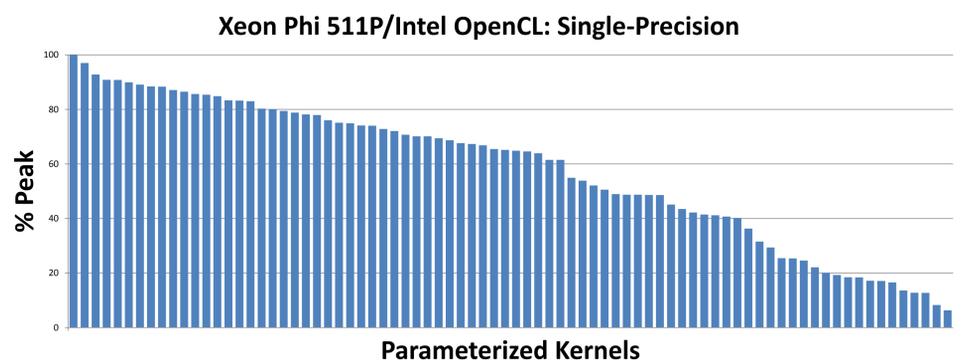
- Gap between hardware capability and software utility.
- Programming model for ubiquitous, diverse, and heterogeneity (x86, ARM, GPU, Phi, Epiphany) parallel architectures.
- Scalability to large number of cores (MPI+X).



N-body performance on wide range of architectures

ARL Facilities and Capabilities Available to Support Collaborative Research

- Available accelerator-based systems
 - 456 Tesla GPU cluster
 - Xeon Phi/Kepler cluster
- Low power platforms
 - Calxeda ARM server
 - Parallella Epiphany processor
 - ARM big.LITTLE development board
- Publications
 - D. A. Richie, J. A. Ross, J. Ruloff, S. J. Park, L. Pollock, and D. R. Shires, "Investigation of Parallel Programmability and Performance of a Calxeda ARM Server Using OpenCL," in *The Sixth Workshop on Unconventional High Performance Computing 2013 (UCHPC 2013)*, 2013.
 - J.A. Ross, D.A. Richie, S.J. Park, D.R. Shires, and L.L. Pollock, "A Case Study of OpenCL on an Android Mobile GPU," *IEEE High Performance Extreme Computing (HPEC'14)*, 2014.
 - D. Richie and J. Ross, "Cycle-Accurate 8080 Emulation on Using an ARM11 Processor with Dynamic Binary Translation," *ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE)*, 2014.



- Performance results show that newer architectures are sensitive to user programming and implementation

Complementary Expertise/ Facilities/ Capabilities Sought in Collaboration

- Auto-tuning techniques to address performance portability where optimal performance is tied to utilizing specific hardware features.
- Expertise in programming model for heterogeneous parallel computing.
- Architecture research for data intensive sciences.