Vancouver: Designing a Next-Generation Software Infrastructure for Productive Heterogeneous Exascale Computing

Jeffrey S. Vetter (PI), Oak Ridge National Laboratory
Wen-mei Hwu, University of Illinois at Urbana-Champaign
Allen D. Malony, University of Oregon
Rich Vuduc, Georgia Institute of Technology

Scalable Heterogeneous Computing (SHC) systems present several challenges: low programmer productivity, no portability, lack of integrated, standard tools and libraries, and high performance stability sensitivity. However, the most evident challenge is that few programming constructs are able to span the range from the fine-grain parallelism supported by these heterogeneous computing devices to the large-scale parallelism required for the Exascale. For example, OpenCL can address the former, and the Message Passing Interface (MPI) can address the latter, but there are few languages or software tools that address both levels simultaneously. These challenges also hold for performance tools, debuggers, resource managers, and libraries. Taken together, these issues will impede the adoption of SHC architectures by erecting a very high entry barrier to application teams and their scientific productivity.

The Vancouver team proposes to address these challenges to performance and productivity with a three-tiered approach for a next-generation software infrastructure designed for Exascale computing:

High-level systems and abstractions
We propose new programming models that support global address spaces and multiple levels of parallelism.

Programming, development, and performance tools
We propose to develop static and dynamic tools for code analysis, inspection, and transformation with integrated performance analysis and prediction.

Low-level libraries and runtime systems
We propose development of libraries and infrastructure for benchmarking, auto-tuning, data movement, and task scheduling.

We believe that, combined, these enhancements will greatly improve the productivity, and thus viability, of heterogeneous systems for Exascale high performance computing. More concretely, we will test, validate, and apply the infrastructure and tools we develop on systems and applications of interest to DOE, and we will deliver them as an open-source software stack. In fact, many of our team's existing tools are already available on DOE systems. Furthermore, we will build upon our proven track record of working closely with application teams to port and optimize their applications onto SHC architectures. We believe that these close collaborations will be absolutely necessary in the early stages of the community's transition to SHC systems, and we propose an aggressive outreach plan to help facilitate the success of this transition.

OAK RIDGE National Laboratory
ILLINOIS UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
UNIVERSITY OF OREGON
Georgia Tech College of Computing Computational Science and Engineering
U.S. DEPARTMENT OF ENERGY Office of Science

Our research agenda is categorized into seven inter-related areas:

Language tools that use static analysis and transformations to generate efficient heterogeneous executables.

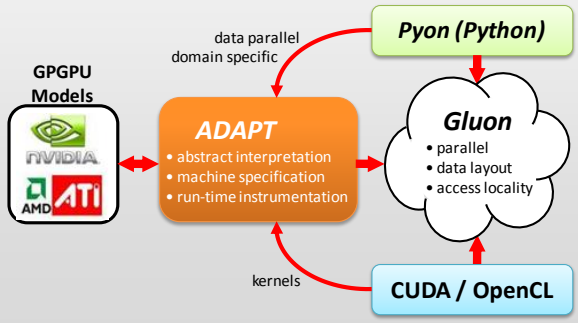

Partitioned Global Address Space programming models that facilitate easy programming of many nodes by providing a global address view of heterogeneous memory systems.

- Testing Pyon/Gluon in context of multiple Global Arrays parallel operations to investigate ability to generate and optimize Global Array API accesses
- Developing and implementing a set of interfaces for spawning asynchronous GPU computation that will work with any PGAS language
- Uses OpenCL and CUDA
- Integrates with Maestro and CUDA to support asynchronous execution

Autotuning irregular algorithms.
Developing model-guided techniques for heterogeneous systems, using Vancouver tools

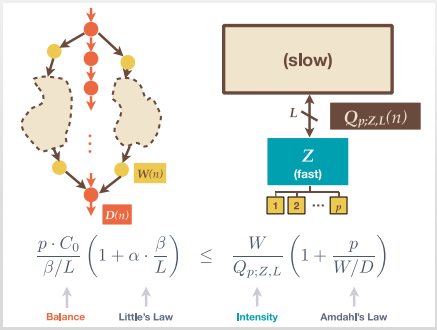

Runtime systems that orchestrate data movement with little or no input from the application.

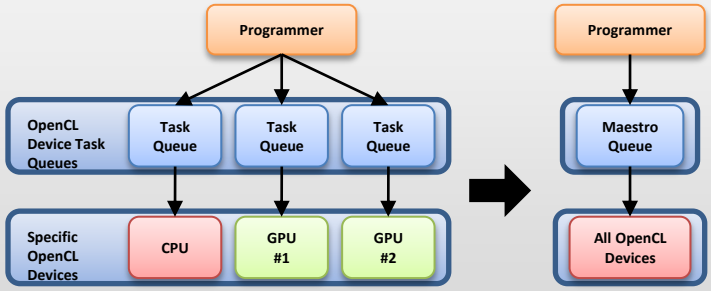

Maestro unifies the OpenCL task queues for several devices into a single high level queue that is independent of the underlying hardware

Performance tools that provide an integrated view of all the performance information used in a scalable heterogeneous application.

- Developing and evaluating TAU support for OpenCL similar to existing TAU CUDA support
- Building support for code instrumentation and transformation of CUDA and OpenCL code, including intermediate PTX code
- Building tools support for Maestro and the Vancouver PGAS prototype
- Building performance analysis, performance database, and data mining support

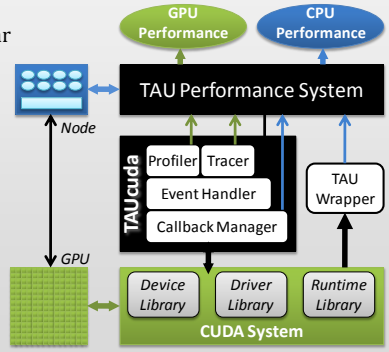

Benchmarks that provide quantitative guidance about costs of key kernels and data movement operations.

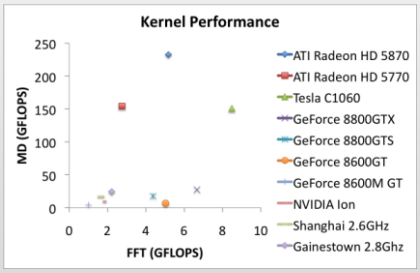

Scalable HeterOgeneous Computing (SHOC) benchmark suite Molecular Dynamics results

Performance prediction tools that account for different instruction set architectures (ISAs) and data orchestration costs

For more information, contact:
Jeffrey S. Vetter, vetter@ornl.gov, 865-356-1649
Wen-mei Hwu, hwu@crhc.uiuc.edu, 217-244-8270
Allen D. Malony, malony@cs.uoregon.edu, 541-346-4407
Rich Vuduc, richie@cc.gatech.edu, 404-385-3355