RAPIDS

Resource and **A**pplication **P**roductivity through computation, **I**nformation, and **D**ata **S**cience

Application Engagement & Community Outreach
Tiger Teams
Liaisons | *Outreach*

Data Understanding | Platform Readiness | **Scientific Data Management**

Platform Readiness

Jeffrey Vetter, Pul Hovland, Samuel Williams, Costin Iancu, Seyong Lee, Swann Perarnau, Philip C. Roth, Kevin Huck

Motivation

ASCR Computing Systems are growing more **complex** with
- heterogeneous compute and
- deep memory hierarchies, and
- ambiguous, non-portable programming models and libraries.

Summit

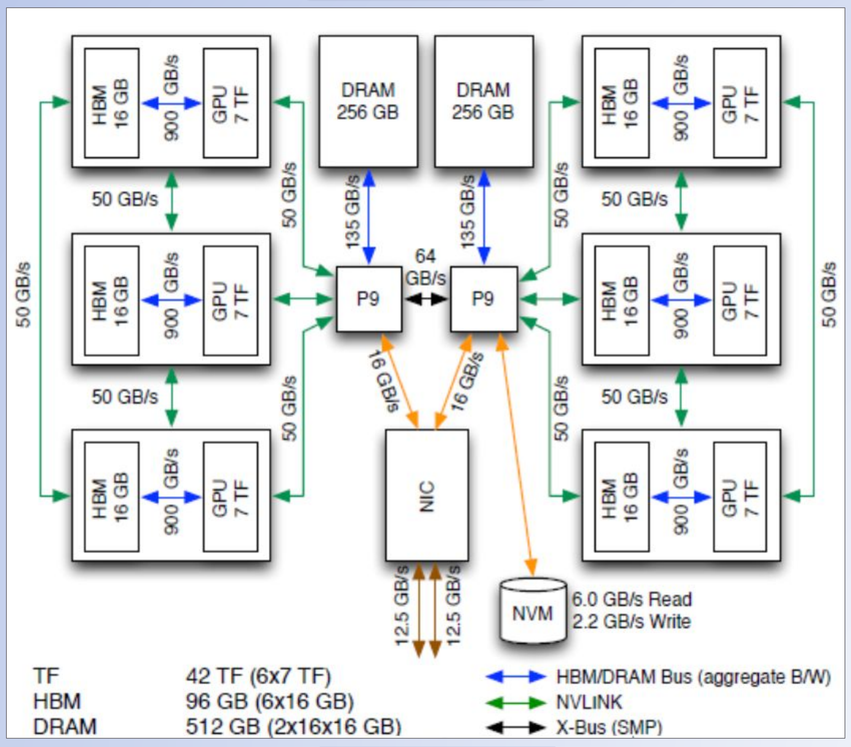

Cori

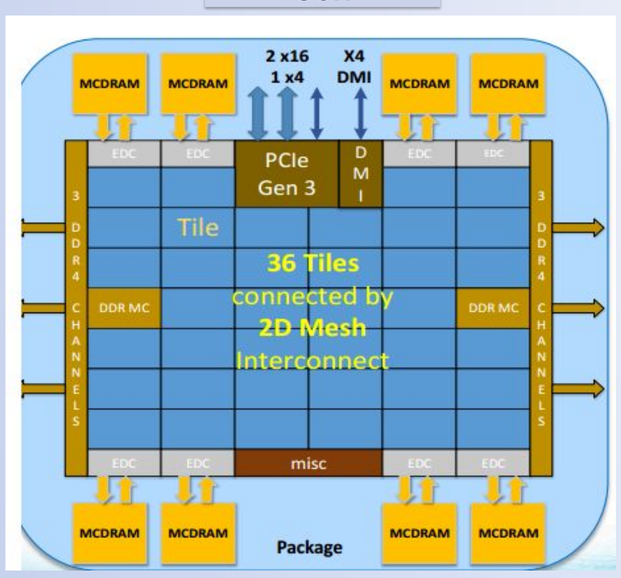

Focus Areas:

- Preparing applications for current and upcoming architectures through best-in-class expertise and tools.
- Portable programming for heterogeneous and many-core systems, deep memory hierarchies
- Code generation and autotuning for computation and communication
- Performance modeling and analysis for identifying optimization opportunities
- Correctness of programs (e.g., when moving to new platforms)
- Tools: CHiLL, CIVL, various compilers (ROSE, OpenARC, LLVM), Roofline toolkit, Orio, Papyrus, SCR, TAU

Software Products

OpenACC+OpenMP / OpenARC

- Open-sourced, OpenACC/OpenMP compiler supporting NVIDIA/AMD GPUs, Intel Xeon Phis, and Altera FPGAs.
- https://ft.ornl.gov/research/openarc
- Contact: Seyong Lee, lees2@ornl.gov

OpenACC on LLVM / Clacc

- CLACC: OpenACC support for Clang/LLVM
- Automated translation from OpenACC source to OpenMP source.
- https://ft.ornl.gov/research/clacc
- Contact: Seyong Lee, lees2@ornl.gov

Deep Memory Hierarchies

- Papyrus: a novel programming system for aggregate distributed NVMs.
- https://ft.ornl.gov/research/papyrus
- AML: building blocks for deep-memory aware algorithms.
- https://xgitlab.cels.anl.gov/argo/aml

MPI Message Reordering

- Higher MPI messaging performance
- Performance portable, topology-independent
- Applicable to wide variety of applications.
- Contact: Costin Iancu (cciancu@lbl.gov)

Code Generation and Autotuning

- Compiler optimization of sparse matrix computations
- Autotuning of OpenMP or OpenACC pragmas underway
- https://github.com/CtopCsUtahEdu/chill.git

Roofline Modeling

- Model for identifying performance bottlenecks in applications
- applied to CPUs/KNL/GPUs
- http://crd.lbl.gov/roofline
- Contact: Sam Williams (SWWilliams@lbl.gov)

Program Correctness (CIVL)

- Formal verification of parallel programs to verify equivalence of two implementations or verifying safety properties (e.g., deadlock-free)
- C and Fortran + MPI, OpenMP, and/or CUDA
- http://vsl.cis.udel.edu/civl, siegel@udel.edu

TAU Performance System®

- Profile, trace, and sampling measurement, analysis, and visualization toolkit
- MPI, OpenMP, OpenACC, Python, CUDA, PAPI
- http://tau.uoregon.edu
- Contact: Kevin Huck (khuck@cs.uoregon.edu)

Application Engagement

- Outreach to SciDAC and other applications teams for assessment of their platform readiness challenges.
- Present tutorials and hackathons.
- Coordinate with the AE on tiger teams.

Hierarchical Roofline in Advisor

- RAPIDS/NERSC/Intel developed & analyzed cache-simulator Roofline
- Integrated into Intel Advisor.
- Measures data movement through cache hierarchy & identifies bottlenecks.

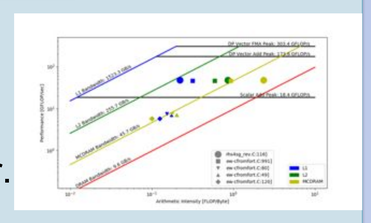

Hierarchical Roofline from Advisor, see ISC'18 paper.

MPI Message Reordering

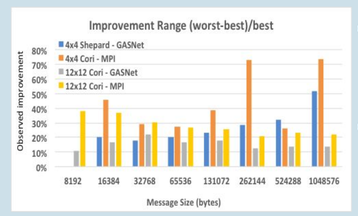

Performance boost from reordering on synthetic stencil benchmarks

- Reordering messages improves performance.
- Applicable to MPI/GASNet, Aries/InfiniBand
- **1.7x on HPGMG, 1.2x on FFTs, 1.2x on Sort**

Roofline on CPU/KNL/GPUs

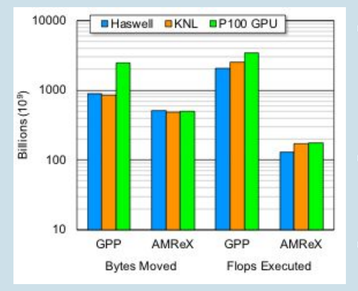

- CRD/NERSC/LANL developed perf. counter Roofline methodology
- Applied to GPP and AMReX proxies.
- Low overhead analysis of MPI+ OMP/CUDA apps.

KORC: Kinetic Orbit Runaway Code

- Performed initial assessment and identified/optimized redundant computations.
- Designed a basic GPU porting strategy using OpenACC/OpenMP.

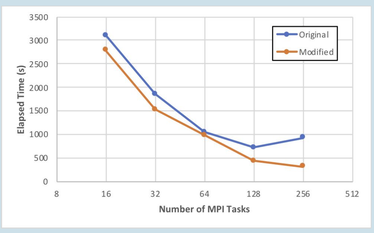

Initial performance improvement of KORC on the OLCF Eos system by common code hoisting optimization

Publications

- K. Ibrahim, S. Williams, L. Oliker, "Roofline Scaling Trajectories: A Method for Parallel Application and Architectural Performance Analysis", HPBench 2018.
- T. Koskela, et al., "A Novel Multi-Level Integrated Roofline Model Approach for Performance Characterization", ISC 2018.
- C. Yang, B. Friesen, T. Kurth, B. Cook, S. Williams, "Toward Automated Application Profiling on Cray Systems", CUG 2018.
- K. Hou, H. Wang, W. Feng, J. Vetter, and S. Lee, "Highly Efficient Compensation-based Parallelism for Wavefront Loops on GPUs", IPDPS 2018.
- W Lavrijsen et al., "Improving Network Throughput with Global Communication Reordering", IPDPS 2018.

Argonne NATIONAL LABORATORY | BROOKHAVEN NATIONAL LABORATORY | BERKELEY LAB | The HDF Group | University of Delaware | Kitware | Lawrence Livermore National Laboratory | Los Alamos NATIONAL LABORATORY | Northwestern University | OAK RIDGE National Laboratory | The Ohio State University | University of Oregon | Rutgers | The University of Utah | U.S. Department of Energy | Office of Science