RAPIDS: The SciDAC Institute for Computer Science and Data

ROBERT ROSS
Institute Director
Argonne National Laboratory
rross@mcs.anl.gov

LENNY OLIKER
Deputy Director
Lawrence Berkeley National Laboratory
loliker@lbl.gov

Diverse Science and Systems

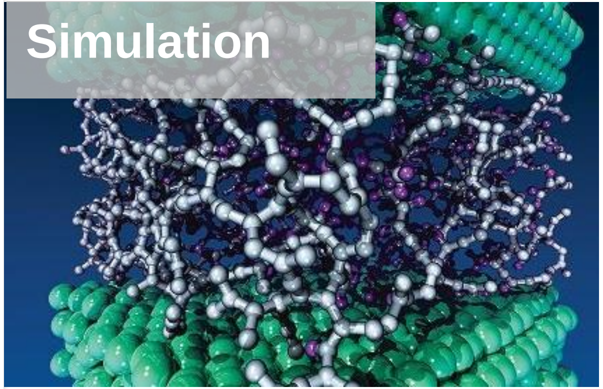

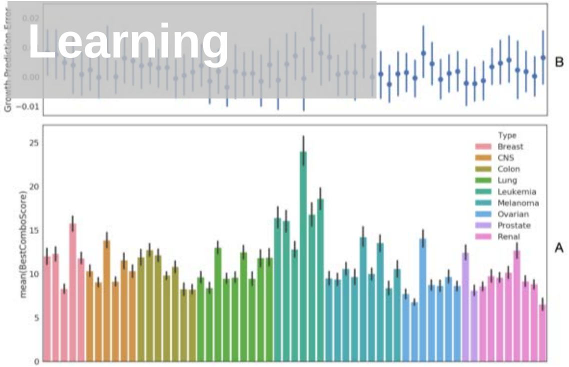

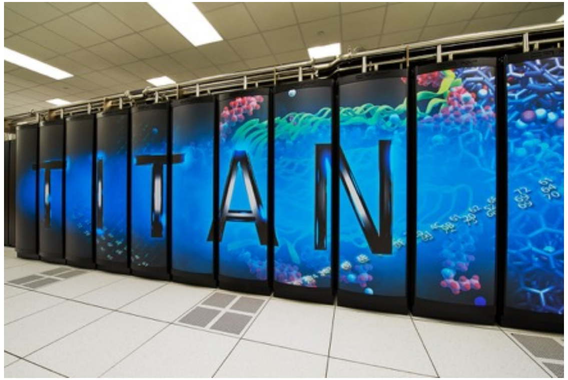

Top image credit B. Helland (ASCR). Bottom left and right images credit ALCF. Bottom center image credit OLCF.

The RAPIDS Institute

Resource and Application Productivity through computation, Information, and Data Science

Objective: Solve computer science and data challenges for Office of Science application teams to achieve science breakthroughs on DOE platforms.

- **Technology Focus Areas**
  - **Data Understanding** – scalable methods, robust infrastructure, machine learning
  - **Scientific Data Management** – I/O libraries, coupling, knowledge management
  - **Platform Readiness** – hybrid programming, deep memory hierarchy, autotuning, correctness
- **Application Engagement**
  - *Tiger Teams* engage experts in multiple technology areas to work with science teams and codes
  - Software productivity: verification and validation, etc.
  - Outreach activities connect with broader community

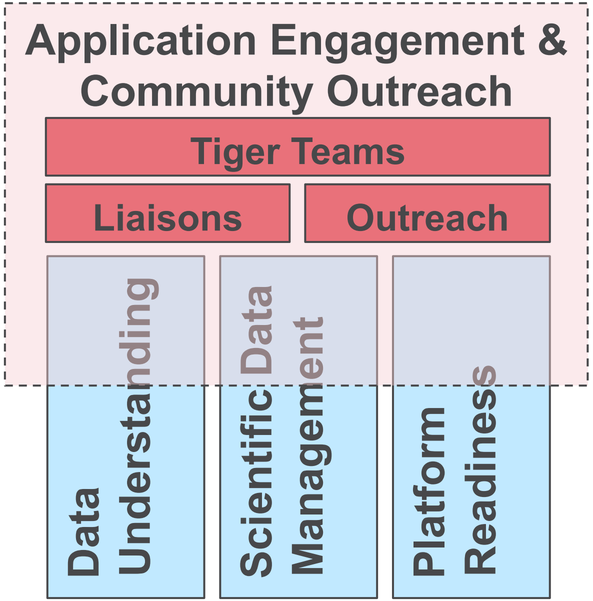

Building on Prior Success

RAPIDS brings together key SciDAC-3 SDAV and SUPER Institute members.

- **Scalable Data Management, Analysis, and Visualization (SDAV) Institute**
  - Data Management – I/O libraries and frameworks, in situ processing and coupling, data compression, indexing
  - Data Analysis – topological & flow field, ensembles, feature driven exploration
  - Data Visualization – analysis frameworks, many-core, distributed viz. and analysis
- **Institute for Sustained Performance, Energy, and Resilience (SUPER)**
  - Performance Engineering – experience with numerous SciDAC-3 SAPs, speedup improvements up to 1000x on SC supercomputers
  - Autotuning – advanced capabilities and tools to optimize SciDAC applications
  - Performance Tools – tuning and analysis tools, Roofline modeling & visualization
- **FASTMath Collaboration** – performance improvements in solvers and applications
- **Machine Leaning and Deep Learning Thrust** – Advance existing expertise to develop domain-specific adaptations targeted for SciDAC application data analysis

Data Understanding

DMITRIY MOROZOV
LBNL

PRASANNA BALAPRAKASH
ANL

Data Understanding

Facilitate understanding of large and complex science data through robust and scalable analysis methods, including learning approaches.

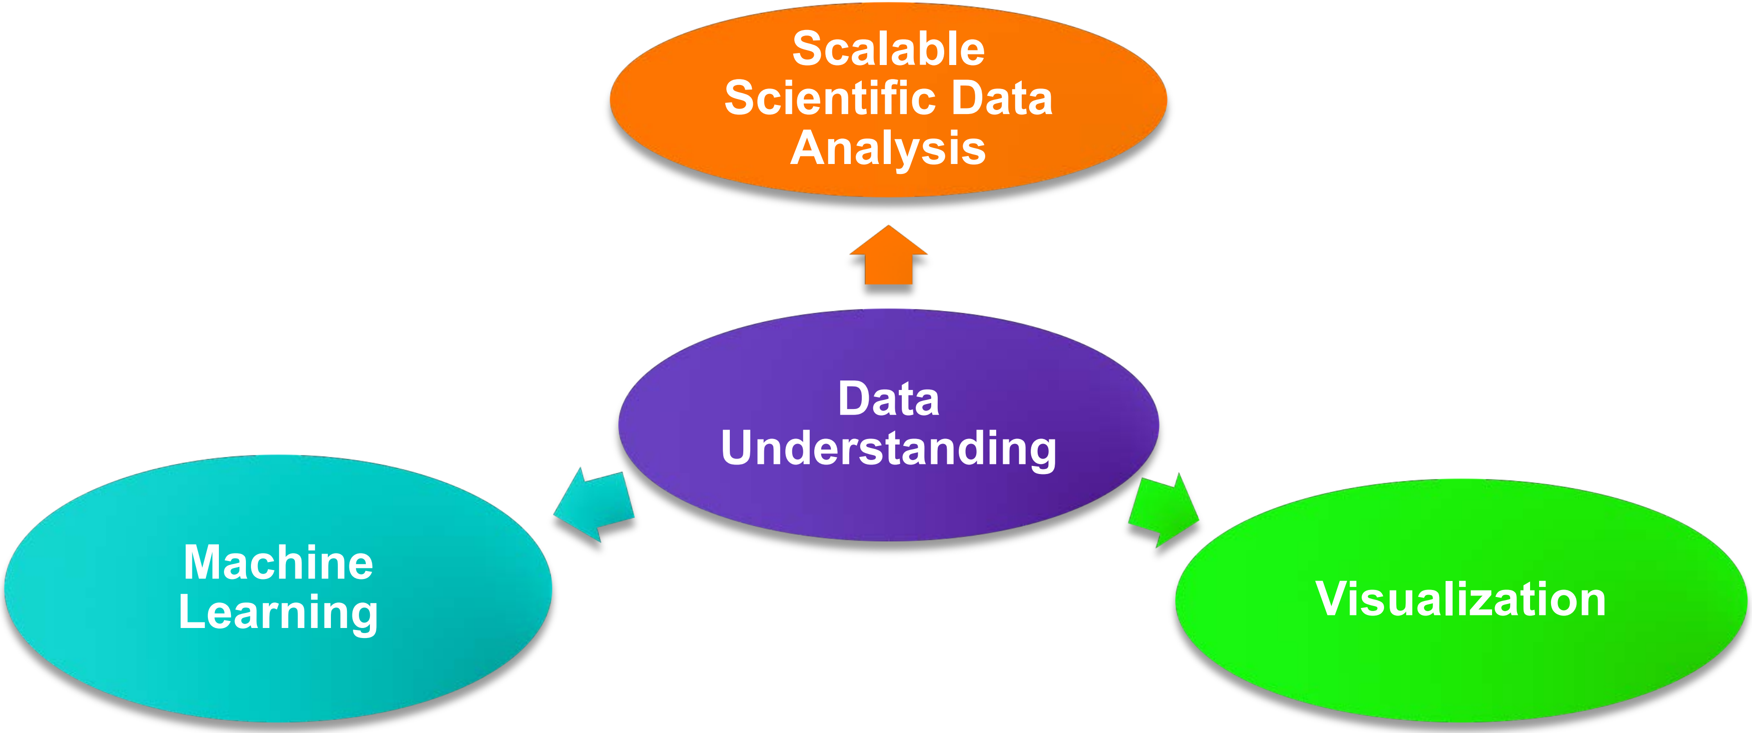

Visualization

RAPIDS

Feature-rich visualization tools that can be run at scale, in situ

- Successful existing tools: **ParaView** and **VisIt,** both built on top of **VTK**, take advantage of massively parallel architectures of modern super-computers
- In situ frameworks, **VisIt/libsim, ParaView/Catalyst, ADIOS, Sensei**, enable using these systems efficiently with the simulations, e.g., to visualize live simulations avoiding the IO bottleneck
- **Scalable infrastructure**: service-oriented data analysis and reduction, co-analysis with performance data
- Major focus on adapting to the deep memory hierarchies and massive on-node hybrid parallelism **(VTK-m)**
- Also useful information visualization techniques **(EDEN)**, techniques for analysis and visualization of high-dimensional datasets

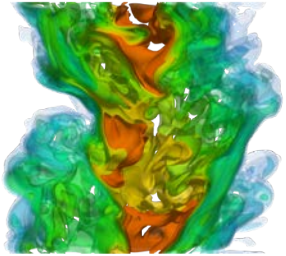

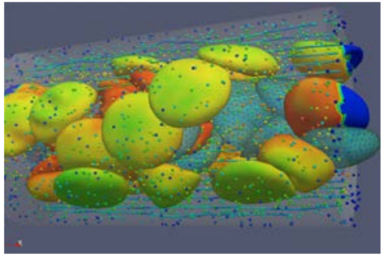

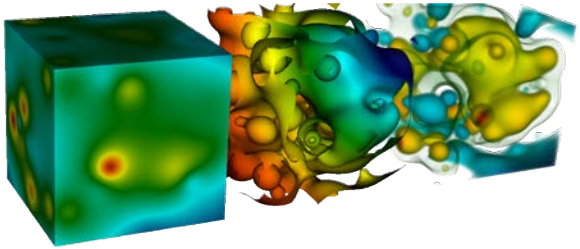

VisIt VTK ParaView VTK-m

Scientific Data Analysis

Scalable methods for finding and analyzing features of importance

- Expertise in **feature detection**, traditionally for visualization and comparative analysis. Moving forward as input to machine learning methods.
- **Geometric analysis (tess):** scalable computation of Delaunay and Voronoi tessellations, e.g., for density estimation in cosmological data
- Statistical analysis of ensemble data **(edda)**:
  - representation of large scale uncertain data
  - analysis of ensemble and uncertain features
  - exploration of parameter space for ensemble simulations
- **Uncertain flows** from ensemble modeling (fluid dynamics, climate, weather)
  - Generalizing flow features for uncertain data
  - Surface Density Estimates to quantify uncertainty
  - Scalable algorithms to stochastically trace particles
- **Topological features** in scalar fields
  - Scalable computation of merge trees, contour trees, persistence diagrams (used in cosmology, combustion, materials science, etc.)
  - Useful both for visualization and for comparison of simulations, to each other and to experiments

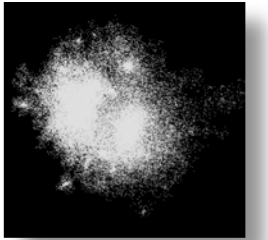

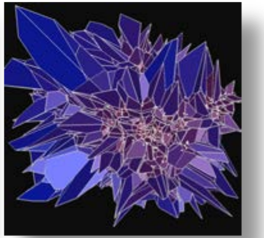

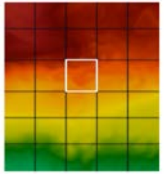

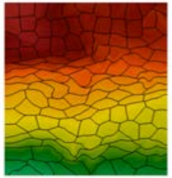

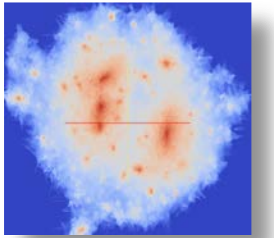

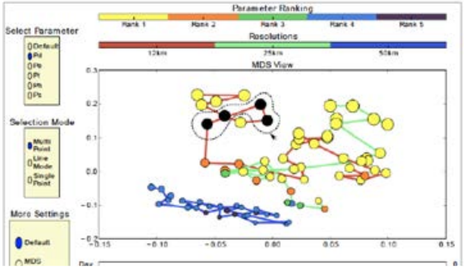

Machine Learning

Domain-specific applications of deep learning, predictive performance models, data- and model-parallel training

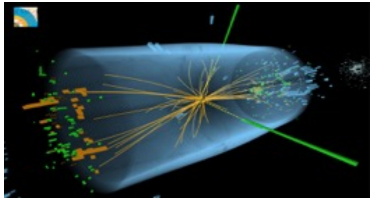

- Supervised learning methods:
  - Deep learning for object classification and identification
  - Large-scale training of convolutional NNs
  - Automatic multiobjective modeling **(AutoMOMML)** to simplify model selection
  - Performance, power, and energy modeling of novel HPC architectures; autotuning computational parameters

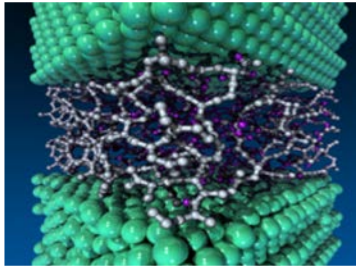

- Unsupervised learning methods:
  - Manifold learning/dimensionality reduction; approximation algorithms to cope with streaming data
  - Streaming spectral clustering
  - Useful for adaptive sampling (e.g., for molecular dynamics trajectories)
- Scalable parallel **graph algorithms (GraphBLAS)**:
  - recast graph algorithms into linear algebra operations
  - communication-avoiding algorithms for key functions
  - building blocks for scalable algorithms
  - apply these techniques to neural networks
- Large-scale sparse inverse covariance matrix estimation **(HP-CONCORD)**

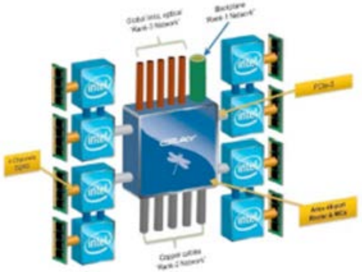

Multivariate, Temporal Visual Analytics for Climate Model Analysis

Scientific Achievement

EDEN enables exploratory data analysis for new DOE E3SM climate simulation and observational data using techniques that combine interactive data visualization and statistical analytics.

Significance and Impact

EDEN gives climate scientists the ability to consider more variables from large scale, land model parameter sensitivity analyses and ultimately improve DOE model accuracy.

Research Details

- Approach helps scientists see that high values of GPP are associated with low leaf carbon to nitrogen ratio values (leafcn) and low critical growing degree days (crit_gdd).
- Based on this insight, climate scientists will generate new ensembles covering smaller ranges of the leafcn and crit_gdd parameter space for more accurate surrogate models.

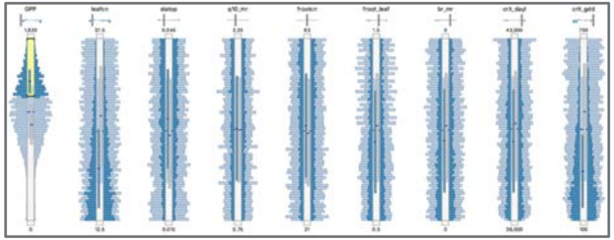

EDEN Parallel Histogram Summary View

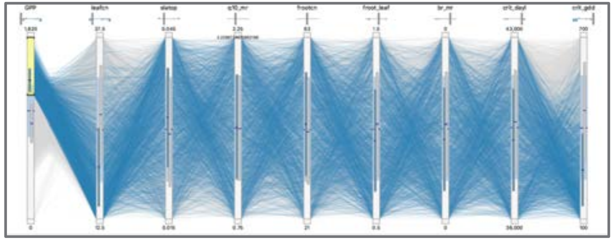

EDEN Parallel Histogram Summary View

The left axis, GPP, is a model output and the other 8 axes are parameters used to run model ensembles (each line is an ensemble run). In both views, the upper range of GPP values are selected revealing an association with low values of crit_gdd and leafcn. New model runs will be executed using more constrained values for these parameters.

Work performed by C. Steed (ORNL) in collaboration with D. Ricciuto (ORNL)

Graphical Model Structure Learning at Unprecedented Scale

Scientific Achievement

Developed HP-CONCORD, a statistically grounded and extremely scalable unsupervised learning method, able to sift through trillions of pair-wise relationships to find the most prominent ones.

Significance and Impact

HP-CONCORD bridges a computational scalability gap between statistically sounds methods and practical usability for some of the largest modern datasets.

Research Details

- Efficient parallel scaling via novel linear algebra communication-avoiding algorithm.
- Allows solving problems orders of magnitude larger and faster than previous best algorithm.
- Proof of concept via brain connectivity analysis, next step is focusing on applying this approach to important SciDAC4 data analysis challenges.

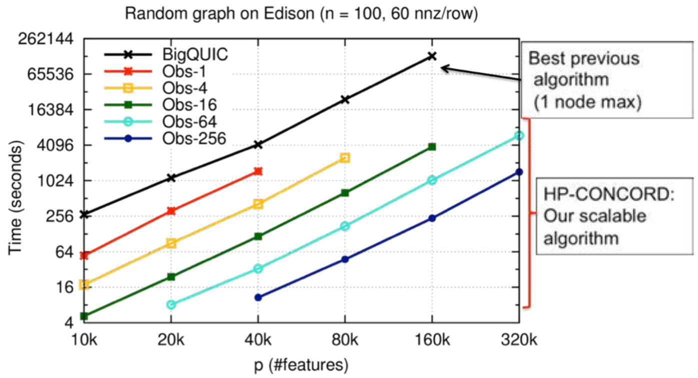

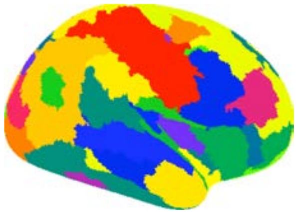

Data-driven right hemisphere parcellation.

(TOP) Extreme scalability of HP-CONCORD using problems up to 1.28 million dimensions (> 800 billion parameters), showing that it can outperform the best previous method by orders of magnitude.
(BOTTOM) Data-driven method demonstrated recovery of underlying connectivity from a resting state fMRI dataset (~91K grayordinates). Results show good agreement with a state-of-the-art clustering from neuroscience literature.

P. Koanantakool, et al., *AISTATS* (2018).

Accelerating Weather Research Forecasting Simulations with Deep Neural Network Surrogates

Scientific Achievement

Developed deep neural network model to accelerate the physical processes simulation in weather research forecasting simulation on leadership-class machines

Significance and Impact

A factor of 10 reduction in the model simulation time would increase the ensemble size into hundreds of individual realizations as opposed to the tens that are feasible now.

Research Details

- Training Dataset: 10 years of model output saved every three hours.
- 1-D Inputs: Surface properties, fluxes, ground temperature; Target: Discover a model for calculating profiles of wind, temperature, moisture given initial conditions and 1-D inputs
- Developed **physics-based deep learning** model

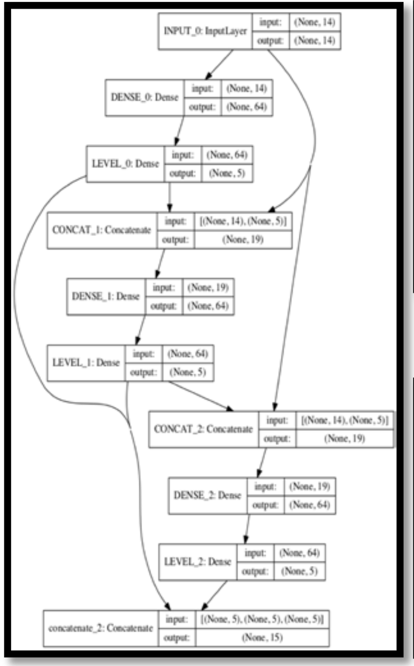

(level 1: close to surface)

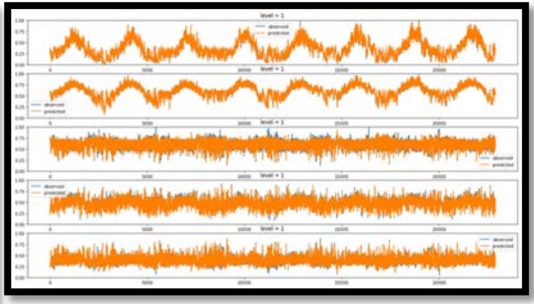

(level 16: ~2 km)

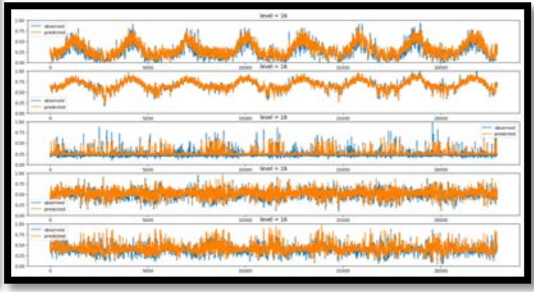

Jiali Wang, Prasanna Balaprakash, Rao Kotamarthi
Deep Neural Network Emulation of a Planetary Boundary Layer Parameterization in Weather Research Forecasting model.

U.S. DEPARTMENT OF ENERGY | Office of Science

RAPIDS

Data Management

SCOTT KLASKY
ORNL

JOHN WU
LBNL

Data Management

Deploy and support efficient methods to move and manage data in a scientific campaign.

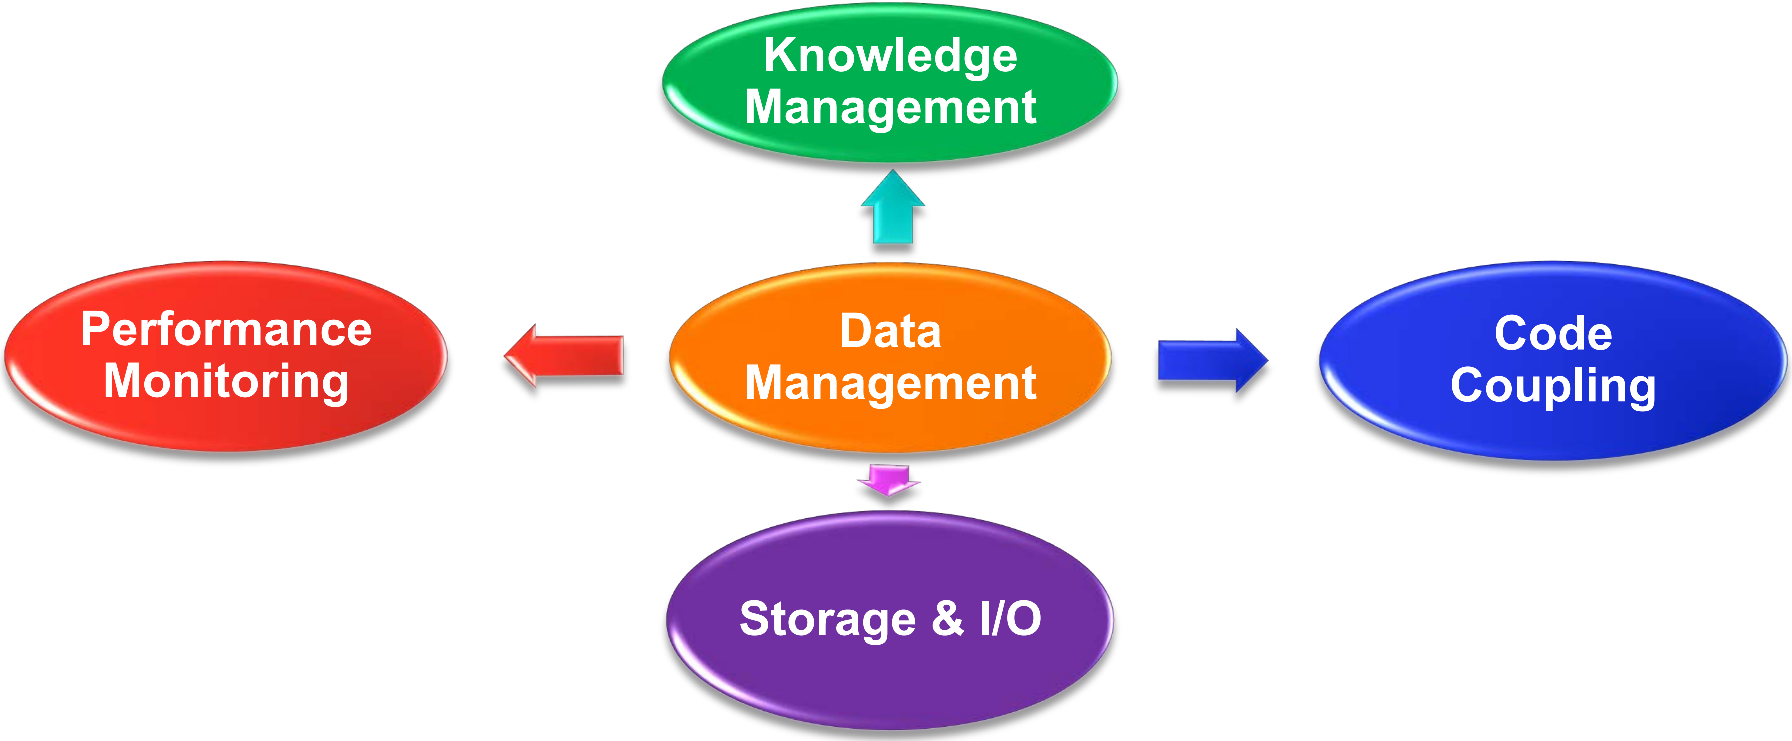

Performance Monitoring

Enabling understanding of I/O performance at scale

- Darshan
  - **“Always on” statistics gathering**
  - Observes I/O patterns of applications running on production HPC platforms, without perturbing execution, with enough detail to gain insight and aid in performance debugging

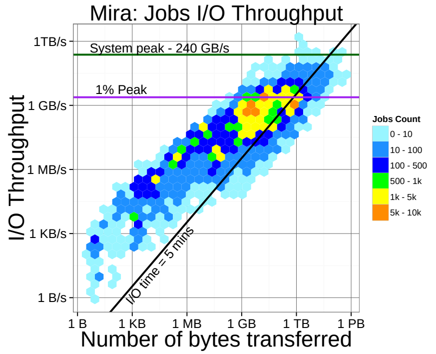

- TAU
  - **Fine-grained tracing of I/O operations at multiple layers**
  - SOSFlow: (Scalable Observation System for Scientific Workflows) provides a flexible, scalable, and programmable framework for observation, introspection, feedback, and control of HPC applications

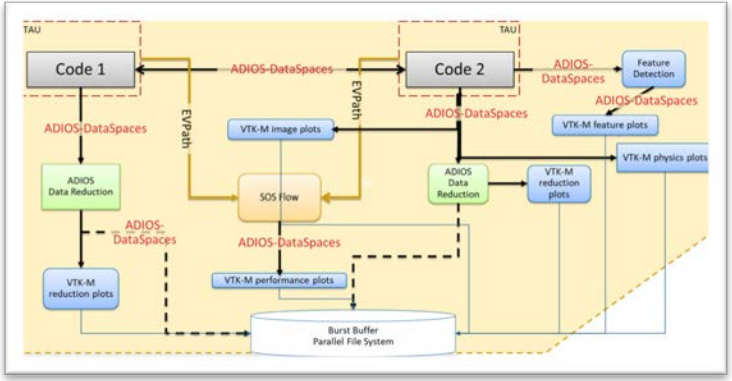

Storage and I/O

Libraries/frameworks to assist in fast and portable I/O.

- **HDF5**
  - A data model, parallel I/O library, and file format for storing and managing data
  - Flexible, self-describing, portable, high performance
- **Parallel netCDF**
  - Provides parallel access to traditional netCDF datasets
  - Includes algorithms for accelerating common patterns such as multi-variable writes
- **ADIOS**
  - A community I/O framework to enable scientific discovery
  - **Online communication** or for persistence to the storage layers
  - Incorporates the state of the art I/O techniques for checkpoint, self describing data, and in situ data movement between codes

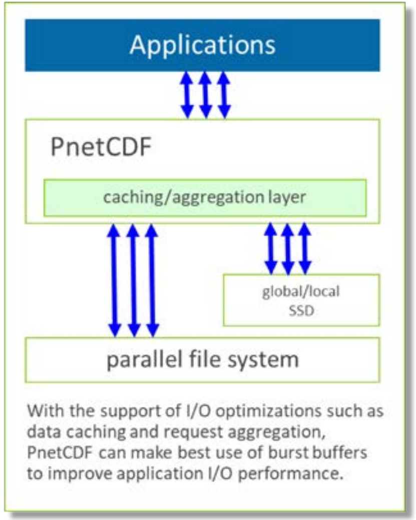

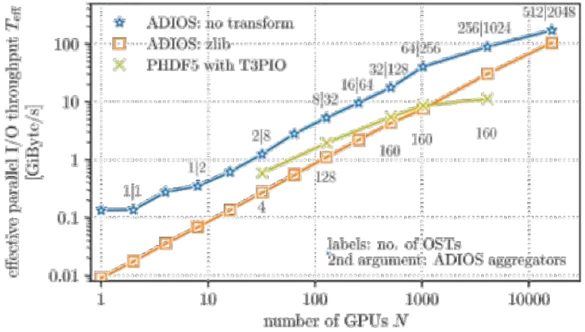

Knowledge Management: FastBit+

Organize and quickly find records across files generated and used during a scientific campaign

- Keeping relationships for “cause-effect” from provenance
- Manage and query the data across a scientific campaign
- Maintain the relationships of input and output from source code, workflows, images, input data, and output data
- Initial implementation will work with data from ADIOS-BP, HDF5, images, and source code

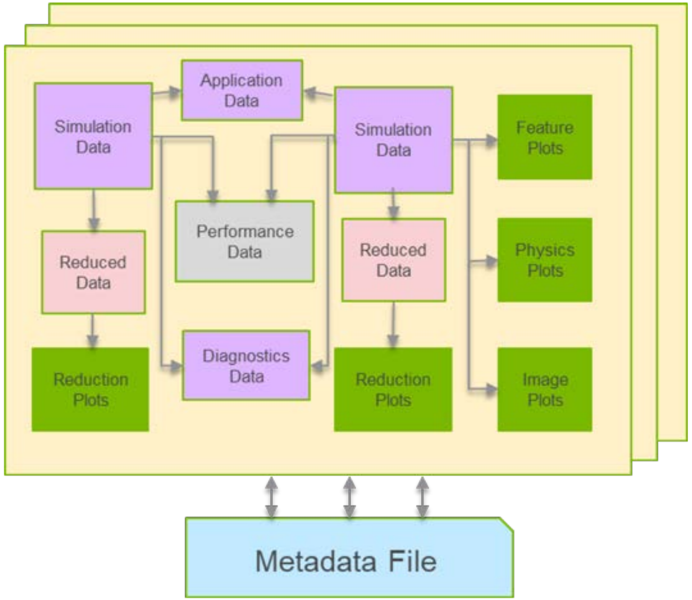

Code Coupling: DataSpaces

In-memory storage distributed across set of cores/nodes, using RAM and/or NVRAM

- Fast I/O to asynchronously couple codes together
- Couple simulation, visualization, analysis, and performance monitoring
- In-staging data processing, querying, sharing, and exchange
  - Virtual shared-space programming abstraction
  - Provides an efficient, high-throughput/low-latency asynchronous data transport

Note: Newest DataSpaces is called SST2.

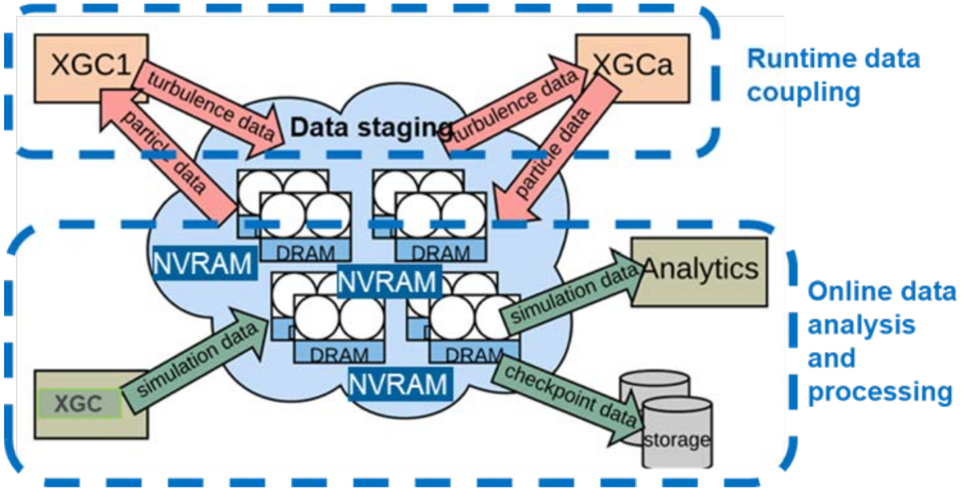

Enabling Global Adjoint Tomography at Scale through Next-generation I/O

Scientific Achievement

Most detailed 3-D model of Earth's interior showing the entire globe from the surface to the core–mantle boundary, a depth of 1,800 miles.

Significance and Impact

First global seismic model where no approximations were used to simulate how seismic waves travel through the Earth. The data sizes required for processing are challenging even for leadership computer facilities.

Research Details

- To improve data movement and flexibility, the Adaptable Seismic Data Format (ASDF) was developed that leverages the Adaptable I/O System (ADIOS) parallel library.
- It allows for recording, reproducing, and analyzing data on large-scale supercomputers
- 1PB of data is produced in a single workflow step, which is fully processed later in another step.

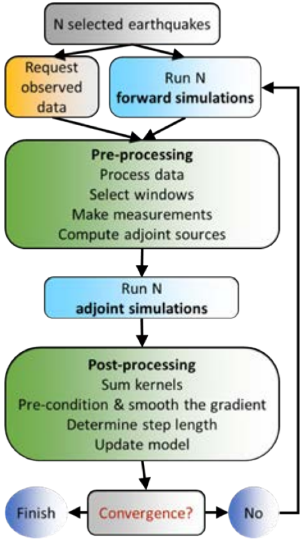

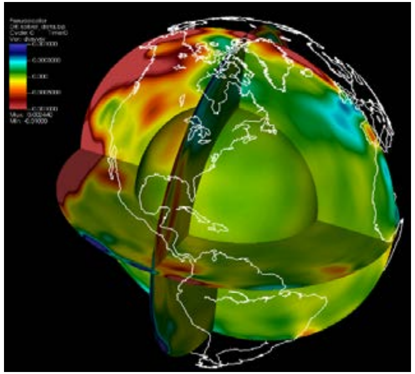

(Left)** Seismic Tomography workflow graph. The heavy computational steps are the Forward and Adjoint Simulations steps. They produce and consume the large data sets, respectively.* ***(Right) A visualization of the Earth's interior with unprecedented details from the seismic tomography process model, which maps the speeds of waves generated after earthquakes. *(Image Credit Dave Pugmire)*

Ebru Bozdag et al. **Global adjoint tomography: first-generation model.** *Geophysical Journal International* 2016 207 (3): 1739-1766.

Autonomic Data Movement for Data Staging-based In-Situ Workflows

Scientific Achievement

Enables machine learning guided data staging for extreme-scale in-situ workflows across multiple layers of staging hierarchy.

Significance and Impact

Stacker reduces the number of read requests to disk (SSD) using intelligent prefetching, improving read performance by ~40% as compared to no prefetching.

Research Details

- Use on-node SSDs to extend in-memory data staging solutions across DRAM and SSDs/burst buffers
- Perform autonomic data movement across the multi-tier staging area using n-gram models to learn and predict the next data access request to enable data prefetching
- Learns inter-application data access patterns, enabling support for multiple concurrent workflows

P. Subedi et al. Accepted for presentation at SC'18.

RAPIDS

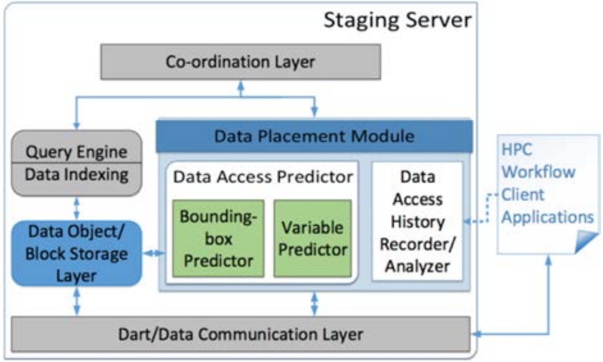

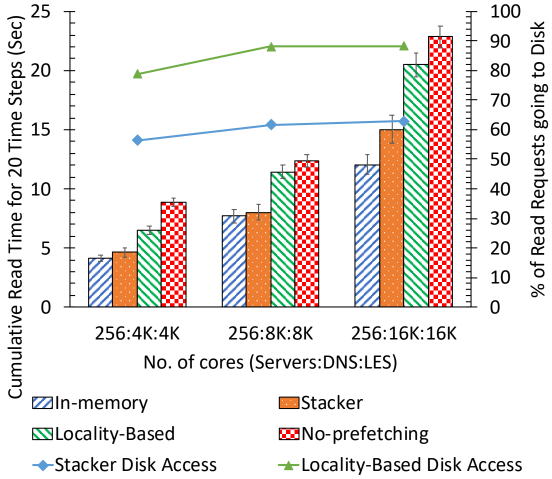

RUTGERS Discovery Informatics Institute

OAK RIDGE National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

RAPIDS

Platform Readiness

JEFFREY VETTER
ORNL

PAUL HOVLAND
ANL

Dealing with Heterogeneity, Deep Memory, and Modern Interconnects

- **Heterogeneity**
  - Many-core, GPU, other accelerators
  - Many languages
- **Deep Memory**
  - Many types (HBM, DDR, PCRAM, ReRAM, …)
  - Many configurations
- **Interconnects**
  - Torus, fat tree, dragonfly(+), …

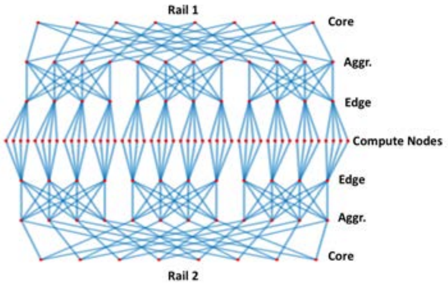

Platform Readiness

Preparing scientific codes for current and upcoming system through application of best-in-class expertise and tools.

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

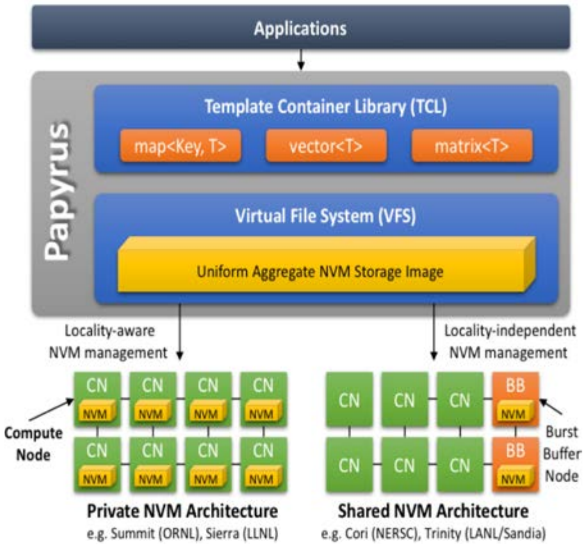

Papyrus provides abstractions for large shared data structures using map, vector, and matrix modalities

Roofline Performance Modeling

RAPIDS

ASCR Base & LDRD

Developed Roofline concept
2006-2011:

- Easy-to-understand, visual performance model
- Offers insights to programmers and architects on improving parallel software and hardware.

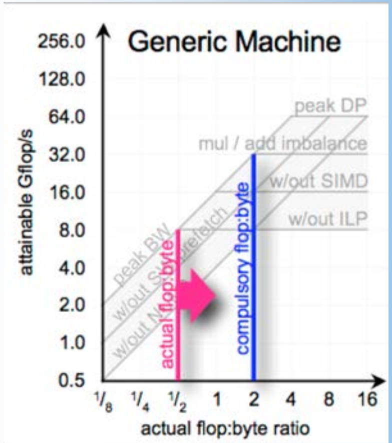

Proof of concept successfully applied to numerous computational kernels and emerging computing systems.

SciDAC3 Development

Roofline augmentation under SciDAC3
2013-2017:

- Collaboration with FASTMath SciDAC Institute
- Developed Empirical Roofline Toolkit (ERT) with public release 03/2015, with Roofline Visualizer
- Created community tool for automatic hardware introspection and analysis

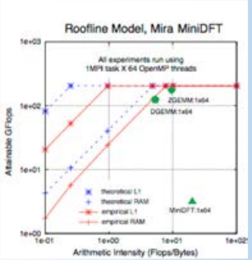

Automated Roofline code used to diagnose performance problems for DOE and SciDAC codes.

Outcome & Impact

- Roofline has become a broadly used performance modeling methodology across DOE
- Intel has embraced the approach and integrated it into its production Intel® Advisor
- **Collaboration with NERSC to instrument and analyze execution of real applications on machines such as Edison and Cori**

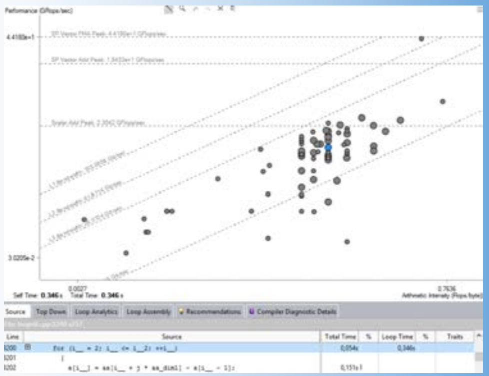

Snapshot of existing Intel Roofline tool in practice.

Performance Observation, Analytics, and Tuning for Heterogeneous Platforms with TAU

RAPIDS

- Heterogeneous software stacks
  - Languages: OpenMP, OpenACC, CUDA
  - Libraries/Metaprogramming: Kokkos, RAJA
  - Hybrid: MPI+X
- Runtimes
  - OpenMP, MPI, I/O, asynchronous multitasking
- Compilers and autotuners
  - LLVM, Chill, Oreo, Active Harmony, OpenARC
- Heterogeneous hardware measurement
  - Memory, Power, Network

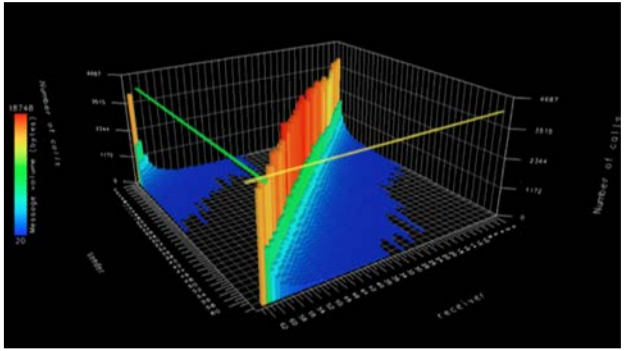

Model-based Autotuning (CHiLL)

- Integrate static (binary and source) analysis into autotuning for GPUs and CPUs
  - Can greatly reduce the empirical search space
  - Can preserve and reuse past autotuning results
- Focus on current and near-future platforms
  - (Portable) vectorization through code generation
  - Memory movement optimization in shared memory environments
    - Especially irregular memory access patterns
  - Effective use of accelerators without maintaining multiple code versions

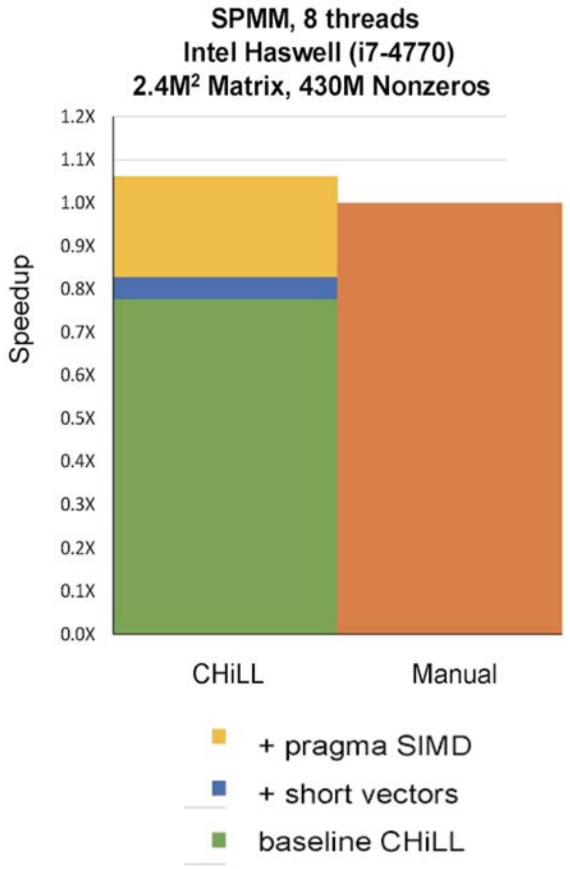

Static Verification of HPC Programs (CIVL)

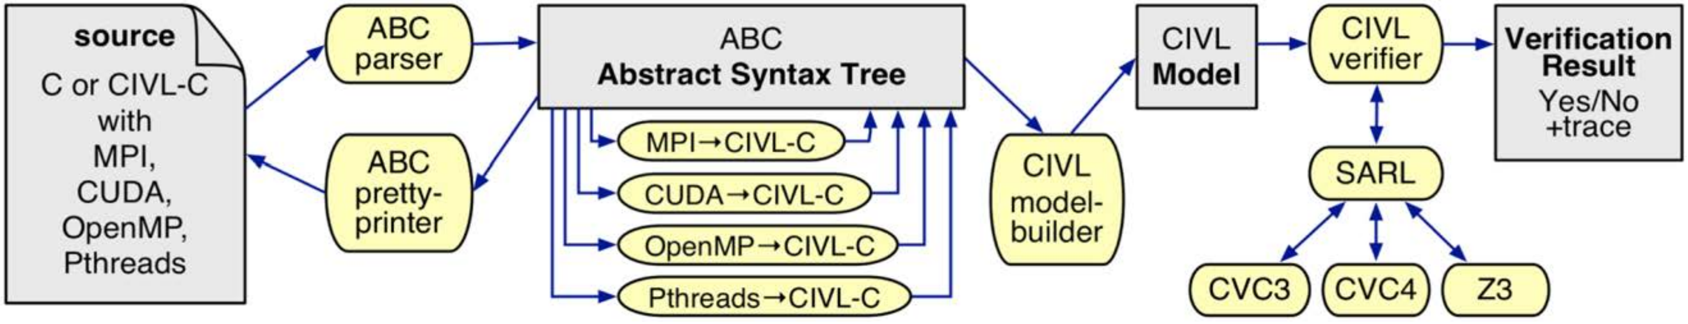

- Source may include CIVL-C primitives: **input**, **output**, **assert**, **assume**
- All concurrency translated to **CIVL-C**
- **Fortran** support in progress; same AST
- Program may be composed of multiple translation units (including Fortran+C)
- Verifier uses **symbolic execution** to check properties for all possible inputs (within specified bounds)
- Absence of: assertion violations, deadlock, illegal pointer operations, out-of-bound indexes, violations of MPI semantics, …
- Verify **equivalence** of 2 implementations

Improving Network Throughput with Global Communication Reordering

Scientific Achievement

Demonstrated up to 70% throughput improvement for production applications using topology-independent, runtime communication reordering.

Significance and Impact

Performance portability, topology-independence, and run-time implementation facilitate use in existing communication libraries for a wide variety of applications.

Research Details

The intuition behind the reordering algorithm is as follows:

1. In order to minimize deadlines, long operations should be issued as early as possible.
2. High variability is a sign of congestion and ordering in ranks that take longer than expected should be perturbed.
3. Communication is likely to be subject to congestion if the load on both endpoints (source and destination) is high.

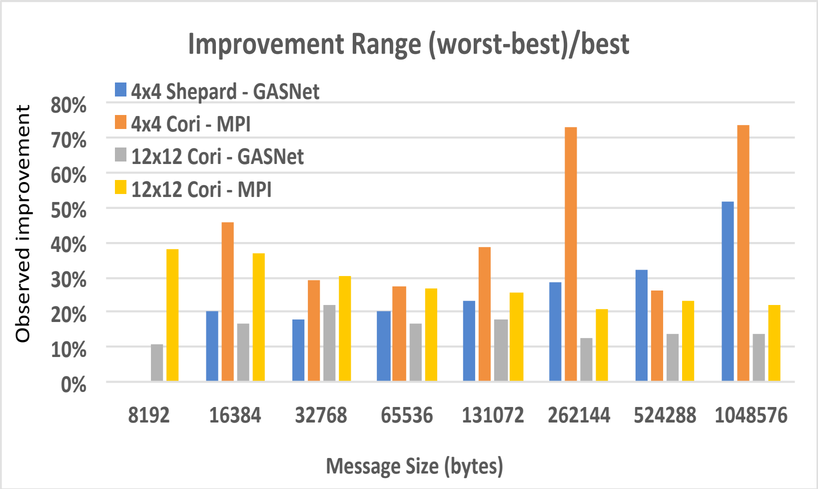

Range of improvement on InfiniBand and Cray Aries, 4x4 and 12x12 stencil; each 2 processes x 16 threads (32 cores) per node. Shepard: up to 1152 cores. Cori: up to 4608 cores.

NERSC procurement benchmarks:
HPGMG: 70% improvement
FFT: 20% improvement
Sorting: 20% improvement

W Lavrijsen et al. Improving Network Throughput with Global Communication Reordering. IPDPS 2018.

Deep Stack Program Optimization

Scientific Achievement

Demonstrated dynamic analyses able to hide communication and synchronization latency in HPC apps by optimizing applications together with their runtime(s).

Significance and Impact

The approach is able to perform optimizations beyond the capabilities of any contemporary HPC compiler and provides a path forward for co-design and specialization techniques for the HPC software stack of the future.

Research Details

- Run program with instrumentation for memory, communication, synchronization e.g. load/store, Put/Get, barrier
- Trace analysis for data dependences and CFG information
- Automatic rewriting of patterns of communication, computation, and synchronization
  - Non-blocking point-to-point (optimal overlap)
  - Non-blocking collectives
  - Bespoke synchronization patterns (e.g. producer-consumer)

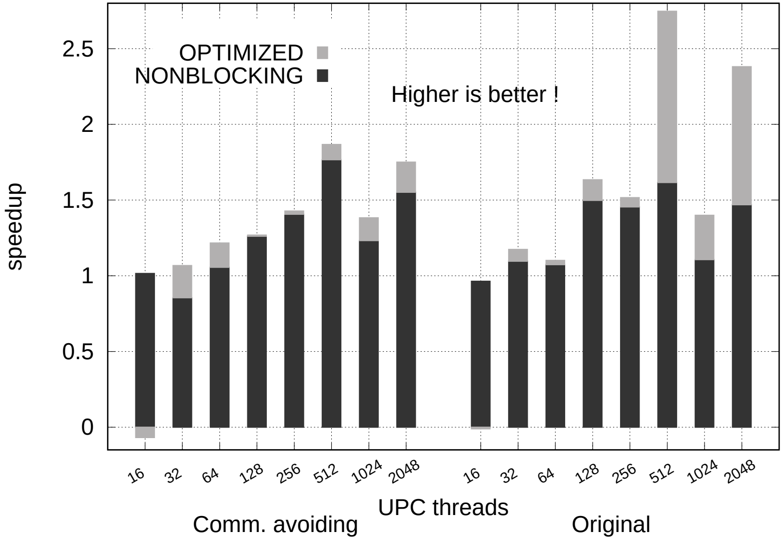

HPGMG performance improvements:

- **Good improvements (64%) when a large number of operations are in flight.**
- **Good improvements (33%) for complex synchronization.**

Maximizing Communication Overlap with Dynamic Program Analysis. E Saillard et al. HPC Asia 2018. Best paper finalist.

(Application) Engagement

ANSHU DUBEY
ANL

SAM WILLIAMS
LBNL

In Many Cases, Already Connected!

RAPIDS

Title	PI	Prog.	RAPIDS Member(s)
Coupling Approaches for Next-Generation Architectures (CANGA)	P. Jones	BER	Peterka, Ross
Prob. Sea-Level Projections from Ice Sheet and Earth System Models	S. Price	BER	Patchett
Advancing Catalysis Modeling	M. H. Gordon	BES	Williams
AToM: Advanced Tokamak Modeling Environment	J. Candy	FES	Bernholdt
Plasma Surface Interactions: Predicting the Performance and Impact of Dynamic PFC Surfaces (PSI-2)	B. Wirth	FES	Bernholdt, Roth, Pugmire
Center for Tokamak Transients Simulations (CTTS)	S. Jardin	FES	Williams
Integrated Simulation of Energetic Particles in Burning Plasmas (ISEP)	Z. Lin	FES	Williams, Klasky, Pugmire
High-fidelity Boundary Plasma Simulation (HBPS)	C. S. Chang	FES	Klasky, Podhorszki
Tokamak Disruption Simulation	X. Tang	FES	Patchett
Inference at Extreme Scale	S. Habib	HEP	Yoo, Morozov, Balaprakash
HEP Data Analytics on HPC	J. Kowalkowski	HEP	Peterka, Ross
HPC Framework for Event Generation at Colliders	S. Hoeche	HEP	Hovland
HEP Event Reconstruction with Cutting Edge Computing Architectures	G. Cerati	HEP	Norris
Simulation of Fission Gas in Uranium Oxide Nuclear Fuel	D. Andersson	NE	Bernholdt, Roth
Towards Exascale Astrophysics of Mergers and SuperNova	W. R. Hix	NP	Dubey
Nuclear Low Energy Initiative (NUCLEI)	J. Carlson	NP	Norris

New Connections

Goal: Help application teams realize highest possible scientific impact.

Applications: New Connections

In Progress

Partnership for Multiscale Gyrokinetic Turbulence
Performance optimization of GENE code on Intel-MKL platform

In Discussion

Development of Terrestrial Dynamical Cores
Software architecture and process design

Prospective

HEP Data Analytics on HPC
Data management

Community Project for Accelerator Science and Simulation 4 (ComPASS4)

Computational Framework for Unbiased Studies of Correlated Electron Systems

Center for Integrated Simulation of Fusion Relevant RF Actuators

FASTMath Collaborations

PETSc and Xolotl

SuperLU and Oreo

Image compression using deep learning techniques

Benefits to applications through FASTMath and RAPIDS Connection

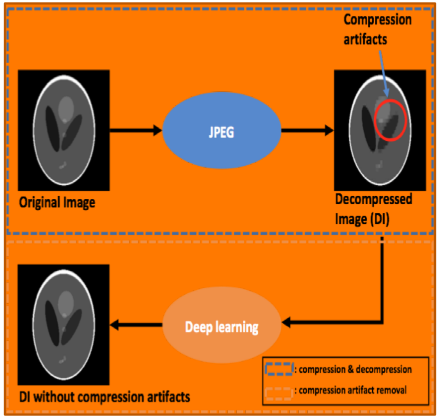

Tiger Teams

RAPIDS

Focused engagements to resolve specific application challenges

- **Engagement with one team for 3-6 months**
  - Needed expertise drawn from all the focus areas as needed
  - Plan ahead for scope and resources
  - Define expected outcomes

- **Two to three Tiger Team activities per year**
  - Prioritized by available expertise and potential impact on the application
  - Resources budgeted in RAPIDS appropriately

If you think we can help, please ask!

Training, Tutorials, and Webinars

Outreach activities

- Allow us to **disseminate** RAPIDS technologies and approaches **to a broader audience**
- Tutorials will be offered at various conferences
  - Also cover software productivity issues
- Other training may be hosted by facilities or institutions, may be online
- Best practices documentation and consultation
- Hackathons and coding camps

Data Understanding Tutorials/Training

- November ‘17 @ SC
- July ‘18 @ PEARC
- August ‘18 @ ATPESC
- October ‘18 @ VIS
- November ‘18 @ SC

Data Management Tutorials/Training

- November ‘17 @ SC
- April ‘18 @ RIKEN
- June ‘18 @ Wuxi
- August ‘18 @ ATPESC
- November ‘18 @ SC

Platform Readiness Tutorials/Training

- November ‘17 @NERSC
- November ‘17 @ SC
- February ‘18 @ ECP
- June ‘18 @ ISC
- November ‘18 @ SC

Not Just *Application* Engagement: Liaisons

RAPIDS

What do they do?

- <u>**Help**</u> **make connections** between RAPIDS activities/people and facilities and other entities
- Promote general awareness of opportunities and needs in RAPIDS and other organizations

Facility/Project	Liaison
ALCF	Kevin Harms, ANL
NERSC	Oliver Ruebel, LBNL
OLCF	David Bernholdt, ORNL
FASTMath Institute	Sam Williams, LBNL
ECP	Jim Ahrens, LANL

Accelerating Fusion/Fission Simulations

RAPIDS

Scientific Achievement

Improved science throughput for several applications used in SciDAC fusion and fission projects

Significance and Impact

Can simulate much larger Xolotl reaction networks than before; established baseline for KORC scientific throughput, and then improved it.

Research Details

- Reduced memory footprint of Xolotl version used for FY2018 FES theory milestone experiments by up to 88% on tested configurations
  – Involved contribution to PETSc (**FASTMath**) code repository
- Developed XDSpace, a Kokkos-based library for multi-resolution decomposition of discrete space
  – Initially targeted to Xolotl reaction network initialization and use
  – Identified optimization enabling Kinetic Orbit Runaway Code (KORC) to run up to ~2x faster

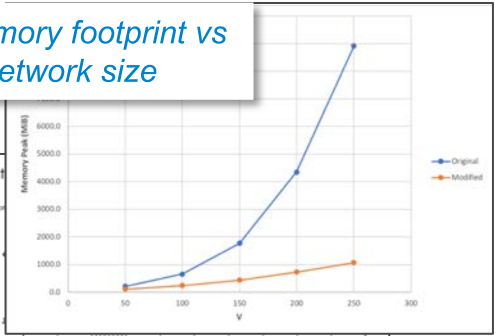

Xolotl memory footprint vs reaction network size

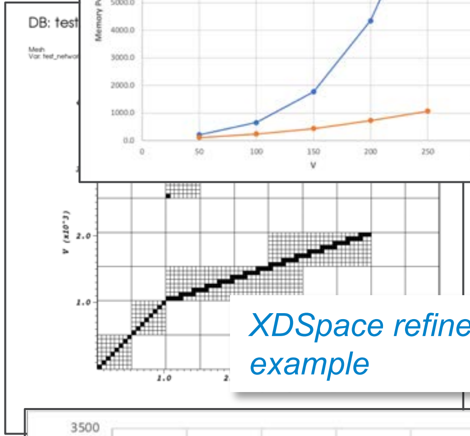

XDSpace refinement example

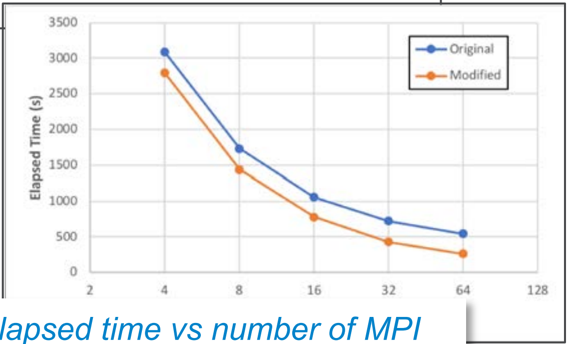

KORC elapsed time vs number of MPI tasks (64 total threads for all runs)

P.C. Roth, with S. Blondel (PSI2, FGS) and L. Carbajal Gomez (SCREAM)

Image Enhancements in Compressed Sensing using Deep Neural Networks

Scientific Achievement

Developed deep learning approaches for fast and accurate image enhancements to enable faster data transfer across networks and reduce storage

Significance and Impact

- **Compression techniques such as JPEG can significantly reduce the size of images, but are prone to blocking artifacts and blurring**
- **Traditional approaches to address these issues are computationally expensive**
- **Deep learning models can reduce the time by an order of magnitude given enough training data**

Research Details

– The model is trained with JPEG compressed image as input and the original uncompressed image as output

– The trained model can be used to enhance new images and compared with the currently used compressed sensing

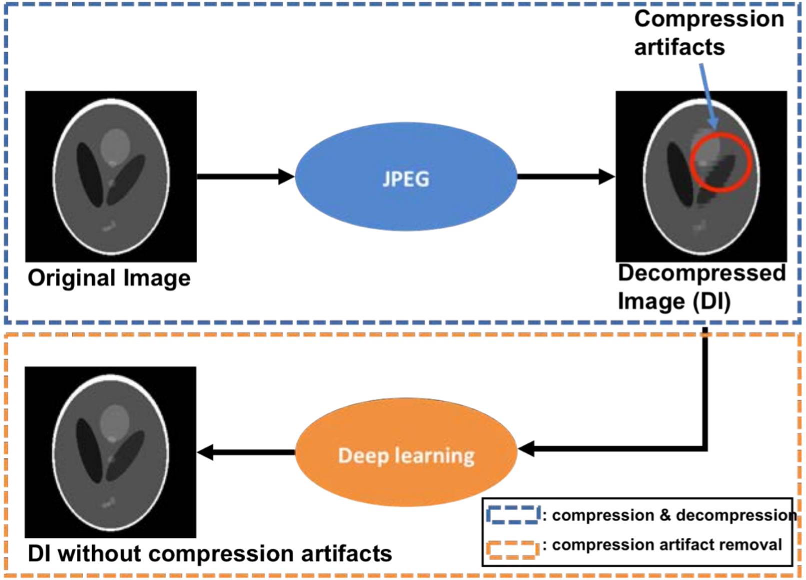

RAPIDS: JH Park, S. Yoo (BNL), S. Madireddy, P. Balaprakash (ANL), W. Liao & S. Lee (Northwestern)

FASTMath: Richard Archibald (ORNL)

Argonne NATIONAL LABORATORY

Machine Learning Based Parallel I/O Predictive Modeling: A Case Study on Lustre File Systems

RAPIDS

Scientific Achievement

We developed a sensitivity-based robust Gaussian process (GP) regression approach that explicitly treats the variability in the data, groups applications with similar characteristics, and automatically reduces the effect of outliers.

Significance and Impact

The robust GP approach provides significant predictive accuracy improvements compared to baseline and high-performing machine learning approaches especially in the presence of outliers, and quantifies the uncertainty in the model prediction.

Research Details

- This approach is demonstrated on I/O performance data obtained on Lustre file system at NERSC, where I/O performance is modeled as a function of application and file system characteristics.
- Detailed insights are drawn into features that are most significant in explaining the variability.

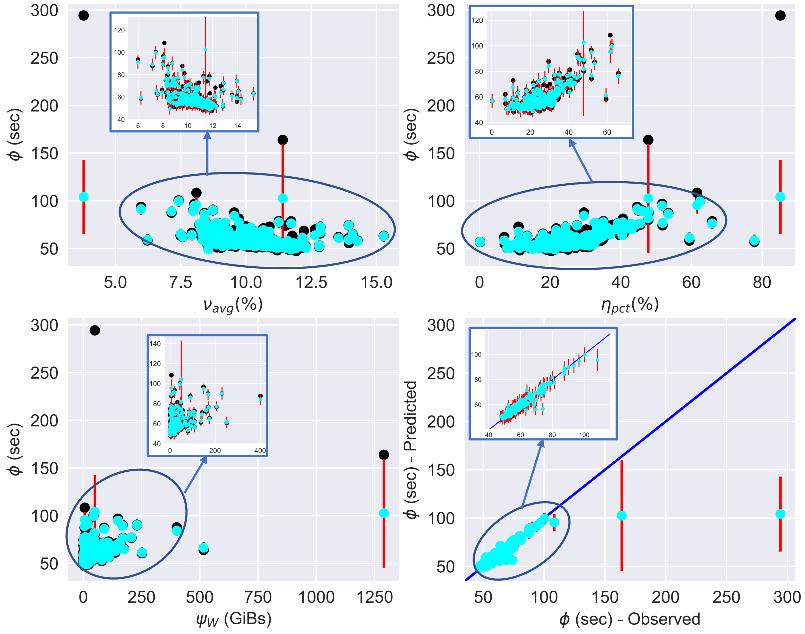

Model prediction for I/O time and its uncertainty (90% credible intervals) shown as a function of a few input features, as well as w.r.t. the experimentally observed I/O time. This approach identifies outliers and ignores them automatically as shown.

S. Madireddy, et al., *ISC-HPC* (2018).

Argonne NATIONAL LABORATORY

Thanks to the RAPIDS Team!

… and more …

RAPIDS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program.

Argonne National Laboratory · Berkeley Lab · University of Delaware · Lawrence Livermore National Laboratory · Oak Ridge National Laboratory · University of Oregon · Rutgers · Brookhaven National Laboratory · The HDF Group · Kitware · Los Alamos National Laboratory · Northwestern University · The Ohio State University · The University of Utah

RAPIDS

For general questions:

Rob Ross <rross@mcs.anl.gov>

Lenny Oliker <LOliker@lbl.gov>

For engagement discussion:

Anshu Dubey <adubey@anl.gov>

Sam Williams <swwilliams@lbl.gov>

On the web:

http://www.rapids-scidac.org

… or just reach out to the RAPIDS person that you already know!

Argonne National Laboratory
Brookhaven National Laboratory
Berkeley Lab
University of Delaware
The HDF Group
Kitware
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Oak Ridge National Laboratory
Northwestern University
The Ohio State University
University of Oregon
Rutgers
The University of Utah