RAPIDS2: The SciDAC Institute for Computer Science, Data, and Artificial Intelligence

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

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

RAPIDS2: The SciDAC Institute for Computer Science, Data, and Artificial Intelligence

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

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

RAPIDS2 targets key computer science needs in four areas

Data Understanding

- Ensemble analysis
- Feature detection
- Production visualization
- In situ analysis

Artificial Intelligence

- Representation learning
- Surrogate modeling
- Automation

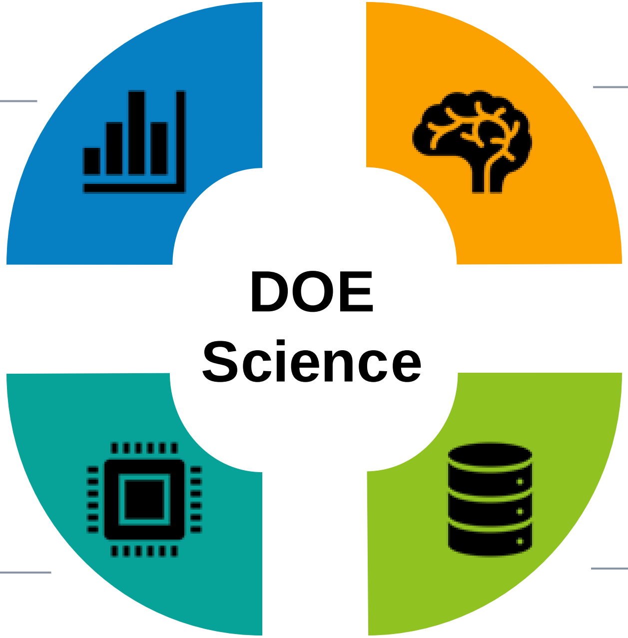

Platform Readiness

- Heterogeneous programming
- Autotuning
- Performance modeling and analysis
- Correctness

Scientific Data Management

- Storage and I/O
- Data Reduction
- Knowledge management
- Workflow automation

Engaging with SciDAC Partnerships

What's an Engagement?

- Engagements take a variety of forms
  - Short term: apply tools, etc.
  - Long term: co-design solutions
  - Can be multifaceted and/or multi-disciplinary
  - In collaboration with FASTMath
  - With new pilot programs
- Also engage with non-SciDAC teams, DOE facilities

Deep Bayesian Climate Emulator

WIth the ImPACTS Partnership (BER)

RAPIDS

Scientific Achievement

We present a probabilistic modeling framework for near-term climate forecasting based on deep learning modules. In addition to providing useful measures of predictive uncertainty, Bayesian versions of deep learning models outperform their deterministic counterparts in terms of predictive skill.

Significance and Impact

Following concerted national and international efforts over the past 70 years to model climate, comprehensive climate models have emerged as a powerful tool in helping unravel and better understand the myriad processes underlying climate and climate change. We incorporate Bayesian inference into cutting-edge network architecture, giving current deep learning models uncertainty estimation power.

Technical Approach

- We developed a Bayesian convolutional encoder–decoder deep network for uncertainty quantification.
- Our study considers Stein variational inference for exploring the high-dimensional posterior distribution of the network parameters.

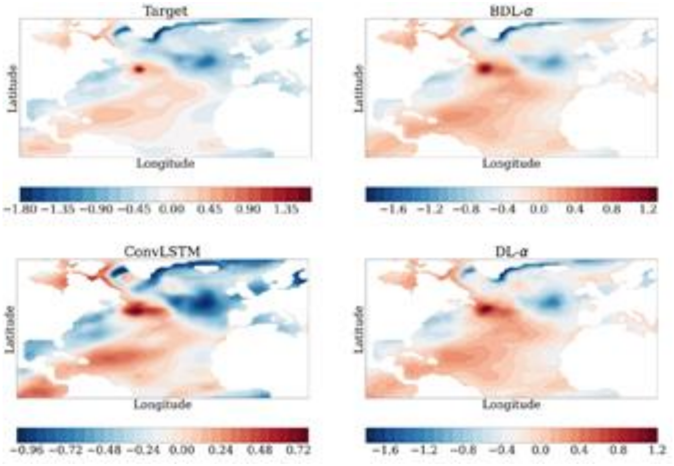

Figure 1. Prediction results

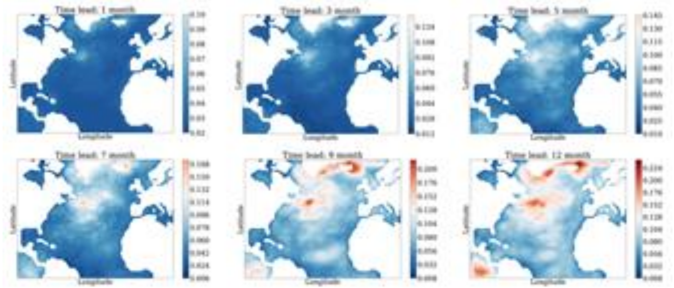

Figure 2. Predictive uncertainty quantification results

U.S. DEPARTMENT OF ENERGY | Office of Science

1. Luo, X., Nadiga, B.T., Park, J.H., Ren, Y., Xu, W. and Yoo, S., 2022. A Bayesian Deep Learning Approach to Near-Term Climate Prediction. Journal of Advances in Modeling Earth Systems, 14(10), p.e2022MS003058.

Brookhaven National Laboratory

Deep Bayesian Climate Emulator

WIth the ImPACTS Partnership (BER)

Scientific Achievement

We present a probabilistic modeling framework for near-term climate forecasting based on deep learning modules. In addition to providing useful measures of predictive uncertainty, Bayesian versions of deep learning models outperform their deterministic counterparts in terms of predictive skill.

Significance and Impact

Following concerted national and international efforts over the past 70 years to model climate, comprehensive climate models have emerged as a powerful tool in helping unravel and better understand the myriad processes underlying climate and climate change. We incorporate Bayesian inference into cutting-edge network architecture, giving current deep learning models uncertainty estimation power.

Technical Approach

- We developed a Bayesian convolutional encoder–decoder deep network for uncertainty quantification.
- Our study considers Stein variational inference for exploring the high-dimensional posterior distribution of the network parameters.

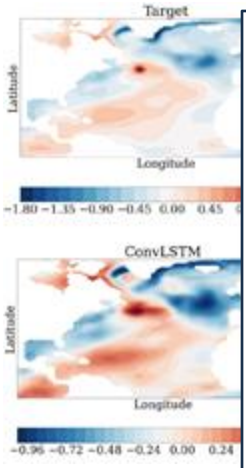

Fig

Contact: Shinjae Yoo

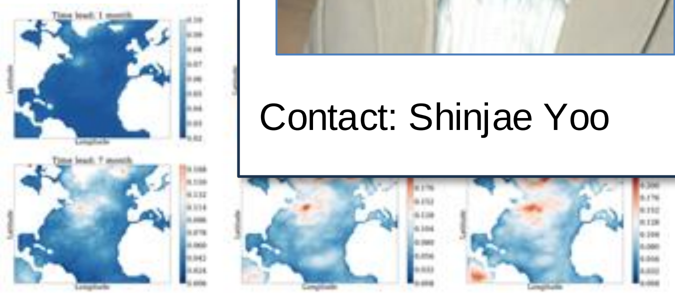

Figure 2. Predictive uncertainty quantification results

1. Luo, X., Nadiga, B.T., Park, J.H., Ren, Y., Xu, W. and Yoo, S., 2022. A Bayesian Deep Learning Approach to Near-Term Climate Prediction. Journal of Advances in Modeling Earth Systems, 14(10), p.e2022MS003058.

ML-based performance portability for Density Functional Theory

With the DECODE Partnership (BES)

Scientific Achievement

Developed a framework for optimizing the search of runtime configurations and code variants to improve the performance portability on HPC platforms for RT-TDDFT codes.

Significance and Impact

Improving the search speed for optimal performance configuration is critical for performance portability for applications with large space of algorithmic variants and runtime configurations. The developed methodology, despite focusing on RT-TDDFT computation, can be used with a wide set of applications.

Research Details

- Optimized search of the configuration space can leverage a wide set of techniques including Bayesian optimization and model-based estimation.
- Developed a framework to do transfer-learning, where the search on an HPC platform is used to reduce the number of search attempts on another platform.
- Focused on finding optimized configurations for the calculation involved in the conjugate gradient calculation part of rt-TDDFT.
- **Improved the performance by up to 9x,** reducing the search attempts by up to 86%.

Adrian P. Dieguez∗, Min Choi, Xinran Zhu, Bryan M. Wong and Khaled Z Ibrahim "ML-based Performance Portability for Time-Dependent Density Functional Theory in HPC Environments"2022 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS), Supercomputing 2022.

ASCR-BES DECODE SciDAC-5

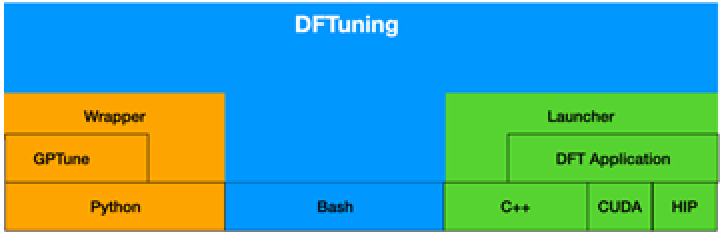

The RT-DFT Tuning stack for optimizing the performance on HPC machines.

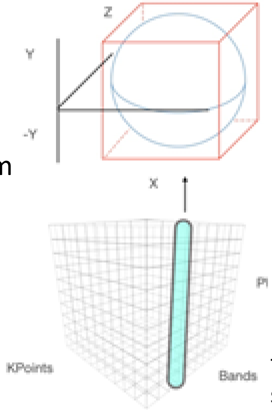

3D FFT mapping in rt-TDDFT

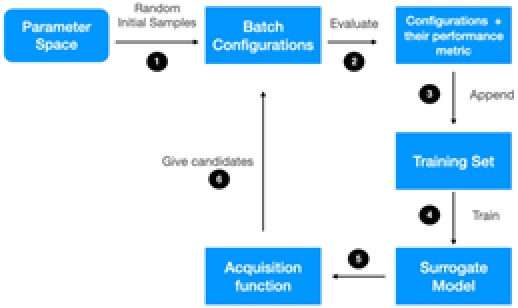

The internal process for searching the optimal space using Bayesian Optimization. Transfer learning leverages configuration on one platform to find optimal configuration on another.

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

ML-based performance portability for Density Functional Theory

With the DECODE Partnership (BES)

RAPIDS

Scientific Achievement

Developed a framework for optimizing the search of runtime configurations and code variants to improve the performance portability on HPC platforms for RT-TDDFT codes.

Significance and Impact

Improving the search speed for optimal performance configuration is critical for performance portability for applications with large space of algorithmic variants and runtime configurations. The developed methodology, despite focusing on RT-TDDFT computation, can be used with a wide set of applications.

Research Details

- Optimized search of the configuration space can leverage a wide set of techniques including Bayesian optimization and model-based estimation.
- Developed a framework to do transfer-learning, where the search on an HPC platform is used to reduce the number of search attempts on another platform.
- Focused on finding optimized configurations for the calculation involved in the conjugate gradient calculation part of rt-TDDFT.
- **Improved the performance by up to 9x,** reducing the search attempts by up to 86%.

Adrian P. Dieguez∗, Min Choi, Xinran Zhu, Bryan M. Wong and Khaled Z Ibrahim “ML-based Performance Portability for Time-Dependent Density Functional Theory in HPC Environments"2022 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS), Supercomputing 2022.

ASCR-BES DECODE SciDAC-5

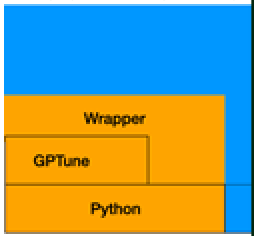

The RT-DFT Tuning sta
on HPC machines.

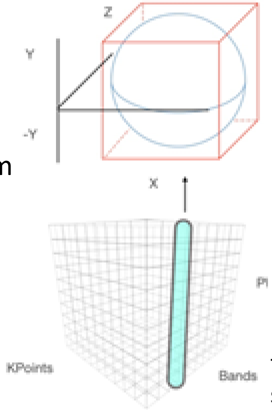

3D FFT mapping in rt-TDDFT

Contact: Khaled Ibrahim

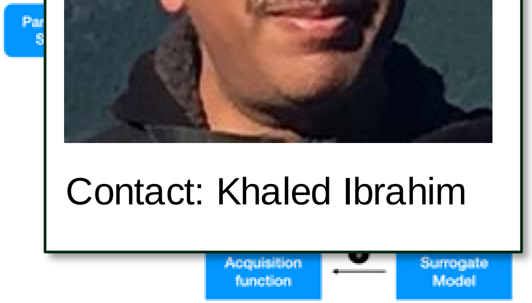

The internal process for searching the optimal space using Bayesian Optimization. Transfer learning leverages configuration on one platform to find optimal configuration on another.

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Improving I/O in the Sherpa and Pythia event generators

With the HPC Framework for Event Generation at Colliders Partnership (HEP)

RAPIDS

Scientific Achievement

Experiments at the Large Hadron Collider generate complex final states that must be simulated at high precision to identify the physics principles that determine how subatomic particles interact. Current simulations are too complex to achieve the precision goals set for the 2030s. We improved performance and scalability.

Significance and Impact

An I/O bottleneck limited the number of particle collisions we could simulate. It also broke strong scaling in parallel computations. By dramatically reducing the I/O overhead, we can now carry out the Monte-Carlo integration with a greater number of events. This reduces the statistical uncertainty of the simulations, and leads to theoretical predictions with higher confidence level.

Technical Approach

- Modified support libraries to expose collective HDF5 operations
- Adjusted use of HDF5 to allow further collective I/O optimizations
- Off-line converter to work around one HDF5 performance bug
- No machine-specific tuning required: approach is appropriate for any parallel file system

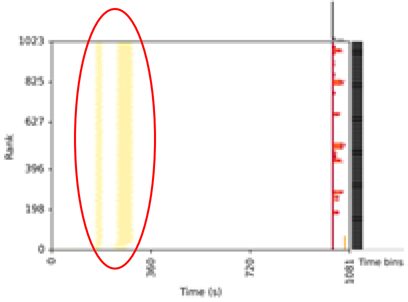

Initially code exhibited significant time spent in I/O.

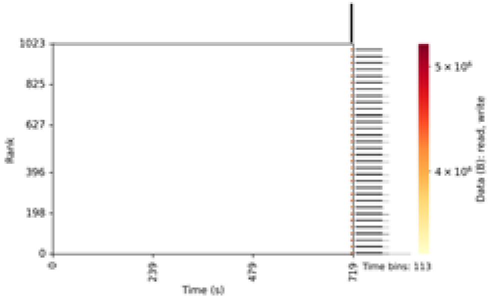

After optimization I/O time reduced to barely measurable; overall runtime reduced by 34%

I/O cost as the simulation scaled up alarmed scientists and prevented running higher resolution experiments. By enabling HDF5 collective data and metadata optimizations, and by making a minor adjustment to the arrangement of variables in the HDF5 file, the simulation could create HDF5 files significantly faster and simulate more events.

U.S. DEPARTMENT OF ENERGY | Office of Science

Argonne NATIONAL LABORATORY | Fermilab

Improving I/O in the Sherpa and Pythia event generators

With the HPC Framework for Event Generation at Colliders Partnership (HEP)

RAPIDS

Scientific Achievement

Experiments at the Large Hadron Collider generate complex final states that must be simulated at high precision to identify the physics principles which determine how subatomic particles interact. Current simulations are too complex to achieve the precision goals set for the 2030s. We improved performance and scalability.

Significance and Impact

An I/O bottleneck limited the number of particle collisions we could simulate. It also broke strong scaling in parallel computations. By dramatically reducing the I/O overhead, we can now carry out the Monte-Carlo integration with a greater number of events. This reduces the statistical uncertainty of the simulations, and leads to theoretical predictions with higher confidence level.

Technical Approach

- Modified support libraries to expose collective HDF5 operations
- Adjusted use of HDF5 to allow further collective I/O optimizations
- Off-line converter to work around one HDF5 performance bug
- No machine-specific tuning required: approach is appropriate for any parallel file system

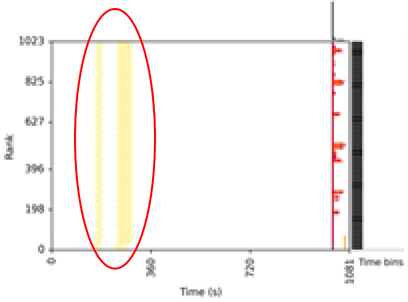

Initially code exhibited significant time spent in I/O.

I/O cost as the simulat ... ts and prevented running h ... By enabling HDF5 coll ... ta optimizations, and by m ... he arrangement of variable ... on could create HDF5 files ... te more events.

Contact: Rob Ross

U.S. DEPARTMENT OF ENERGY | Office of Science

Argonne NATIONAL LABORATORY Fermilab

Porting and tuning HBPS applications with TAU and APEX

With the HBPS Partnership (FES) (1/2)

Scientific Achievement

TAU and APEX's scalable performance & visualization tools from The University of Oregon OACISS Institute help High-Fidelity Boundary Plasma Simulation (HBPS) scientists achieve performance portability on the latest generation of DOE leadership-class systems.

Significance and Impact

TAU's performance measurement and analysis is helping XGC and GENE developers design new algorithms and parallelization strategies to fully utilize GPUs on Perlmutter and Polaris as well as on Frontier and Aurora development machines. TAU and APEX are also helping the HBPS team to improve OpenMP performance on CPUs. Improved computational efficiency and reduced memory usage allow for explorations of larger fusion energy problems with fewer compute nodes.

Technical Approach

- Measured, analyzed and tuned XGC computational kernels for AMD GPUs as well as NVIDIA GPUs using TAU and APEX.
- Helped tune XGC OpenMP regions to replace atomic operations with faster reductions.
- APEX is enabling portable comparisons of the GENE simulation to fully utilize new systems utilizing NVIDIA, AMD, or Intel GPUs.

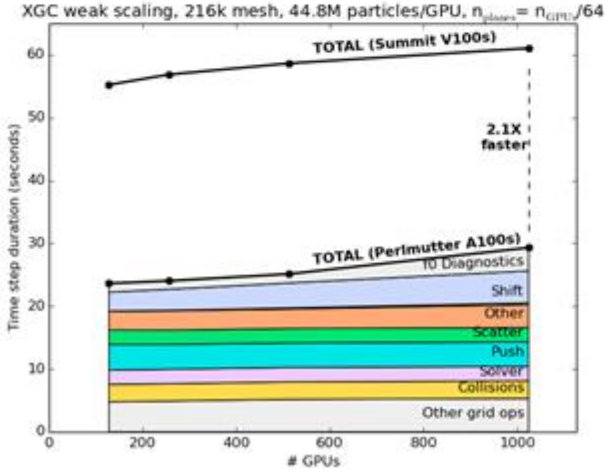

Figure: Scaling performance and timer breakdown of the XGC simulation on Perlmutter (NERSC) relative to Summit (OLCF), using the same number of GPUs (roughly 2.1x faster). Because electromagnetic simulations are a science priority going forward, the electron push kernel is less time dominant since it is sub-cycled fewer times. Scaling performance on Polaris (ALCF) is similar to Perlmutter. Credit: A. Scheinberg, Jubilee Development

Porting and tuning HBPS applications with TAU and APEX

With the HBPS Partnership (FES) (1/2)

Scientific Achievement

TAU and APEX's scalable performance & visualization tools from The University of Oregon OACISS Institute help High-Fidelity Boundary Plasma Simulation (HBPS) scientists achieve performance portability on the latest generation of DOE leadership-class systems.

Significance and Impact

TAU's performance measurement and analysis is helping XGC and GENE developers design new algorithms and parallelization strategies to fully utilize GPUs on Perlmutter and Polaris as well as on Frontier and Aurora development machines. TAU and APEX are also helping the HBPS team to improve OpenMP performance on CPUs. Improved computational efficiency and reduced memory usage allow for explorations of larger fusion energy problems with fewer compute nodes.

Technical Approach

- Measured, analyzed and tuned XGC computational kernels for AMD GPUs as well as NVIDIA GPUs using TAU and APEX.
- Helped tune XGC OpenMP regions to replace atomic operations with faster reductions.
- APEX is enabling portable comparisons of the GENE simulation to fully utilize new systems utilizing NVIDIA, AMD, or Intel GPUs.

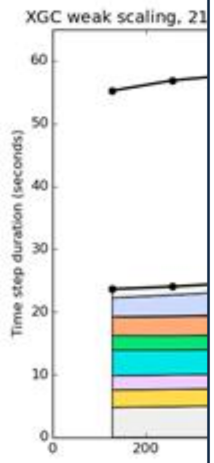

Figure: Scaling perfor on Perlmutter (NERSC number of GPUs (roug simulations are a scien is less time dominant s performance on Polari Scheinberg, Jubilee D

Contact: Kevin Huck

U.S. DEPARTMENT OF ENERGY | Office of Science

UNIVERSITY OF OREGON PPPL Jubilee Development

End-to-end simulation monitoring workflow

With the HBPS Partnership (FES) (2/2)

RAPIDS

Scientific Achievement

We have developed an end-to-end workflow to monitor simulation results at runtime. This workflow enables scientists to collaboratively track results produced by high performance simulations in order to ensure successful runs. In collaboration with the HBPS SciDAC partnership, we have demonstrated this workflow for the XGC simulation modeling edge plasma physics.

Significance and Impact

The high performance simulation workflows involve lengthy (non-interactive) runs on supercomputers followed by data movement and post-processing. The workflow that we developed empowers the science teams to produce visual output from the simulation at runtime and present them in an accessible way to a group of experts. This enables the correction of issues early in the simulation cycle, significantly increasing productivity.

Technical Approach

- The workflow begins with an in situ analysis engine, EFFIS, to analyze and visualize results as they are produced.
- The data extracts generated are then captured via ADIOS by an interactive dashboard, eSimmon, that can be accessed over the Web.
- The workflow enables the use of various VTK-m and ParaView for data extraction and interactive visualization over the Web.

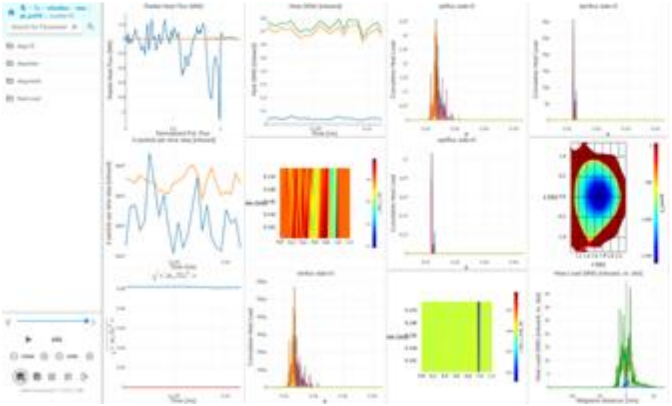

The eSimmon dashboard showing data extracts generated by EFFIS during an XGC simulation.

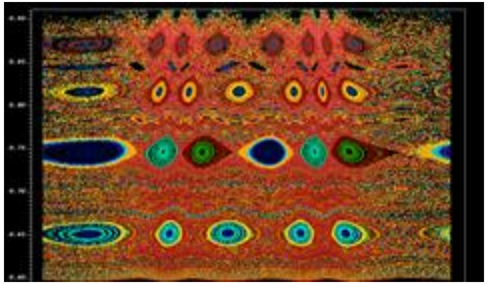

A Poincaré plot showing the cross section of many particle traces that iterate several times around a fusion reactor. These plots are generated in situ using VTK-m and displayed by the eSimmon dashboard

U.S. DEPARTMENT OF ENERGY | Office of Science

Oak Ridge National Laboratory

kitware

End-to-end simulation monitoring workflow

With the HBPS Partnership (FES) (2/2)

Scientific Achievement

We have developed an end-to-end workflow to monitor simulation results at runtime. This workflow enables scientists to collaboratively track results produced by high performance simulations in order to ensure successful runs. In collaboration with the HBPS SciDAC partnership, we have demonstrated this workflow for the XGC simulation modeling edge plasma physics.

Significance and Impact

The high performance simulation workflows involve lengthy (non-interactive) runs on supercomputers followed by data movement and post-processing. The workflow that we developed empowers the science teams to produce visual output from the simulation at runtime and present them in an accessible way to a group of experts. This enables the correction of issues early in the simulation cycle, significantly increasing productivity.

Technical Approach

- The workflow begins with an in situ analysis engine, EFFIS, to analyze and visualize results as they are produced.
- The data extracts generated are then captured via ADIOS by an interactive dashboard, eSimmon, that can be accessed over the Web.
- The workflow enables the use of various VTK-m and ParaView for data extraction and interactive visualization over the Web.

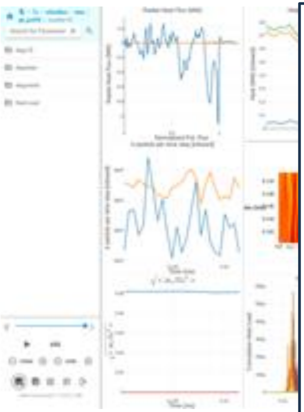

The eSimmon dashboard show an XGC simulation.

Contact: Berk Geveci

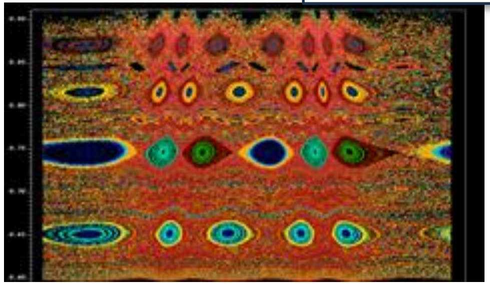

A Poincaré plot showing the cross section of many particle traces that iterate several times around a fusion reactor. These plots are generated in situ using VTK-m and displayed by the eSimmon dashboard

U.S. DEPARTMENT OF ENERGY | Office of Science

OAK RIDGE National Laboratory

kitware

Scientific Achievement

We developed a heterogeneous graph neural network (GNN) to predict optimal load in the maximum loadability problem and are working towards optimal blocker placement.

Significance and Impact

Geomagnetic disturbances resulting from intense solar activity caused by coronal mass ejections pose risks to the electrical grid by generating geomagnetically induced currents (GICs). Optimal blocker placement will significantly reduce the impact of this problem. Our GNN based machine learning prediction approach allows us to find high-quality solutions in shorter computation time than traditional solvers (e.g., Julia optimizers like PowerModelsGMD.jl).

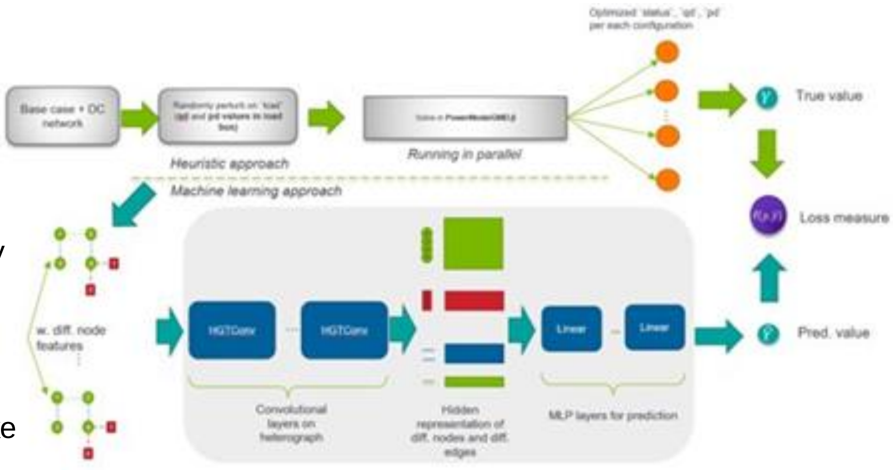

Technical Approach

- Modeled AC and DC networks with various node and edge features
- GNN predictive models for maximum loadability regression and for classification for GIC blocker placement
- Training for the loadability problem on 300 perturbed input graphs is faster than heuristic solver on a single graph
- DeepHyper to optimize hyperparameters and improve accuracy

The heuristic approach involves generating random perturbations on a given power grid and feeding those perturbations into one of the optimizers in the PowerModelsGMD.jl package, which finds the optimal real ("pd") and reactive ("qd") power demands. The machine learning approach on the bottom involves feeding the same power grid data into the heterogeneous GNN, passing through several convolutional and Multi-Layer Perception (MLP) layers. The true values from the optimizer and the predicted values from the machine learning model will then be used to compute the loss to evaluate the model's performance.

Machine Learning for Space Weather Mitigation

An ASCR – Office of Electricity Pilot Project

Scientific Achievement

We developed a heterogeneous graph neural network (GNN) to predict optimal load in the maximum loadability problem and are working towards optimal blocker placement.

Significance and Impact

Geomagnetic disturbances resulting from intense solar activity caused by coronal mass ejections pose risks to the electrical grid by generating geomagnetically induced currents (GICs). Optimal blocker placement will significantly reduce the impact of this problem. Our GNN based machine learning prediction approach allows us to find high-quality solutions in shorter computation time than traditional solvers (e.g., Julia optimizers like PowerModelsGMD.jl).

Technical Approach

- Modeled AC and DC networks with various node and edge features
- GNN predictive models for maximum loadability regression and for classification for GIC blocker placement
- Training for the loadability problem on 300 perturbed input graphs is faster than heuristic solver on a single graph
- DeepHyper to optimize hyperparameters and improve accuracy

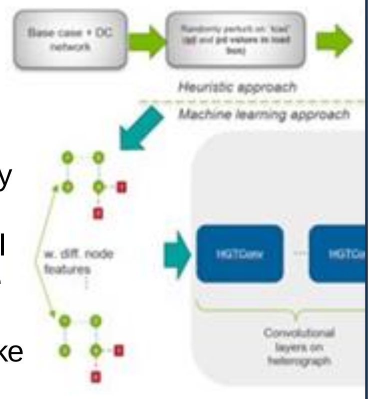

Contact: Prasanna Balaprakash

The heuristic approach i… iven power grid and feeding t… the PowerModelsGMD.jl pac… reactive ("qd") power de… bottom involves feeding … us GNN, passing through s… (MLP) layers. The true v… es from the machine learning model will then be used to compute the loss to evaluate the model's performance.

Adapting to Change

Artificial Intelligence

AI and ML for accelerated discovery for science applications

Representation

- Bayesian Neural embedding for better uncertainty quantification and prior knowledge integration
- Physics constrained AE compressed large scale volumetric datasets by O(10^2) with negligible errors
- Explainable AI (xAI) models developed and applied to climate and material science.

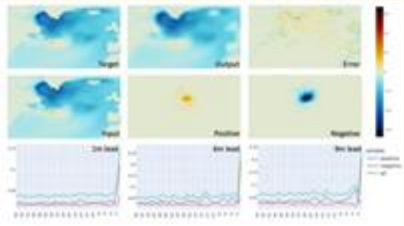

Acceleration

- Developed (physics constrained) surrogate models for various SciML applications with orders of magnitude speed up
- Evaluated and developed scaling algorithms such as graph neural network, kernel approximation, etc.
- AI acceleration on edge device using FPGA or AI accelerators

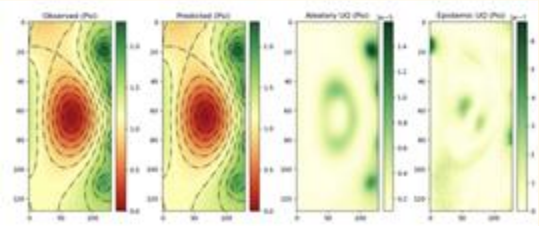

Automation

- Developed a novel AI/ML optimization method to tune storage service parameters for SciDAC HEP workflow resulting in (2x to 10x speedup)
- Developed an adaptive training algorithm that dynamically schedules the work among multiple GPUs to minimize imbalance

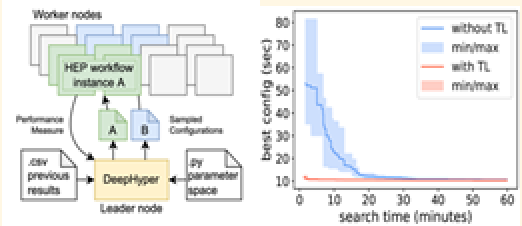

Machine Learning Based Volumetric Compression

Crosscutting AI and SDM areas

RAPIDS

Scientific Achievement

We have developed neural networks based constrained autoencoders for compressing datasets that also maintain user provided linear or non-linear equality constraints. These autoencoders can be used to compress large scale volumetric datasets by two orders of magnitude while achieve near floating point errors on Quantities of interest (QoI).

Significance and Impact

In scientific data compression down-stream quantities derived from the original data are crucial for post-analysis in order to analyze physical phenomena, and therefore these Quantities of Interest (QoI) are required to be preserved during compression.

Technical Approach

It is difficult to enforce constraints in neural networks since top layer predictors are usually nonlinear and we wish to use standard convex optimization methods and strong duality. To overcome this, we introduce a new saddle-point Lagrangian with auxiliary predictor variables on which constraints are imposed. Eliminating the auxiliary variables leads to a dual minimization problem on the Lagrange multipliers to satisfy the linear constraints. A back projection approach is then used to extend these ideas for non-linear constraints.

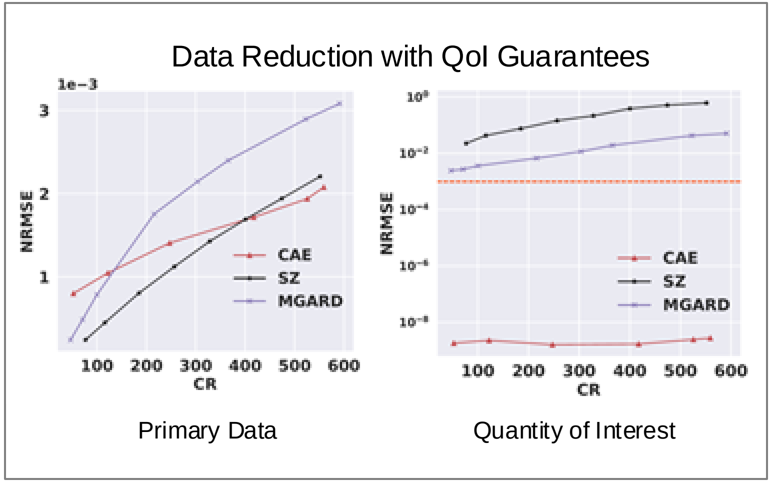

Normalized Root Mean Squared Error of constrained auto encoder (CAE) and comparison with MGARD and SZ. The left figure is the error on the primary data and the right figure is the error on one of the Quantities of Interest

Publications

- Lee, J. et al. "Error-bounded learned scientific data compression with preservation of derived quantities." Applied Sciences, 12(13):6718, Jul 2022.
- Lee, J. et al. "Constrained Autoencoders: Incorporating equality constraints in learned scientific data compression." 2023 Data Compression Conference (DCC), Snowbird, UT, USA, 2023. (an extended version submitted to IEEE Transactions on Neural Networks and Learning Systems, 2023)

U.S. DEPARTMENT OF ENERGY | Office of Science

UF UNIVERSITY of FLORIDA

OAK RIDGE National Laboratory

Machine Learning Based Volumetric Compression

Crosscutting AI and SDM areas

RAPIDS

Scientific Achievement

We have developed neural networks based constrained autoencoders for compressing datasets that also maintain user provided linear or non-linear equality constraints. These autoencoders can be used to compress large scale volumetric datasets by two orders of magnitude while achieve near floating point errors on Quantities of interest (QoI).

Significance and Impact

In scientific data compression down-stream quantities derived from the original data are crucial for post-analysis in order to analyze physical phenomena, and therefore these Quantities of Interest (QoI) are required to be preserved during compression.

Technical Approach

It is difficult to enforce constraints in neural networks since top layer predictors are usually nonlinear and we wish to use standard convex optimization methods and strong duality. To overcome this, we introduce a new saddle-point Lagrangian with auxiliary predictor variables on which constraints are imposed. Eliminating the auxiliary variables leads to a dual minimization problem on the Lagrange multipliers to satisfy the linear constraints. A back projection approach is then used to extend these ideas for non-linear constraints.

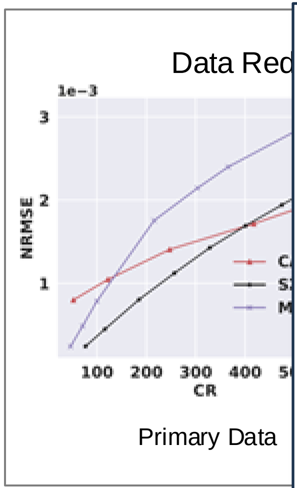

Normalized Root Mean S
and comparison with MG
primary data and the righ
Interest

Contact: Scott Klasky

Publications

- Lee, J. et al. "Error-bounded ... preservation of derived quantities." Applied Sciences, 12(13):6718, Jul 2022.
- Lee, J. et al. "Constrained Autoencoders: Incorporating equality constraints in learned scientific data compression." 2023 Data Compression Conference (DCC), Snowbird, UT, USA, 2023. (an extended version submitted to IEEE Transactions on Neural Networks and Learning Systems, 2023)

U.S. DEPARTMENT OF ENERGY | Office of Science
UF UNIVERSITY of FLORIDA
OAK RIDGE National Laboratory

Data Understanding

Facilitating the understanding of large and complex science data through robust and scalable analysis and visualization methods

Analysis

- Accelerated merge tree computation on GPUs by 10x
- Developed asynchronous, load-balanced union-find algorithm
- Developed new latent representation for particle features
- Developed a hierarchical super-resolution method for multiscale scientific data

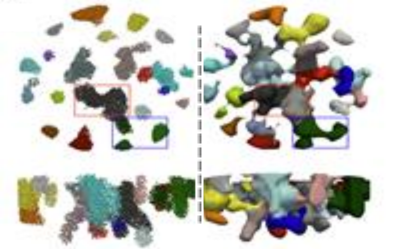

Visualization

- Accelerated Poincare plot computation on GPUs by 50x
- Investigated effects of lossy data reduction on volume rendering
- Improved distributed rendering performance with image compression
- Produced a science video about sea-level change
- Developed a surrogate to visually explore simulation parameters

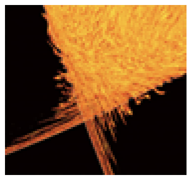

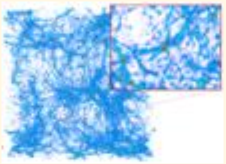

Infrastructure

- Extended DIY with asynchronous communication + direct GPU messaging
- Automated optimization of generator tuning parameters
- Developed end-to-end simulation monitoring workflow for XGC
- Demonstrated VTK-m performance portability on new architectures
- Developed a general purpose data model for visualization of streaming data

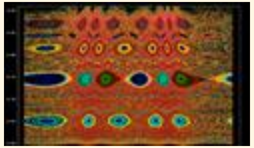

FTK2: Meshing the Deforming Spacetime for Feature Tracking

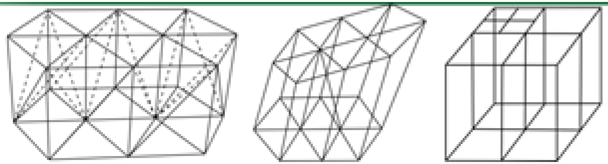

Spacetime meshing for non-changing meshes (left), time-varying meshes (middle), and AMR (right)

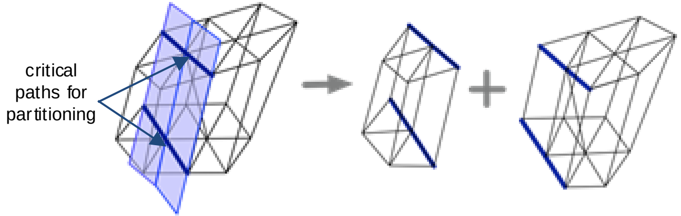

Divide: divide the spacetime by faces generated by shortest paths on both “lower” and “upper” meshes

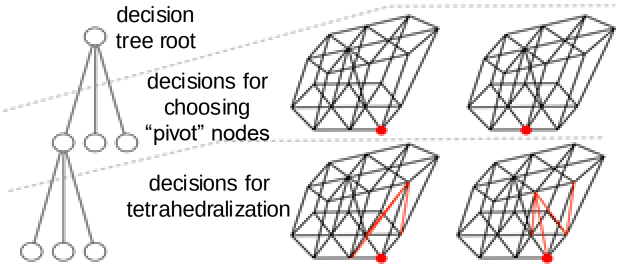

Conquer: search on decision tree to eliminate mesh nodes one by one by removing tetrahedra around them

Scientific Achievement

A novel 4D meshing capability that fills the gap between spatial and temporal varying mesh discretization for advanced analysis and visualization

Significance and Impact

Representing the deforming spacetime with simplicial complexes makes it possible to robustly track and visualize geometric and topological features such as critical points, vortex core lines, and levelsets in scientific data

Technical Approach

- **Driver applications**
  (1) magnetic confinement fusion where scientific variables interpolate over magnetic field lines in space and time
  (2) rotational mechanics where space twists over rotating frame-of-reference
  (3) adaptive mesh refinement (AMR) that changes discretization over time
- **Sewing the “lower” and “upper” meshes** by divide-and-conquer
  (1) 4D domain decomposition based on spacetime face/edge connectivities
  (2) Subdomain decomposition based on shortest-path search
  (3) Decision-tree based search to find optimal triangulations of subdomains
- **Software integration** into FTK2, the next generation of FTK (the feature tracking kit): https://github.com/hguo/ftk2

Platform Readiness

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

Performance Modeling/Analysis

- Measured, analyzed and tuned XGC kernels for GPU using TAU and APEX
- Developed a data dependency aware performance model for sparse triangular solvers

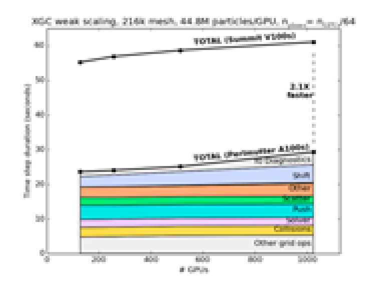

Portable Programming

- IRIS: portable runtime system exploiting multiple heterogeneous programming systems
- Explored use of OpenACC and OpenMP to portably accelerate the Fortran-based EFIT code (FES)
- Developed an algorithm-centric performance portability metric

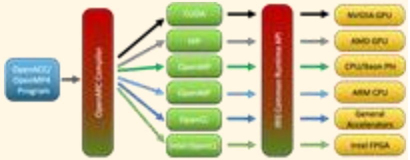

Autotuning

- Flexible autotuning framework including Bayesian optimization and model-based estimation
- OpenMP language extensions for runtime adaptation
- Framework for transfer-learning in autotuning

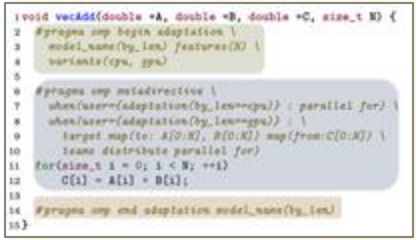

Correctness

- Support for verification of Fortran in CIVL framework through correctness-preserving transformations to CIVL language
- New model checking technique that can be used to verify race-freedom for all sequentially-consistent executions

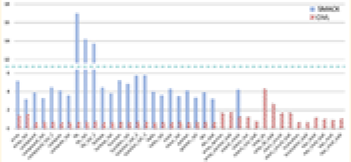

Verification of Data Race Freedom for OpenMP Programs

RAPIDS

Scientific Achievement

Parallel programs often contain subtle defects known as ***data races***. A data race occurs when two concurrently executing threads access a variable without proper synchronization. These programs may crash or produce results that appear believable but are inaccurate. Programmers need better tools to detect races or verify that a program is race-free. We have developed a novel model checking approach, implemented in the CIVL model checker, that can verify absence of races or produce precise diagnostic information when a race exists, for a significant subset of OpenMP.

Significance and Impact

The new tool has been applied to over one hundred benchmark C/OpenMP programs and is remarkably accurate. The programs include those from the DataRaceBench suite which fall into the supported language subset, as well as programs using “unstructured” parallelism. With further development, it will become an increasingly practical tool for developers.

Technical Approach

- Model checking and symbolic execution are used to analyze the program within a “small scope”
- A novel partial order reduction dramatically reduces the number of interleavings which need to be explored, leading to greater scalability

```
int main() {
  int i, len=100, a[100];
  for (i=0; i<len; i++) a[i]=i;
#pragma omp parallel for
  for (i=0; i<len-1; i++) a[i+1]=a[i]+1;
  printf("a[50]=%d\n", a[50]);
  return 0;
}


> civl verify -input_omp_thread_max=10 \
  DRB029-truedep1-orig-yes.c

a[50]=50
Data-race detected:
  {a[1], a[3], ..., a[97]}
read by thread 1 intersects
  {a[1], a[3], a[5], ..., a[99]}
written by thread 0.
```

Example: CIVL detects a data race in a DataRaceBench benchmark, reporting precise diagnostic information to user.

U.S. DEPARTMENT OF ENERGY | Office of Science

Wenhao Wu, Stephen Siegel (University of Delaware); Jan Hückelheim, Paul Hovland (Argonne). To appear in CAV2023.

University of Delaware | Argonne National Laboratory

Scientific Data Management

Enabling efficient movement and management of data in a scientific campaign

Storage and I/O

- Understanding and optimizations for I/O and storage on DOE resources. E.g. Frontier and Perlmutter
- In situ data management for code coupling and connections to analysis and visualization
- Optimal and scalable I/O for streaming I/O from experiments to HPC facilities
- Integration and optimization of HPC I/O techniques into ML software

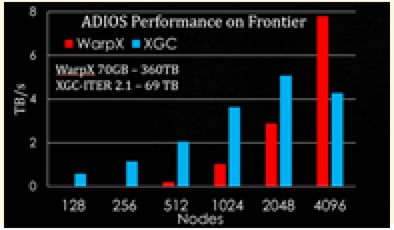

Knowledge Management

- Partner with applications for efficient data access across the data lifecycle
- Enhance indexing and querying for data and metadata in scientific applications
- Schemas to integrate applications with visualization frameworks

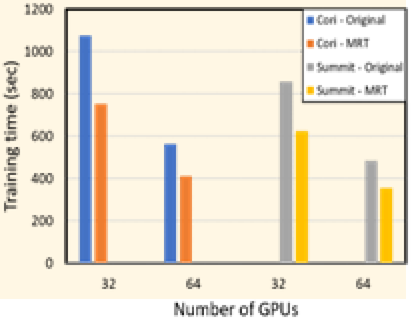

Workflow Automation

- Optimize workflow infrastructure optimized for current and next-generation HPC systems
- Automatic parallelizing custom analysis functions for complex data
- Collaborative web-frameworks to monitor simulations
- Efficient command & control capabilities

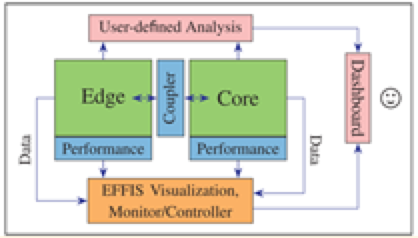

Data Reduction

- Integrate ML, constrained optimization and statistical similarity with lossy reduction
- Investigate the use of adaptive reduction techniques
- Incorporate reduction for storage and real-time workflows

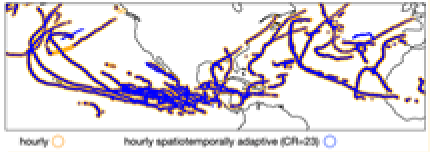

Intelligent Draining to Reduce I/O Contention in Staging-based Workflows

Scientific Achievement

We created the RISE framework, an extension to existing data staging middleware that intelligently manages the transfer of staged data in order to avoid contention between data staging operation commonly in use in extreme-scale scientific workflows. RISE identifies applications' data access patterns and moves data towards data consumers only when the network is expected to be idle, reducing the impact of asynchronous background data movement upon critical data read/write requests.

Significance and Impact

We demonstrated that intelligent draining and data placement reduced I/O contention using common coupled workflow data and execution patterns as well as evaluated performance on a coupled DNS-LES simulation using S3D. We observed an improvement in staging write-responsiveness of 22% vs typical out of process storage without any contention avoidance.

Technical Approach

- Add writer-local staging area, with channels for dynamic relocation (draining) to staging server processes.
- Extend indexing to hybrid transfer: direct transfer of undrained objects from writers to readers and indirect transfer via a staging space.
- RISE uses a Markov model to predict periods of staging quiescence.
- When sufficient quiescent periods are anticipated, data objects are drained from writer-local staging.

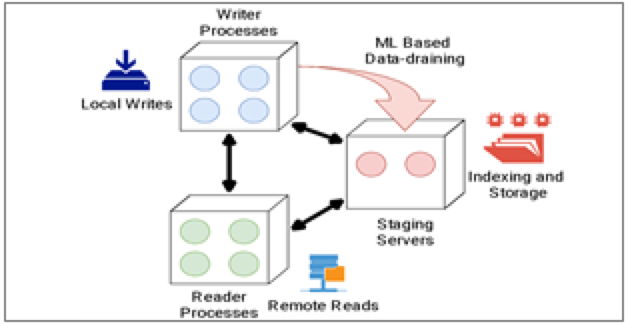

RISE (above) makes data immediately available to readers and also offloads the data in the background to the staging servers. This provides comparable write responsive performance comparable to writer-local storage without the long-term memory costs and without the network costs of Naive-drain (below).

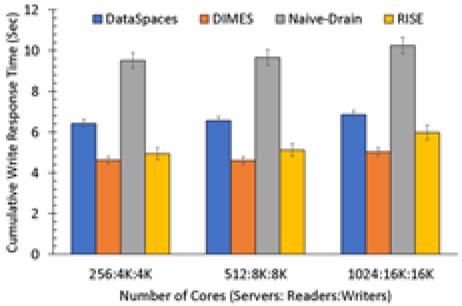

U.S. DEPARTMENT OF ENERGY | Office of Science

Subedi, P., Davis, P. E., & Parashar, M. (2021, September). Rise: Reducing I/O Contention in Staging-Based Extreme-Scale In-situ Workflows. In *2021 IEEE International Conference on Cluster Computing (CLUSTER)* (pp. 146-156). IEEE.

THE UNIVERSITY OF UTAH

Wrapping Up…

Next Steps

- Help domain scientists **productively use exascale platforms** via code portability, improved standards, and better tools for understanding mixed CPU-GPU performance,
- **Develop and apply ML and AI methods** for science applications and in other areas such as compression, data understanding, and data movement, and
- Further **support complex workflows** through improved data movement, better dashboards, and a richer collection of in situ data analysis techniques.

The RAPIDS Team

Green outline means they're here in person! Say hi!