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

Applications Engagement and Community Outreach

Anshu Dubey, Sam Williams, Philip C. Roth, Seyong Lee, Tom Peterka, Hongzhang Shan, Chad Steed

Overview

Interface between applications and other RAPIDS focus areas

Members of other areas matrixed into AE

Modes of Engagement

Collaborations

Establish working relationships between projects

- *Existing working relationships are continuing*
- *Working relationships are being established where RAPIDS members are part of partnership projects*
- *We are exploring and establishing new relationships*

Tiger Teams

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*
- ***Three to four Tiger Team activities per year***
  - *Prioritization based on available expertise and potential impact on the application*
  - *Resources budgeted in RAPIDS appropriately.*

Consultation and Training

Outreach activities

- ***Trainings, tutorials, and webinars***
  - *Disseminate RAPIDS technologies and approaches to a broader audience*
  - *Tutorials at various conferences*
- *Best practices documentation and consultation*
- *Hackathons and coding camps*

New Connections

In Progress

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

Being Developed

Development of Terrestrial Dynamical Cores
Software architecture and process design

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

Application Collaborations

HEP Event Reconstruction with Cutting Edge Computing Architectures
Performance analysis and optimization, autotuning, help design algorithms

Simulation of Fission Gas in Uranium Oxide Nuclear Fuel
Performance optimization of Xolotl cluster dynamics code using RAPIDS tools and performance monitoring of Xolotl/MOOSE coupled simulations

Coupling Approaches for Next-Generation Architectures (CANGA)
Use Decaf to loosely couple scalable parallel in situ particle tracing with MPAS-O ocean model.

HEP Data Analytics on HPC
Use DIY to parallelize and scale NOvA neutrino event generation and Feldman-Cousins correction; redesign HEP object storage with HEPnOS. I/O collaboration.

Towards Exascale Astrophysics of Mergers and SuperNova
Performance portable kernels, productivity, measurements

Nuclear Low Energy Initiative
Performance analysis and optimization, autotuning, help design algorithms

Plasma Surface Interactions: Predicting the Performance and Impact of Dynamic PFC Surfaces (PSI-2)
Performance optimization of Xolotl cluster dynamics code, GITR impurity transport code, and coupled simulations for FY18 DOE PMI Theory Milestones

Catalysis
Performance analysis and optimization of VASP, CP2k, etc.. on KNL, and GPUs. Optimizations for mitigation network communication at scale.

CTTS
Performance analysis and optimization of SuperLU, SpTS, M3D-C1, and NIMROD on KNL and GPUs.

ISEP
Performance analysis and optimization of GTC and GTC-P on KNL and GPUs.

Inference At Extreme Scales
Efficient search of the parameter space,
scalable hyperparameter search for deep learning methods

Fastmath Collaborations

PETSc and Xolotl

SuperLU and Oreo

Image compression using deep learning techniques

Liaisons

- Poster Session @ FastMath Kickoff
- Roofline engagement with NERSC
- Facility liaisons

Outreach

- SciDAC, LCF, NERSC, …
- Roofline engagement with vendors
- Conducted Roofline Tutorials/Training...
  - November '17 @NERSC
  - November '17 @ SC
  - February '18 @ ECP

Early Results

Reduction of Xolotl Memory Requirement

Xolotl cluster dynamics code exhibited out-of-memory errors with smaller problem sizes than expected. Used RAPIDS expertise, internal Xolotl timers, and DOE ASCR-funded tools to identify problem during PETSc solver (FASTmath) initialization. Proposed API change to PETSc to resolve problem. Reduced Xolotl peak memory footprint by 88% with previous largest problem size possible on OLCF Eos Cray XC30.
(P.C. Roth with S. Blondel, FGS/PSI2)

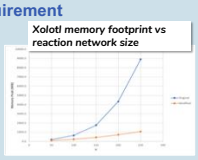

KORC Performance Optimization

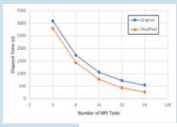

Conducted initial triage of Kinetic Orbital Runaway electron Code (KORC) with RAPIDS expertise and DOE ASCR-funded tools. Identified opportunities to accelerate code plus early optimization opportunity. Applying that optimization resulted in ~2x performance improvement. (P.C. Roth, S. Lee with L. Carbajal Gomez, SCREAM)

KORC run time vs number of MPI tasks on OLCF Eos Cray XC30

Multivariate, Temporal Visual Analytics for Climate Model

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

EDEN Parallel Histogram Summary View

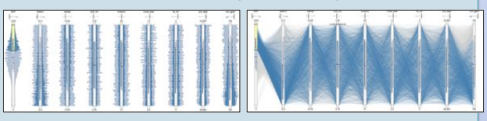

The screenshots show one insight found during parameter sensitivity analysis for realistic values of GPP, a model output that measures photosynthesis in plants.The plots helped scientists see that high values of GPP are associated with low leaf carbon to nitrogen ratio values (leafcn) and low critical 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.

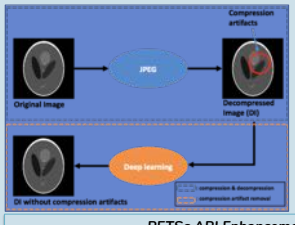

- Compression techniques such as JPEG can significantly reduce the size of images, but there are two main challenges for scientific image compression: blocking artifacts and blurring.
- We developed deep learning techniques to remove artifacts.

PETSc API Enhancement
Proposed new API function to reduce large memory requirement during PETSc initialization within Xolotl; accepted to PETSc repository

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