This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

SAND2022-4643C

ALPINE: In Situ Visualization, Analysis, and Infrastructure for ECP Science

James Ahrens (PI), Los Alamos National Laboratory   Hank Childs (Deputy PI), University of Oregon
Terece L. Turton (PC), Los Alamos National Laboratory

Los Alamos National Laboratory: Jim Ahrens, Ayan Biswas, Soumya Dutta, Terece Turton
Lawrence Livermore National Laboratory: Eric Brugger, Timo Bremer, Cyrus Harrison, Nicole Marsaglia, Sergei Shudler
Lawrence Berkeley National Laboratory: Gunther Weber, Oliver Rübel
Oak Ridge National Laboratory: Dave Pugmire, Jieyang Chen
Argonne National Laboratory: Silvio Rizzi, Joe Insley, Victor Mateevitsi
Sandia National Laboratories: Janine Bennett, Marco Arienti, Jorge Salinas.
Kitware: Berk Geveci, Utkarsh Ayachit, Dan Lipsa, Corey Wetterer-Nelson
University of Oregon: Hank Childs, Yuya Kawakami
University of Utah: Valerio Pascucci, Xuan Huang, Steve Petruzza
University of Leeds: Hamish Carr, Petar Hristov.

ALPINE TEAM MEMBERS

Overview

ALPINE is leveraging current analysis and visualization software while developing new capabilities to deliver **in situ visualization, infrastructure, and data analysis to ECP Applications.** The project focuses on these main development areas:

- Deliver **visualization and analysis algorithms** that are critical for ECP Applications as the dominant analysis paradigm shifts from *post hoc* to *in situ*.
- Deliver an exascale-capable **infrastructure** for the development of in situ algorithms and deployment into existing applications, libraries, and tools.
- **Integrate** our algorithms and infrastructure with ECP Application software and workflows.
- **Integrate** ECP Software Technologies into our infrastructure to meet the needs of ECP Applications.

Ongoing Infrastructure, Algorithm Development & Integration Efforts

Optimal Viewpoint Metrics based on data properties can automate in situ visualization decisions to only capture interesting views. Metrics left to right: data entropy, depth entropy, max depth, projected area.

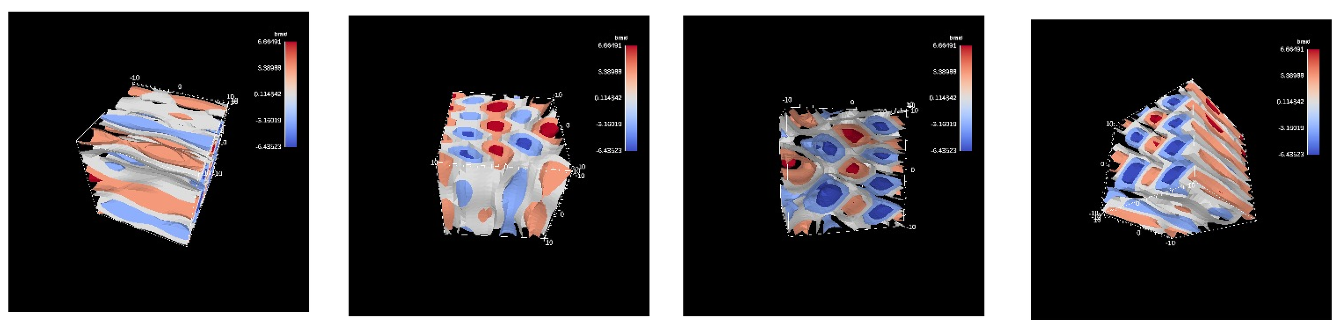

NekRS ⟷ Ascent

Ascent-NekRS initial testing on low-order, linear elements. Using the Ascent-NekRS integration, volume-rendered flow-field images were generated at multiple time-steps for a simple test case, *couette* flow, which has an analytical solution.

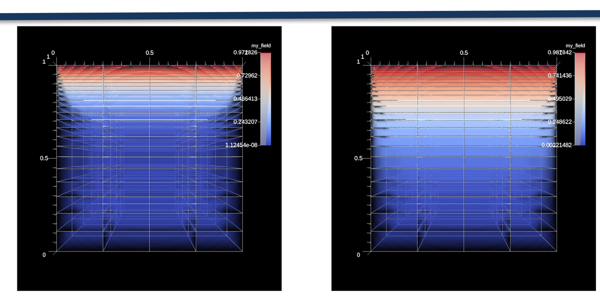

Image rendering of a simple, couette flow at early (t=0.1) and later time step (t = 1.0). The field shown represents velocity magnitude.

Task-based hierarchical feature extraction algorithm based on segmented merge trees. The algorithm is implemented using a multi-runtime abstraction layer, BabelFlow, which can be used to execute arbitrary analysis and visualization dataflows using different task-based runtime systems.

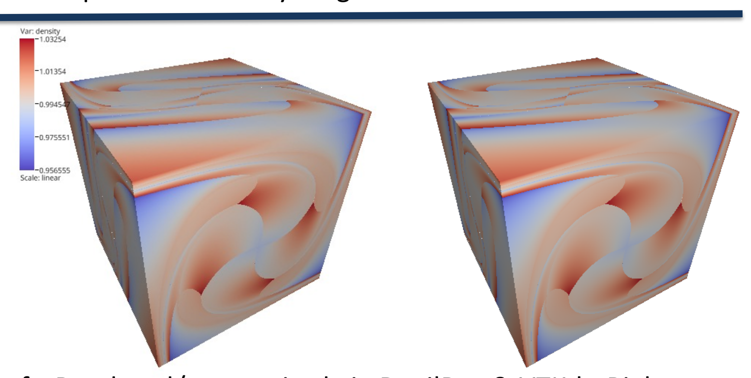

Left: Rendered/composited via DevilRay & VTK-h. Right: Rendered/composited via DevilRay & BabelFlow's RadixK algorithm.

Ascent<>AMR-Wind Integration ported to Crusher

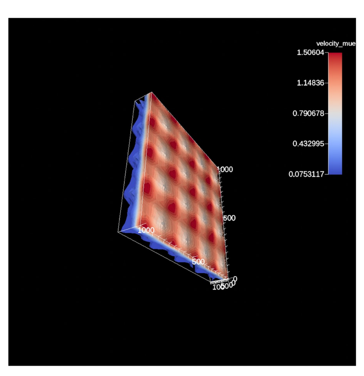

ALPINE Infrastructure on Frontier EAS

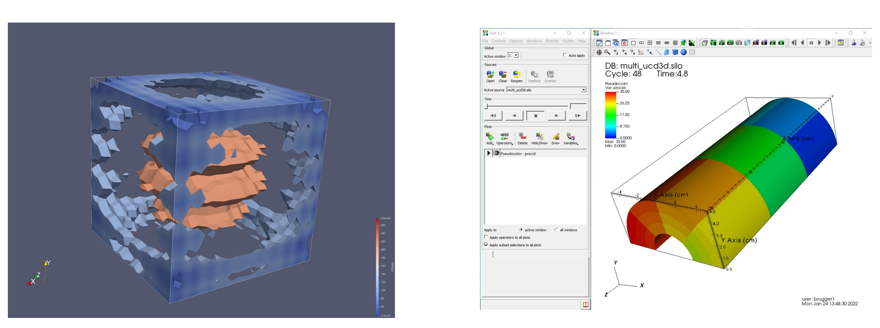

Images from ParaView & VisIt Spock porting activities. ParaView and VisIt compiled on the CPU using OSMesa. Left, a pvbatch image from a rendering test. Right: a VisIt client-server demo of a parallel run. Ascent has been ported to Crusher GPUs.

ALPINE LCF Teams: Delivering on Exascale platforms

ALPINE teams at Oak Ridge and Argonne are providing support for ALPINE infrastructure and algorithms on early access systems to meet the data and visualization needs of ECP applications.

Software Stack and Infrastructure

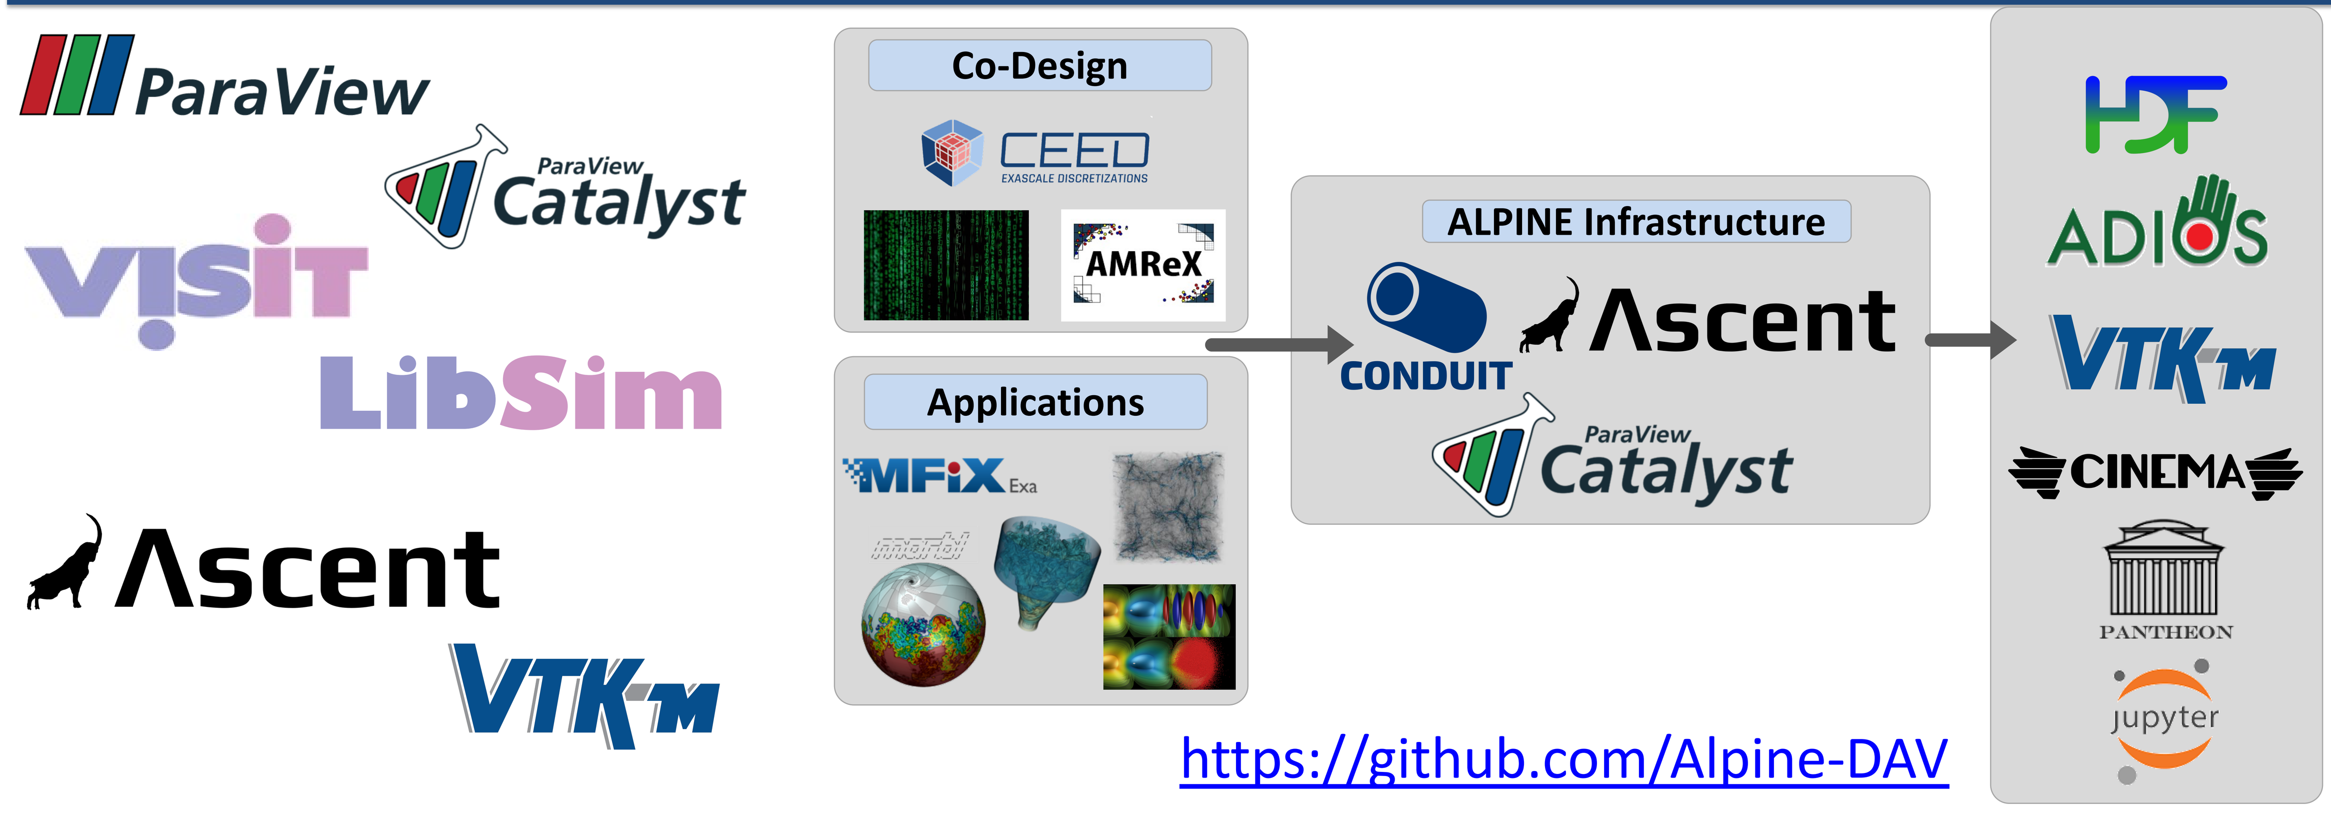

Example Integrations & Workflows

MFIX Exa ⟷ AMReX ⟷ ParaView Catalyst ⟷ VTK-m ⟷ In situ algorithm ⟷ post hoc analysis

Goal: in situ data reduction & feature detection with post hoc interactive analysis via Cinema

In situ **statistical feature detection** detects features in particle data sets using statistical data modeling and probabilistic similarity measures.

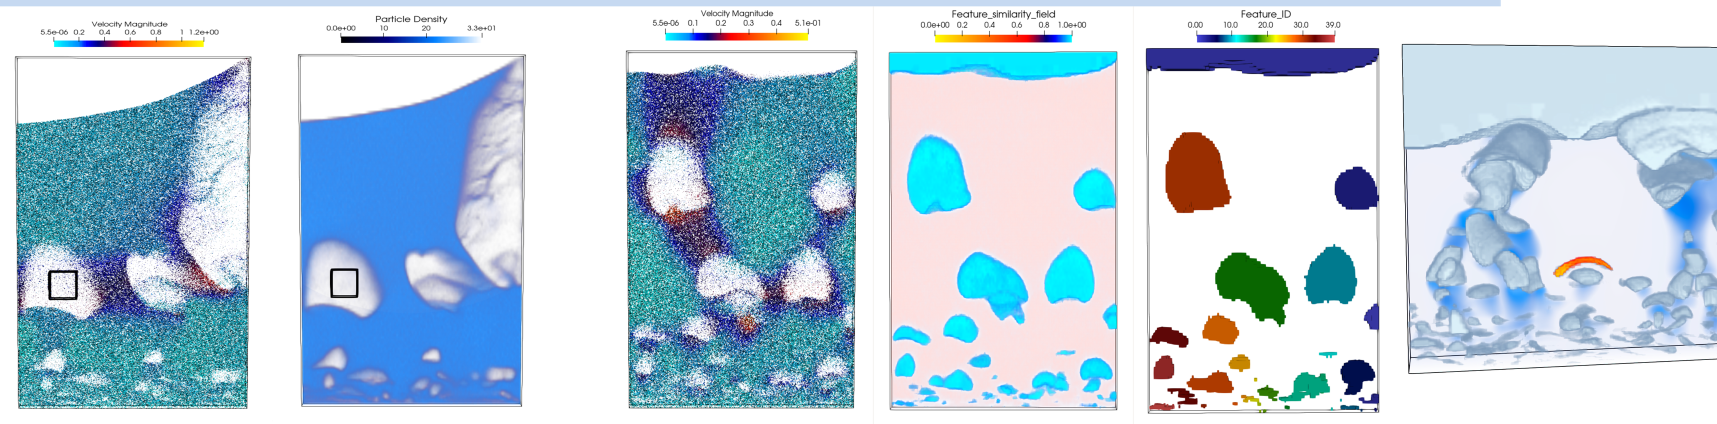

Above: Figures illustrate process to identify "voids" or bubbles in MFIX-Exa particle data. L to R: (1) Void selection: box shows region of interest with low particle density. (2) Voids as seen in density field. (3) Voids in particle field. (4) Feature similarity field – how similar to region of interest. (5) Voids are isolated by thresholding for analysis and characterization. (6) Production visualization.

Combustion-Pele ⟷ AMReX ⟷ Ascent ⟷ ExaLearn **In situ anomaly detection**

Goal: in situ anomaly detection to trigger downstream analysis for auto-ignition events

ALPINE, ExaLearn and Pele collaboration: AI/ML techniques independently developed in ExaLearn are being deployed in situ with Ascent technology from ALPINE to detect auto-ignition regions in a PeleLM simulation.

Demonstrated: GPU-scalable in situ detection of auto-ignition regions using co-kurtosis tensor-based anomaly detection.

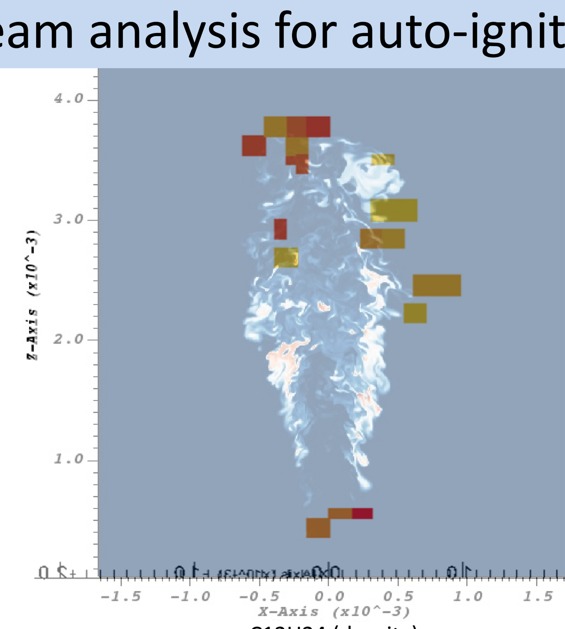

C12H24 (density)

Left: Simulation snapshot of an intermediate species overlaid with AMR boxes flagged by the anomaly detection algorithm.
Right: production visualization via Ascent.
Image left courtesy of M. Rieth, J. Chen (Pele), M. Arienti, M. Larsen, J. Bennett (ALPINE), H. Kolla (ExaLearn).
image right courtesy of M. Larsen (ALPINE).

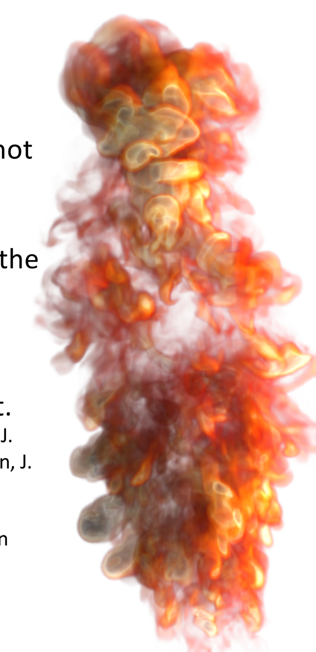

WarpX ⟷ AMReX ⟷ Ascent / VisIt / ParaView ⟷ VTK-m ⟷ In situ algorithm ⟷ post hoc analysis

Goal: Exploratory visualization and analysis to understand physics of plasma-based particle accelerators

Topological analysis is used to identify most relevant contours & create isosurface visualizations in situ; saving resulting images for post hoc analysis. Images are saved to a Cinema DB in a format that supports arbitrary combination of contours during post hoc visualization. Right: Most relevant contours in WarpX simulation selected using two importance measures: persistence, volume.

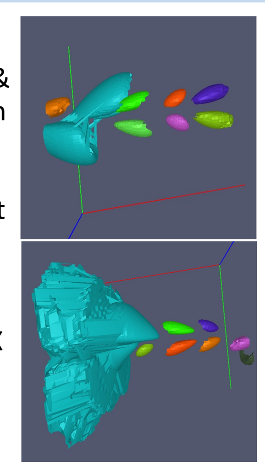

Below: visualization of particles and fields produced in situ can be used to generate movies showing complex physics behavior over time.

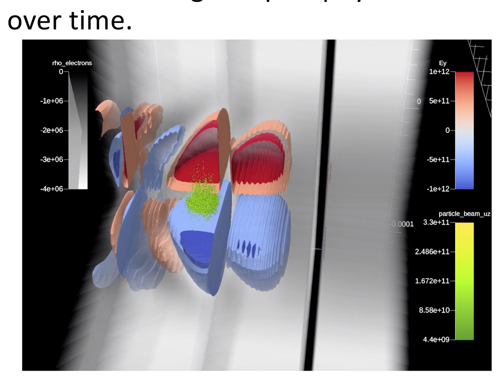

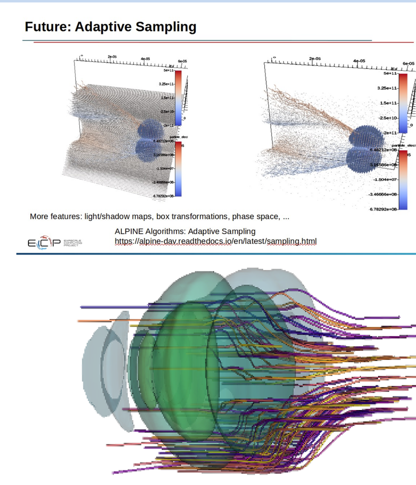

Left top: Sampling on transverse momentum reveals particles behaving unexpectedly.
Image curtesy of A. Huebl (WarpX)

Left bottom: VTK-m WarpX-specific particle advection algorithm visualized post hoc in VisIt.
Image courtesy of H. Childs, A. Yenpure (VTK-m)

Right: Post hoc analysis workflow via ParaView to visualize EM field streamlines in Cinema Jupyter notebook.
Image courtesy of R Bujack, TL. Turton, DH Rogers (Cinema)

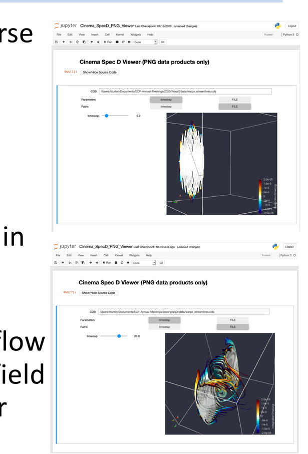

Nyx ⟷ AMReX ⟷ Ascent ⟷ VTK-m ⟷ In situ algorithms ⟷ post hoc viz

Goal: In situ data reduction and slice visualization with post hoc reconstruction

Data-driven sampling enables probabilistic identification of interesting regions in the data automatically, prioritizing important regions. Applied in situ to Nyx, important halo regions are preserved.

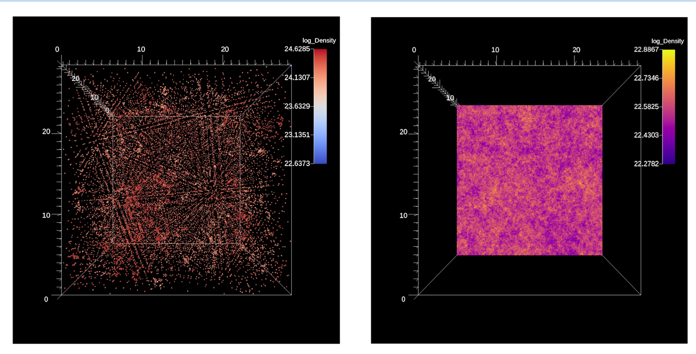

Left: data-driven sampling applied to Nyx data in situ. Right: slicing algorithm produces in situ visualization for post hoc exploration and validation.

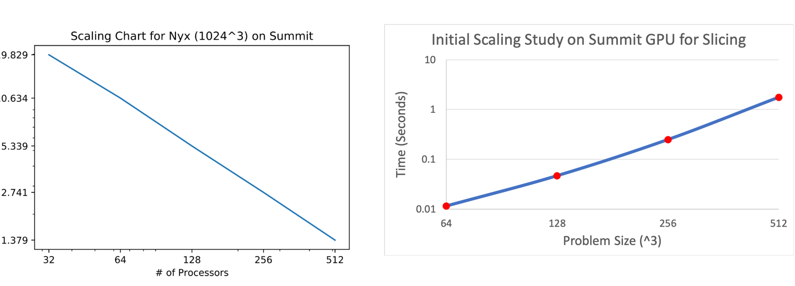

The Nyx <> Ascent <> sampling/slicing pipelines have been run on Summit GPUs with excellent scaling results. Left: sampling algorithm, right: slicing algorithm.

Los Alamos National Laboratory · Berkeley Lab · Lawrence Livermore National Laboratory · Oak Ridge National Laboratory · Sandia National Laboratories · Argonne National Laboratory · Kitware · University of Oregon · University of Utah · University of Leeds

National Nuclear Security Administration · U.S. Department of Energy · Office of Science

This material released under: SANDxxxx-xxxx, LA-UR-22-xxxx

This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001).Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

ECP EXASCALE COMPUTING PROJECT