News and Operational Highlights

OCTOBER 23, 2024

HYPERION ON VENADO

First performance results of next-generation Hyperion HED/ICF code on Venado

The first Hyperion performance results are in for the Venado supercomputer at LANL. Hyperion is a highly portable, next-generation, multiphysics high-energy-density-physics (HEDP) simulation code, currently targeting inertial confinement fusion (ICF) hohlraum modeling. A radiation-plasma code (as opposed to a standard rad-hydro code), Hyperion features kinetic plasmas of arbitrary collisionality and number of species (a unique capability in the complex), radiation transport, fast machine-learned non-local thermodynamic equilibrium (NLTE) atomic physics, and laser-ray tracing. Near-term plans are to add radiation transport, electromagnetic effects, and nonlinear laser-plasma interactions (LPI). Hyperion's long-term aim is to address complex-wide gaps in our ability to model hohlraums for ICF and HEDP applications and to further LANL's leadership in predictive modeling of HED experiments.

From inception, Hyperion targeted GPU computing, making it a first-tier Venado code for ICF multiphysics simulations. Hyperion is designed to embrace novel heterogenous architectures and harness present and future compute power for multiphysics simulations. The code is currently underpinned by sophisticated GPU-ready frameworks for multiphysics coupling (AMP) and particle discretizations (Cabana and Kokkos, both products of the Exascale Computing Project), which future-proof the code against hardware evolution.

Hyperion was recently ported to Venado and scaled up to half of the supercomputer (1024 GPUs). Early scaling results are very promising. Figures 1 and 2 depict weak- and strong-scaling parallel studies, demonstrating almost ideal scaling in up to 1024 GPUs and underscoring Hyperion's potential for exploiting modern exascale architectures towards cutting-edge multiphysics simulations of ICF hohlraums and other HEDP experiments. A cross-machine performance comparison (Figure 3) demonstrates significant performance gains (of 5 times and 16 times, respectively) of the Grace-Hopper NVIDIA architecture of VENADO versus Chicoma (another GPU-enabled cluster at LANL) and Rocinante (CPU only).

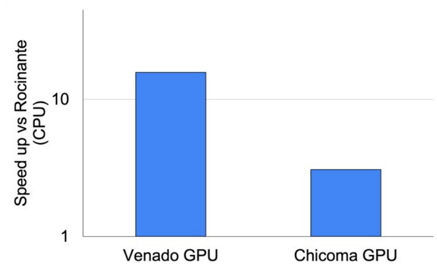

Figure 3 Cross-machine performance comparison of Hyperion. Study employed 16 nodes with 64 MPI ranks (4 ranks/node). GPU machines use 4 GPUs/node (1 rank = 1 GPU). Rocinante CPU uses 28 OMP threads/rank.

We have begun exercising Hyperion on ICF hohlraum exemplars using Venado. Future work will focus on deploying the NLTE module for the radiation-diffusion module, deploying adaptive-mesh refinement and novel particle load-balancing strategies for increased versatility in parallel domain decomposition, maturing the radiation-transport module, adding electromagnetic effects, and extending to 3D. Near-term applications will focus on advanced, high-performance hohlraums.

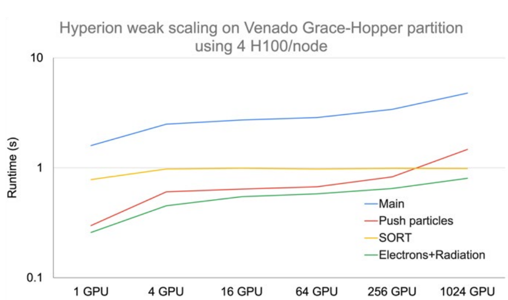

Figure 1 Weak-parallel-scaling study for Hyperion on Venado up to 1024 GPUs.

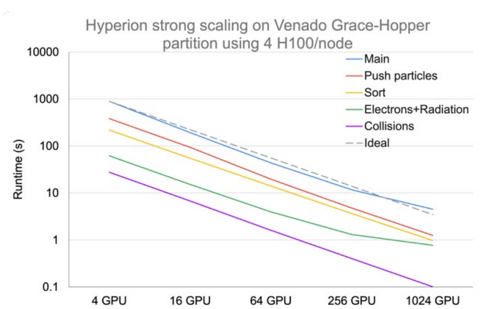

Figure 2 Strong-parallel-scaling study for Hyperion on Venado up to 1024 GPUs.

For more information contact **Jimmy Fung** (fung@lanl.gov), IC Program Manager. We acknowledge the work of L. Chacon, B. Philip, S. Anderson, J. Bonilla, G. Chen, C. Fontes, B. Gifford, L. Green, B. Haines, H. Hammer, A. Simakov, A. Stainer, and L. Yin.

RISTRA PROJECT

Towards Large-Scale Materials Performance Studies

Understanding how materials behave under high strain rates and other extreme conditions is essential to stewarding our aging stockpile. New manufacturing processes, such as additive manufacturing, impart crystallographic textures that differ from the textures of traditional processes like casting. These differences can change a material's response to strain, leading to alternative design decisions when either replacing old parts or designing new ones.

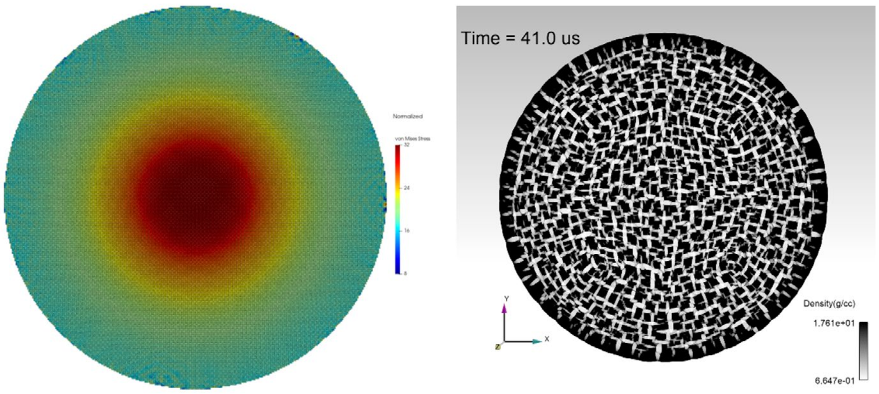

Figure 1 Left: Stress distribution in a slice through an expanding infinite cylinder using the particle-based DDMP method. Right: Density distribution showing fragmentation in a similarly sized cylinder with finite thickness of 5 mm with a finite volume method.

The physics governing the continuum level response of a material to extreme loading, including plastic deformation and damage, occur at the scale of the individual grain or smaller. This is inherently a multiscale problem. Material simulations that bridge multiple scales within ASC at LANL have historically been carried out on small samples (on the order of millimeters).

A goal of LANL's ASC/IC Ristra project is to provide HPC- and GPU-ready simulation capabilities to tackle these mesoscale and multiscale materials questions for the design community. Ristra's Moya code is being refocused to support a variety of models and discretization techniques to support an R&D framework for materials investigations at part-sized scales. This involves leveraging LANL's CSSE FleCSI framework for distributed parallelism coupled with Sandia National Laboratories' Kokkos framework for on-node parallelism like GPUs.

Figure 1 shows two variants of a radially expanding disk problem aimed at understanding fracture. The image on the left uses the dual domain particle method (DDMP) for an infinitely thick disk and portrays the distribution of stress in the disk; symmetry is preserved due to recent improvements in mesh-particle compatibility via the Local Stress Difference algorithm[2]. The image on the right shows a standard staggered-grid, Lagrangian finite volume discretization on a disk of U_6Nb with a finite thickness of 5 mm. This simulation used a Johnson-Cook strength and damage model with seeded porosity, which leads to disk fragmentation. Both methods used an unstructured mesh.

By supporting multiple discretization's, the Ristra project is well-suited to tackle a variety of materials problems in which modeling choices could be better suited to answer a given question. Although the examples shown use relatively established methods, Ristra is also investing in future methods. One avenue being explored is adding additional physics to the DiscoFlux model, which accounts for dislocation dynamics within polycrystalline materials.

For more information contact **Jimmy Fung** (fung@lanl.gov), IC Program Manager. We acknowledge the work of Kevin Larkin, Duan Zhang, Kyle Perez, Paul Barclay, and Jiajia Waters.

[1]Zhang, D.Z., Perez, K.A., Barclay, P.L. et al. Rapid particle generation from an STL file and related issues in the application of material point methods to complex objects. *Comp. Part. Mech.* **11**, 2291–2305 (2024). https://doi.org/10.1007/s40571-024-00813-z

[2]Perez, K.A., Barclay, P.L., & Zhang, D.Z. Nodal force error and its reduction for material point methods. *J. Comp. Phys.* **498**, 112681 (2024). https://doi.org/10.1016/j.jcp.2023.112681

MACHINE LEARNING INFORMED OPTIMIZATION FOR ICF

Improving Cylindrical Implosion Designs

The study of inertial confinement fusion (ICF) implosions advances our understanding of the science relevant to stockpile stewardship. In an ICF implosion, the rocket effect drives a spherical shell inward, compressing and heating a central deuterium-tritium (DT) gas to reach conditions relevant to runaway thermonuclear fusion. However, the features of an ICF target can induce hydrodynamic instability growth and impact implosion performance. As a result, validating and improving computational models of ICF implosion dynamics, coupled with experimental ICF data, strengthens confidence in our capability to predict performance.

Assessing hydrodynamic instability growth in standard spherical systems poses challenges. Imaging through a spherical shell requires analysis methods with potentially ambiguous interpretations, and cone-in-shell experiments are limited to low convergence. Therefore, a cylindrical platform that allows for direct diagnostic access while maintaining geometric convergence effects is needed. Previous cylindrical implosion experiments [1-2] have used annular plastic cylinders with a central low-density foam fill as targets and examined the growth of pre-imposed periodic features, using foam density to control the deceleration history and the final convergence of the target. Still, running high-fidelity simulations to fully capture the instability growth in these geometries remains computationally expensive, further complicating the design of these targets for experimental validation.

In a recent paper [3], LANL scientists leveraged supervised machine learning (ML) techniques, specifically Gaussian processes, Deep Jointly Informed Neural Networks, and their Bayesian counterpart, Bayesian Deep Jointly Informed Neural Networks to optimize cylindrical target design over a large parameter space. Initial efforts trained these surrogate models on a dataset generated from 1D runs of Los Alamos National Laboratory's Eulerian radiation-hydrodynamics code xRAGE simulations, providing low-fidelity predictions of thermonuclear yield as a function of various design parameters (e.g., ablator thickness, DT fuel density).

ICF hotspot conditions are better reproduced by replacing the central foam fill of the target with DT gas, allowing detailed studies on how mixing of shell material into the DT fuel impacts thermonuclear burn. Bayesian optimization was used to iteratively refine design parameters, selecting those that maximized yield while adhering to predefined constraints. When expanded to include laser pulse-shape parameters, as shown in Fig. 1, ML surrogate models identified a design that improved compression efficiency and increased neutron yield by 68% compared to a square-pulse laser drive.

The team is now integrating 2D simulations into a cost-aware, multi-fidelity optimization framework to more accurately account for instability effects, as past studies show that optimal designs in 1D do not always perform similarly in more complex simulations. Multi-fidelity optimization routines integrate data from various simulation fidelities, ensuring accurate performance prediction at a reduced computational expense.

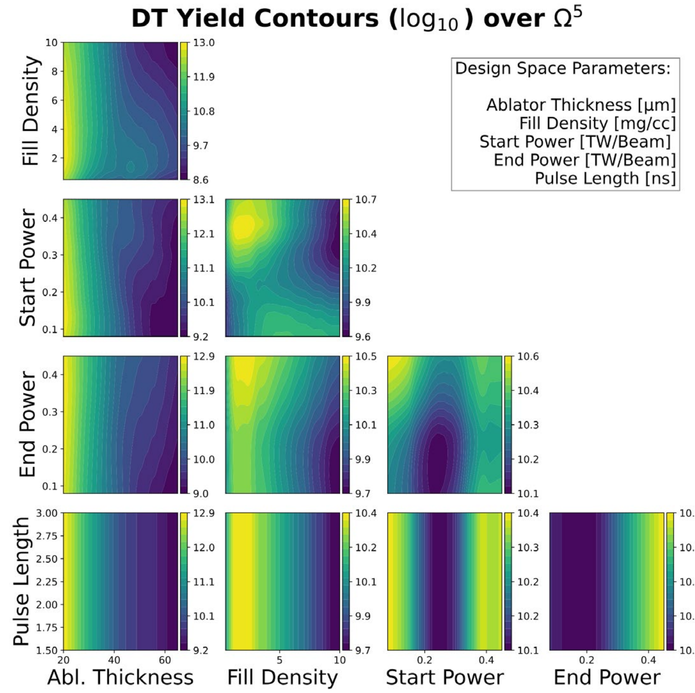

Figure 1 Contour plots of yield from an ML surrogate model over a design space (Ω5) that includes variations in target ablator thickness, target fill density, and laser pulse parameters (duration, initial, and final power). Our recent work focuses on how these factors affect target compression and yield, emphasizing that while higher power typically boosts yield, careful pulse shaping is critical to achieve efficient compression and minimize driver energy requirements

For more information contact **Gowri Srinivasan** (gowri@lanl.gov), V&V Program Manager. We acknowledge the technical contributions of William Gammel, Joshua Sauppe, Paul Bradley, and academic collaborator, Kevin Lin (University of Arizona).

[1] S. Palaniyappan, J. P. Sauppe, B. J. Tobias, C. F. Kawaguchi, K. A. Flippo, A. B. Zylstra, O. L. Landen, D. Shvarts, E. Malka, S. H. Batha, P. A. Bradley, E. N. Loomis, N. N. Vazirani, L. Kot, D. W. Schmidt, T. H. Day, R. Gonzales, and J. L. Kline, “Hydro-scaling of direct-drive cylindrical implosions at the OMEGA and the National Ignition Facility,” *Phys. Plasmas* 27, 042708 (2020).

[2] J. P. Sauppe, S. Palaniyappan, B. J. Tobias, J. L. Kline, K. A. Flippo, O. L. Landen, D. Shvarts, S. H. Batha, P. A. Bradley, E. N. Loomis, N. N. Vazirani, C. F. Kawaguchi, L. Kot, D. W. Schmidt, T. H. Day, A. B. Zylstra, and E. Malka, “Demonstration of scale-invariant Rayleigh-Taylor instability growth in laser-driven cylindrical implosion experiments,” *Phys. Rev. Lett.* 124, 185003 (2020).

[3] W. P. Gammel, J.P. Sauppe, P. Bradley, “A Gaussian process based surrogate approach for the optimization of cylindrical targets,” *Phys. Plasmas* 1 31, 072705 (2024).