Advanced Simulation and Computing

Los Alamos NATIONAL LABORATORY

News and Operational Highlights

APRIL 23, 2025

MULTI-PHYSICS CODE

Solving the heat conduction equation across contact surfaces in multi-material FLAG code

Multi-physics codes require simultaneously modeling different physical phenomena. In particular, ASC's multi-physics code FLAG needs to model heat conduction during contact between different materials. However, modeling heat conduction is difficult because of the way the code represents interacting materials.

In a computational domain consisting of several material regions that are in contact and can slide with respect to each other, the contact capability enables a tangential velocity discontinuity between two contacting bodies. The thin shear boundary layer can be approximated by a discontinuity, or slide surface. This enables the interaction of two disconnected meshes such that gaps between impacting solids can open and close, and therefore friction and boundary layer viscosity can be modeled. Figures 1 and 2 show real-world applications of contact capability.

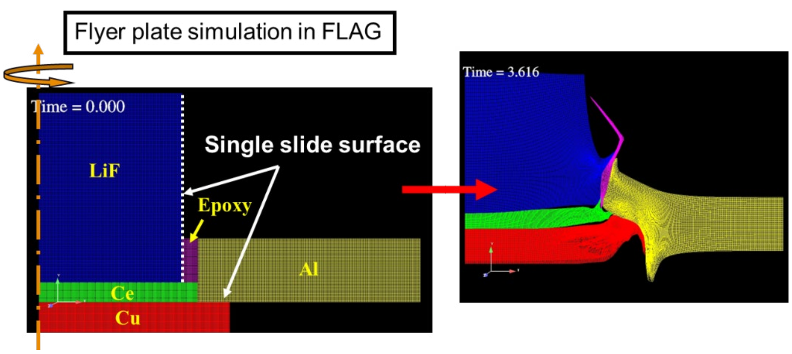

Figure 1 Flyer plate simulation. Left: initial configuration of the materials; right: configuration of the materials at final time.

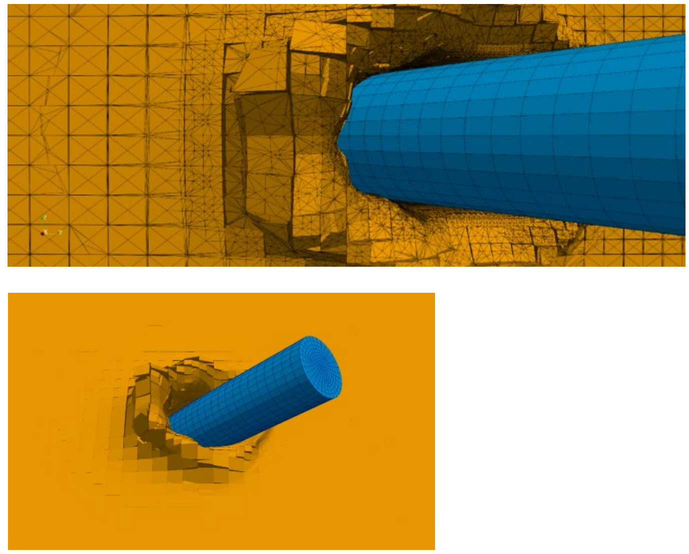

Figure 2 Penetration phenomena. Penetration of aluminum plate with ogive-nose steel projectile.

The main difficulty of modeling heat conduction is because the meshes representing different materials in contact can slightly overlap each other or have gaps (Figure 3).

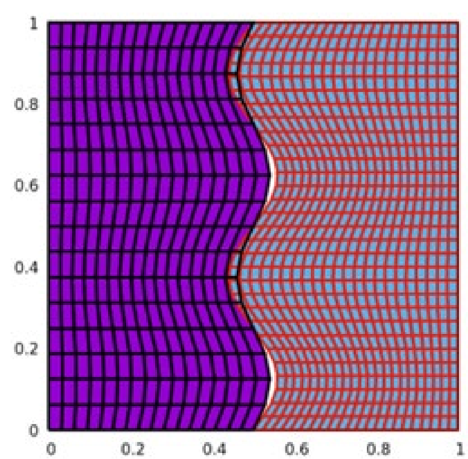

Figure 3 Illustration of typical contact interfaces.

(*continued on following page*)

The standard method for solving heat conduction assumes that there is only one conforming mesh covering the entire computational domain and that there are no gaps or overlaps. The classical formulation of the heat conduction problem for multiple regions can be described as follows:

$$\mathbf{div}\, k_i \,\mathbf{grad}\, T_i = f_i\,, \quad \mathbf{x} \in \Omega_i$$
$$T_i = T_j, \quad \mathbf{x} \in \Omega_i \cap \Omega_j$$
$$(\mathbf{n_i}, k_i \,\mathbf{grad}\, T_i) = -(\mathbf{n_j}, k_j \,\mathbf{grad}\, T_j)\,, \mathbf{x} \in \Omega_i \cap \Omega_j$$

where *Ti* is a temperature in region *Ωi*, *ki* is the conductivity coefficient, and *ni* is normal to the interface between materials. In each domain, temperature satisfies the elliptic equation, and, on the interface, one temperature and normal heat flux is continuous.

LANL developed a new algorithm to discretize the elliptic equation in each region. First, computational geometry methods are used to project interfaces from one mesh to another. Then the continuity of temperature and normal flux is enforced in a weak sense (using a least squares approach) to address gaps and overlaps between contacting surfaces.

To demonstrate this method, LANL solved the equations on the mesh of Figure 3 with two materials that have different heat conductivities: k_1 and k_2. Figure 4 shows the computed temperature field as a color map (left) and surface plot (right).

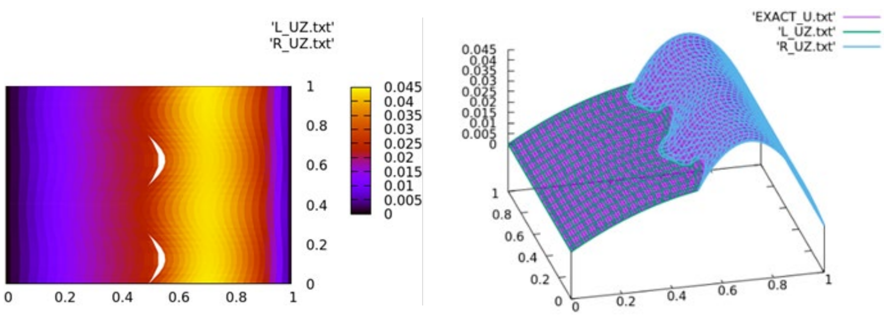

Figure 4 Left: A color map of the heat conduction solution. One can clearly see gaps; however, the color map is smooth, which indicates that the continuity conditions are well maintained. Right: Surface plots of the numerical solutions in each region are superimposed with the exact solution, demonstrating that the numerical solution is accurate.

Currently, we have developed discretizations in 3D and are working to find an optimal algorithm for the computational geometry required to project interfaces and for linear solvers. The first 3D results are presented in Figure 5.

For more information, contact **Jimmy Fung** (fung@lanl.gov), IC Program Manager. We acknowledge the work of Mikhail Shashkov, Konstantin Lipnikov, Mack Kenamond, Nathan Vaughn-Kukura, and support from the ASC NUMIC and RISTRA projects.

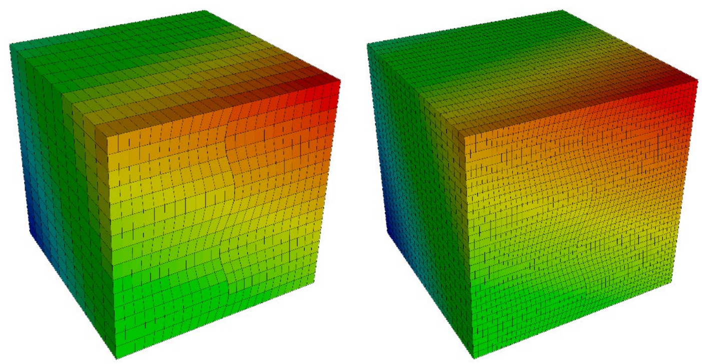

Figure 5 A color map for the solution of the heat conduction problem in 3D in low-resolution (left) and high-resolution (right). This demonstrates that continuity conditions are well maintained.

UNCERTAINTY QUANTIFICATION METHODS

Extending Bayesian Calibration for Improved Learning from Complex Models

In the realm of physical modeling, accurate prediction hinges on robust uncertainty quantification (UQ). Traditional UQ methods, while valuable, can fall short when dealing with complex scenarios where model parameters affect both the shape and timing of simulated curves. This limitation is a stumbling block in fields relying on time-dependent data, such as velocimetry in mission-critical applications.

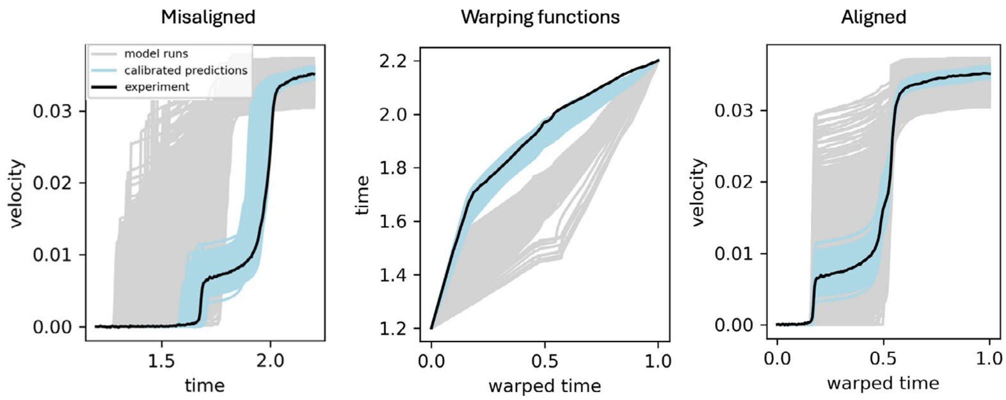

Figure 1 Misaligned flyer plate simulations can be decomposed into aligned flyer simulations and associated warping functions to improve surrogate and discrepancy modeling in Bayesian model calibration. In this case, a significant time-shift discrepancy was used.

To address this critical need, scientists from Los Alamos and Sandia National Laboratories have developed an innovative solution[1]. This new method tackles the challenge of misaligned simulation and measurement curves head-on, offering a sophisticated yet practical approach to Bayesian model calibration. The procedure separates misaligned curves into aligned curves and "warping functions," effectively disentangling shape and timing variability. By utilizing both aligned curves and transformed warping functions, the method achieves a more comprehensive and accurate calibration with the same data than traditional methods. This approach allows for a more nuanced understanding of discrepancies between model and reality since model form error structure can be elicited in a more natural space. Importantly, the method integrates seamlessly with established model calibration tools like Impala developed at LANL by ASC/PEM.

This new approach has already demonstrated its value in flyer plate (Figure 1) and Z machine velocimetry analyses, two critical areas in materials science and high-energy physics. By enabling more accurate learning from the same data set, the method represents an important step forward in our ability to extract meaningful insights from experimental results. The implications of this advancement extend beyond these initial applications. Fields dealing with time-dependent curve data stand to benefit, including materials science and engineering, high-energy physics experiments, seismology and geological studies, climate modeling and prediction, and aerospace engineering. By providing a more robust framework for handling misaligned curves, this method empowers researchers and engineers to learn more from their expensive data. As we continue to tackle increasingly complex challenges in science and engineering, tools like this will be instrumental in advancing our understanding and capabilities across a wide range of disciplines.

For more information, contact **Gowri Srinivasan** (gowri@lanl.gov), V&V Program Manager. We acknowledge the work of Devin Francom (LANL), Derek Tucker, Gabriel Huerta, Kurtis, Shuler, and Daniel Ries (Sandia). This work was also supported by the NA-22 and LDRD programs.

[1] Francom, D., Tucker, J. D., Huerta, G., Shuler, K., & Ries, D. (2025). Elastic Bayesian model calibration. *SIAM/ASA Journal on Uncertainty Quantification*, 13(1) 195-227

APPLICATION PERFORMANCE

Programmable Data Access Accelerator Tackles the Memory Wall in HPC

Nuclear security applications are distinguished by the complexity of coupled multi-physics and the space complexity of high resolution, many materials, and multiple-space representations within a single application. Current state-of-the-art simulations require 1 to 2 petabytes of system memory because of the large spatial domain coupled with the large dynamic range of spatial and state space. This space complexity requires the use of sparse data structures such as adaptive mesh refinement and sparse material representations. Physics algorithms adapted to these data structures result in indirect memory access patterns throughout the simulation.

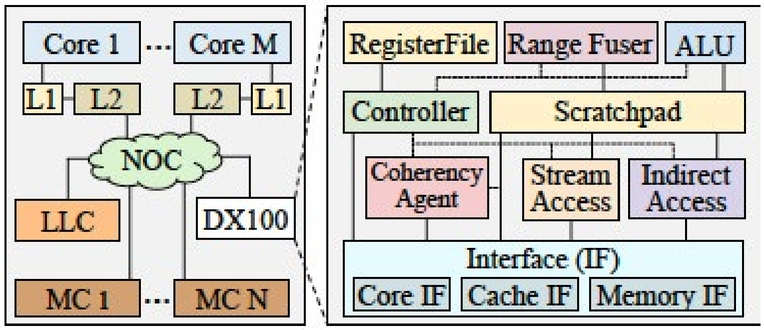

Figure 1 DX100 architecture.

With these workloads in mind, LANL and the University of Michigan have designed and prototyped DX100, a programmable data access accelerator that significantly improves application performance across a broad set of application workloads commonly used in simulation and data analysis. DX100 addresses the most pressing bottleneck exhibited by indirect memory access bound applications: the memory wall. Through a deep iterative codesign approach, DX100 delivers improved efficiency and performance for these workloads (Figure 1).

Through reordering, interleaving, and coalescing memory requests, DX100 can provide up to 17.8 times speedup in memory access performance. DX100 efficiently offloads indirect memory accesses and associated address calculation operations while simultaneously optimizing memory accesses to maximize the effective bandwidth of the memory system. Figure 2 illustrates the DX1000 reordering technique to improve row-buffer hit rate and therefore achieved memory bandwidth.

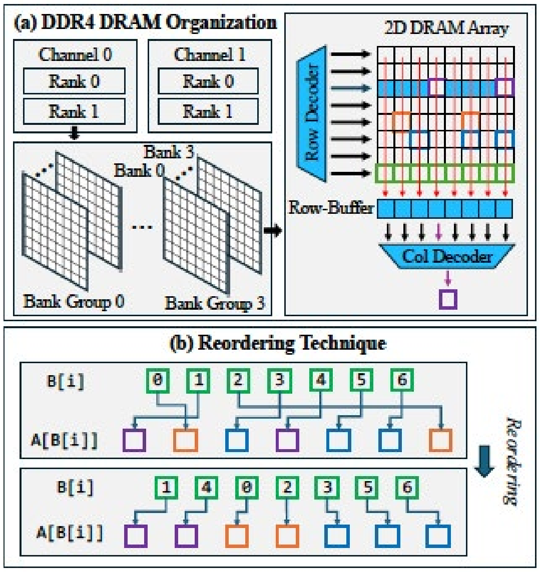

Figure 2 (a) DDR4 DRAM organization; (b) memory access techniques for reordering for improved row buffer hit rate.

To support the DX100 accelerator without significant programming efforts, we developed a set of Multi-Level Intermediate Representation (MLIR) compiler passes that automatically transform legacy code to use the DX100. Experimental evaluations on 12 benchmarks, spanning scientific computing, database, and graph applications, show that DX100 achieves performance improvements of 2.5 times over a multicore baseline and 1.9 times over a state-of-the-art indirect prefetcher.

(*continued on following page*)

Figure 3 illustrates the performance improvements achieved using DX100 across a diverse set of workloads, including access patterns commonly found in both the Eulerian Applications Project (EAP) and the Lagrangian Applications Project (LAP). Key access patterns and kernels found in LAP, including calculating the gradient of a zone-centered field at mesh points (GZP) and calculating the gradient of a zone-centered field at the zone centers (GZZ), achieve 5.3 times and 4.6 times improvement, respectively. Across a broad set of data analysis and modeling/simulation workloads, DX100 provides 2.5 times speedup (geomean) across these benchmarks.

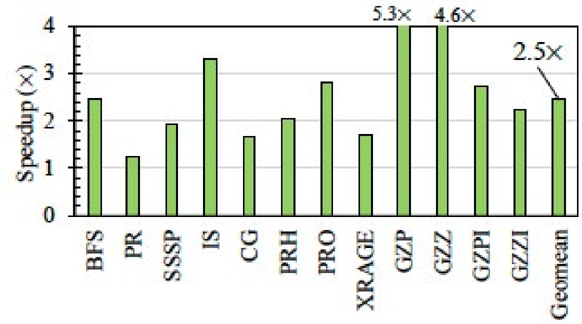

Figure 3 DX100 speedup for different workloads.

Hardware/software codesign is critical to the continued improvement in platforms and codes for the ASC program. The DX100 is an example of the significant improvements in application performance that can be achieved through this approach.

This work will be published in the premier forum for new ideas and experimental results in computer architecture: Khadem, Alireza (UMich) & Kamalavasan Kamalakkannan (LANL), et al. “DX100: A Programmable Data Access Accelerator for Indirection”, To Appear in the 2025 International Symposium on Computer Architecture.

For more information, contact **John Patchett** (patchett@lanl.gov), CSSE Program Manager. We acknowledge the work of our collaborators at the University of Michigan and Institute of Science Tokyo and LANL staff members Kamalavasan Kamalakkannan, Jered Dominguez-Trujillo and Galen Shipman.