

A MESSAGE FROM NEAMS LEADERSHIP

NEAMS program review and progress on FY26 milestones

We would like to start this newsletter by taking a step back and recapping the NEAMS Annual Review Meeting, held February 10–12, 2026, as a hybrid event at Idaho National Laboratory. A total of 235 people registered for the meeting, including 129 for in-person attendance. Participants represented the NEAMS program, other DOE-NE R&D programs, the nuclear industry, the Nuclear Regulatory Commission, and universities. We appreciate everyone's participation and contributions to making the meeting a productive event.

The first day and a half were open to external participants and featured reviews of fiscal year 25 (FY25) accomplishments and future plans, as well as stakeholder feedback through reactor-focused panels and sessions on AI/ML and multiphysics. The panel format fostered productive discussions and valuable feedback. The NEAMS program took note of feedback and developed actions across a range of topics. Examples include maturing capabilities that support light-water reactor power uprates in coordination with the [DOE UPRISE initiative](#); leveraging NEAMS multiphysics capabilities to accurately assess margins, which is a distinguishing feature of the simulation tools; providing transuranic fuel performance capabilities; enabling mechanistic source-term modeling; facilitating code adoption through validation, verification and quality assurance; improving models for inter-assembly heat transfer in metal-cooled fast reactors; benchmarking simulations against upcoming MARVEL microreactor test data; providing capabilities for simulating salt-spill accidents; and predicting fracture and material failure with assessed uncertainties. This feedback was used to update the NEAMS five-year research plan, including refining schedules for planned accomplishments to better serve stakeholder needs, and to inform work packages for FY27.

The remainder of the newsletter highlights selected milestones and accomplishments produced by our five Technical Areas (TAs) during the first three quarters of FY26. The highlights from Multiphysics Applications and Reactor Physics share a common theme: accelerating simulations by enabling codes to run efficiently on graphics processing units (GPUs), which are being rapidly deployed as part of the expansion of AI infrastructure. The resulting increase in computational efficiency can support faster analysis workflows and more efficient generation of synthetic datasets for AI applications. The Thermal Fluids article documents new capabilities for modeling air and moisture ingress in reactor systems, which are important for licensing High-Temperature Gas Reactors. The Fuel Performance article describes the SIFGRS fission gas release model and how it provides a basis for new models of transient fission gas release in UO₂, supporting burnup extension and power uprates of the light-water-reactor fleet. Finally, the Structural Materials and Chemistry highlight demonstrates a new capability for modeling crack propagation in welds during reactor operation.

Stakeholder feedback is critical to the success of the NEAMS program, and we are always looking for better ways to communicate our plans and progress and to solicit and receive input. If you have feedback or ideas about the program or how to best gather input from the community, please do not hesitate to reach out.

-David Andersson, David Henderson and Tanju Sofu

Point of contact: David Andersson (andersson@lanl.gov).

TECHNICAL AREA HIGHLIGHTS

Fuel Performance




Figure 1: Left) Comparison of calculated and measured fission gas release values for the 11 Super-Ramp PKA fuel rods analyzed using the SIFGRS model in BISON. Right) Propagation of atomic-scale uncertainty to grain-boundary diffusivity and fission gas release during recurring power transients. Blue shows posterior uncertainty; red shows the SIFGRS baseline. The figures are reproduced from the paper [1] and the report [2], respectively.

Advances in Fission Gas Modeling in BISON

Fission gas behavior plays a critical role in nuclear fuel performance, directly affecting fission gas release, internal rod pressure, fuel swelling, thermal conductivity, microstructural evolution, and fuel-cladding interactions in light-water reactor (LWR) and advanced reactor fuels. Accurate modeling of these phenomena is essential for assessing fuel safety and operating margins and supports efforts to extend fuel burnup and increase reactor power.

Comprehensive Assessment of the SIFGRS Model

The Simple Integrated Fission Gas Release and Swelling (SIFGRS) model was originally developed for the BISON fuel performance code based on

work by Pastore et al. To strengthen verification and validation and improve connections among lower-length-scale simulations, experiments, and engineering-scale modeling, P. C. Simon and other NEAMS researchers recently completed and published a comprehensive refactoring, review, and assessment of SIFGRS within BISON [1].

SIFGRS has a modular design that addresses three key mechanisms: intragranular processes, intergranular behavior, and fission gas release. Its flexible, mechanism-based architecture enables rapid updates as scientific understanding evolves and facilitates extension to additional fuel types.

An extensive assessment and validation suite demonstrates that SIFGRS accurately predicts fission gas behavior in UO_2 fuel under both normal operating conditions and transient scenarios in

LWRs (Figure 1). Notably, the model reproduces experimental data across a range of test conditions using a single, consistent parameter set, demonstrating its robustness and broad applicability to UO_2 fuel.

The model has also been extended beyond conventional UO_2 fuel to chromia-doped UO_2 , U_3Si_2 , UN, and mixed-oxide (MOX) fuels. Ongoing work is addressing high-burnup conditions and supporting the deployment of advanced nuclear technologies through continued validation, sensitivity analysis, and uncertainty quantification.

New Transient Fission Gas Release Model

Predicting fission gas behavior at very high burnup, where experimental data are limited, remains challenging for SIFGRS, particularly under rare transient or accident conditions. The NEAMS program is using multiscale modeling, sensitivity studies, uncertainty quantification, and continued validation to close these knowledge gaps. A NEAMS milestone completed this year focused on transient fission gas release in UO_2 , with particular emphasis on the cooler periphery of the pellet (where high burnup structure forms) and loss-of-coolant accident (LOCA) scenarios [2].

A key innovation of this work is the mechanistic modeling of grain-boundary microcracking. During temperature increases, nanometer- to micrometer-scale cracks can form between gas bubbles. These microcracks connect the bubbles, enabling transient gas release that increases rod pressure and may affect cladding rupture during an accident.

The new model incorporates gas-defect–bubble reactions that realistically limit xenon-to-vacancy ratios within bubbles. It also includes a sophisticated treatment of microcrack evolution that accounts for the competition between crack growth, through pressure increases during heating, and crack healing over longer timescales.

Initial validation against experimental data has produced promising results. Figure 1 illustrates how the new model can be integrated with uncertainty quantification methods and separate-effects testing to generate distributions for the quantities of interest. Additional validation and testing in multiphysics analyses are underway and are planned for completion by the end of calendar year 2026.

Technical points of contact: P.C Simon (PierreClement.Simon@inl.gov) and Michael Cooper (cooper_m@lanl.gov).

References and Additional Reading: 1. P.-C. A. Simon, K. A. Gamble, D. Schwen, L. H. Harbour, R. T. Sweet, I. Ferguson, L. K. Agesen, N. Capps, C. Matthews, M. W. D. Cooper, D. Andersson, and S. Novascone, "Documentation, Assessment, and Improvements of Fission Gas Modeling Capabilities in BISON." *Journal of Nuclear Materials*, vol. 630, Art. No. 156752, 2026, doi: 10.1016/j.jnucmat.2026.156752; and 2. M. W. D. Cooper, C. O. T. Galvin, P. C. Simon, A. Schneider, C. Matthews, and D. A. Andersson, "Complete Lower-Length-Scale Parameterization of a Mechanistic Transient Gas Release Model for UO_2 with Uncertainty Quantification," Tech. Rep. LA-UR-26-24315, Apr. 2026.

Structural Materials and Chemistry

Demonstrating a MOOSE-based XFEM Workflow for Crack Propagation in Welds

Nuclear power plant components operate in harsh environments that can lead to multiple degradation mechanisms, including crack initiation and growth due to fatigue, stress corrosion cracking (SCC), brittle fracture, and stress-relaxation cracking. Welds represent a region of heightened sensitivity to failure due to stresses induced by the welding process. A recent milestone in the NEAMS Structural Materials and Chemistry technical area demonstrated a MOOSE-based workflow that couples weld-process simulation with fracture propagation in a pressurized-water-reactor control




Figure 2: Simplified axisymmetric illustration of the J-groove weld geometry. Reproduced from [3].

rod drive mechanism (CRDM) J-groove weld (Figure 2). The workflow supports assessment of materials degradation and risk of failure during long-term nuclear power plant operation.

The analysis consists of three sequential stages: (1) a thermomechanical welding simulation in MOOSE with temperature-dependent plasticity, (2) cooldown to ambient conditions to establish the residual-stress field, and (3) restart into a simulation using the MOOSE extended finite element method (XFEM) module for three-dimensional crack propagation driven by SCC.

Several new capabilities and code enhancements were developed to enable the workflow. Prediction of subcritical crack growth in welds require accurate residual stresses. To enhance this capability in MOOSE improvements were introduced that include robust initialization of newly activated weld elements, localized variable projection through new libMesh interfaces that reduce complexity, a do-nothing reinitialization option, patch-recovery-based reinitialization to preserve mesh quality, and a birth-strain treatment that reconciles the constitutive reference configuration with the inherited deformed state. Together, these changes mitigate element distortion, negative Jacobians, and convergence failures while allowing the residual-stress field to be reconstructed with a single equilibrium solve.

On the fracture side, the three-dimensional XFEM cutting-mesh algorithm was enhanced to project crack-front nodes back onto free surfaces, accommodate growth into geometric corners, permit nonuniform crack-front advance, and avoid degenerate zero-area crack-face elements.

The improved workflow was tested on simple verification problems and then applied to the CRDM J-groove geometry for stress-corrosion-cracking conditions. An initial 4 mm × 8 mm semi-elliptical flaw was propagated using a 316L SCC growth-rate model driven by the reconstructed

three-dimensional hoop-stress field. The crack crossed the weld and entered the tube wall in approximately 24 years, breached the tube at about 35 years, and reached the weld–tube–vessel-head triple point at 36.5 years (Figure 2 and 3). The extent of crack advance was greatest where the reconstructed hoop stress was highest.

This work establishes an integrated MOOSE workflow for simulating subcritical crack growth in welded reactor components over service-relevant

time scales. It provides a foundation for fully three-dimensional welding analyses, more complex geometries, and interacting cracks, which will be addressed in future work. Although the present test case focused on a pressurized-water reactor application, the general capability has much broader applicability to advanced reactors [3].

Technical point of contact: Lynn Munday
(Lynn.Munday@inl.gov).




Figure 3: a) Hoop stress around the initial crack. b) The point after 24.3 years at which the crack reaches the tube. c) The crack breaching the tube after 35 years. d) The crack reaching the triple point after 36.5 years. Reproduced from [3].

References and Additional Reading: 3. L. Munday, B. Spencer, T. Hu, B. Barua, and C. H. Yang, "Integrated Simulation of Weld Residual Stress Evolution and Crack Propagation Using XFEM," Idaho National Laboratory, Idaho Falls, ID, USA, Tech. Rep. INL/RTP-26-93507, May 2026.

Reactor Physics




Figure 4: (a) Illustration of the performance benefits of the GPU based sweeper (calculation of the angular neutron flux) compared with the CPU implementation for several different iterative schemes available in Griffin, (b) the unscaled thermal scalar flux solution for the empire benchmark problem used to generate the timing results. Reproduced from [4].

New GPU-Based Solver in Griffin Speeds Up Neutron Transport Calculations

Advances in graphics processing units (GPUs) are reshaping high-performance computing and creating new opportunities to make detailed reactor physics calculations significantly faster. This is important for advanced reactors where simpler diffusion-based approximations for neutronics may have limited applicability, but heterogeneous transport methods have significant computational costs. Recent improvement to Griffin, the deterministic reactor physics code developed by NEAMS in the MOOSE framework, has been published in Nuclear Science and Engineering [4] where the potential of GPU acceleration to substantially reduce the computational cost of neutron transport simulations was demonstrated. This is a challenging task because the specialized

solvers used for heterogeneous neutron transport are not straightforward to parallelize.

Griffin now includes a GPU-accelerated sweeper for discrete ordinates (SN) transport (calculation of the angular neutron flux) with the discontinuous finite element method. Kokkos was selected as the programming model to provide portability across computing architectures while maintaining a single codebase. The new sweep algorithm uses a block-Jacobi approach and increases available parallelism by operating concurrently across neutron directions, mesh wavefronts, energy groups, and local degrees of freedom. A wavefront angular-flux storage strategy was also introduced to reduce GPU memory requirements and enable larger reactor problems to be solved efficiently.

The three-dimensional Empire microreactor benchmark, with a mesh containing approximately

3.2 million elements, was used as a test case. As shown in Figure 4, just four GPUs outperformed a calculation using 1,120 CPU cores by more than a factor of two in transport update time. The memory-efficient wavefront storage approach was essential for fitting the full three-dimensional problem onto the available GPUs. Despite the use of a block-Jacobi algorithm that lags angular-flux information across domain boundaries and can worsen convergence behavior, the resulting increase in iteration count was only marginal. This demonstrates that substantial GPU acceleration can be achieved without significantly degrading convergence behavior.

These results show that GPU computing could make high-fidelity SN neutron transport practical for routine advanced reactor analysis. The work also establishes an important foundation for broader GPU adoption across Griffin and the MOOSE ecosystem. Future development will focus on improved domain partitioning and sweep scheduling for larger distributed systems, adaptive thread-policy selection, improved performance for higher-order calculations, and ultimately porting remaining CPU-based operations, including the multiphysics calculations discussed in another article in this newsletter, to GPUs.

Technical point of contact: Namjae Choi
(namjae.choi@inl.gov)

References and Additional Reading: 4. N. Choi, J. Hanophy, Y. Wang, "Development of a GPU-Based Discrete Ordinates Transport Sweeper in the Griffin Reactor Physics Application," *Nuclear Science and Engineering*, Aug. 2026, doi: 10.1080/00295639.2026.2700698.

Thermal Fluids




Figure 5: Transient response during the simplified HTR-PM air-ingress event, left: air mole fraction during diffusion phase; right: core mass flow rate. Reproduced from [5].

Development of SAM Code Capabilities for Safety Analysis of GCR Air-ingress Events

SAM code capabilities have been extended to

close a critical capability gap for Gas-Cooled Reactors (GCRs): integrated, system-level simulation of air-ingress events following depressurized loss-of-forced-cooling (DLOFC).

The recent work [5] implements and validates three key capabilities that together enable end-to-end accident sequence modeling:

1. Multi-component gas mixture flow model: tracks helium/air composition and composition-dependent density changes that drive buoyancy and flow redistribution.
2. Graphite oxidation model: computes oxygen consumption and exothermic heat release at graphite surfaces.
3. Isentropic critical-flow model: captures choked depressurization through a break for realistic blowdown behavior.

The SAM enhancements were validated against two benchmark experiments. In the NSTF heavy-gas ingress test, SAM reproduced the rapid changes in inlet composition and the resulting suppression of natural circulation when dense argon entered the inlet, followed by the measured recovery of buoyancy-driven flow—demonstrating that the multi-component flow model captures composition-dependent density effects at the system level. In the NACOK graphite oxidation benchmark, SAM predicted the experimentally

observed strong axial nonuniformity and geometry dependence of oxidation and yielded a bottom-level graphite mass loss of about 25% (measured $\approx 24\%$), delivering results comparable to established system codes (SPECTRA and TINTE).

An integrated reactor-scale demonstration of a DLOFC air-ingress transient was exercised for a simplified HTR-PM model. The calculation reproduced the full accident progression: rapid blowdown from 7 MPa to atmospheric pressure, density- and diffusion-driven air ingress over ~ 15 hours while the loop remained hydraulically stalled, the subsequent onset of buoyancy-driven natural circulation that draws fresh air through the hot core, exothermic graphite oxidation with a peak fuel temperature around 62 hours, and eventual passive cooldown as vessel heat rejection overtakes heat sources. This multi-day simulation shows SAM can capture the end-to-end physics of HTGR air-ingress events in a single run. This work establishes SAM as a preliminarily validated, modern tool for HTGR air-ingress safety analysis and paves the way for more detailed, licensing-grade studies.

Technical point of contact: Rui Hu (rhu@anl.gov).

References and Additional Reading: 5. G. Yang, L. Zou, T. Mui, T. Hua, and R. Hu, "Development of SAM Code Capabilities for Safety Analysis of GCR Air-ingress Events," Argonne National Laboratory, Lemont, IL, USA, June 2026.

Multiphysics Applications

MOOSE Accelerates Multiphysics Simulation on GPUs

Modern Department of Energy supercomputers derive much of their computing power from graphics processing units (GPUs), creating both an opportunity and a challenge for scientific software. With support from the Nuclear Energy Advanced Modeling and Simulation (NEAMS) program, the MOOSE acceleration development team is adapting the framework for these systems while

preserving the flexible, modular approach that has made MOOSE a foundation for advanced reactor simulation.

The team is pursuing two complementary paths. The first uses the Kokkos performance-portability library to move selected MOOSE calculations onto GPUs while retaining the established MOOSE and libMesh programming model. Kokkos hides many hardware-specific details, allowing the same implementation to target multiple processor

architectures. This Kokkos-based capability joined the production MOOSE branch at the end of fiscal year 2025, and its coverage now includes major object families used for finite element calculations, boundary conditions, materials, auxiliary calculations, user objects, postprocessors, reporters, and functions.

Fiscal year 2026 development expanded that foundation in several important directions [6].

Linear finite volume calculations can now execute through Kokkos for diffusion, first-order upwind advection, volumetric sources, and Dirichlet and Neumann boundary conditions. These finite volume systems can coexist and interact with finite element systems. This interoperability lets model developers apply the most appropriate




Figure 6: Evolution of the calculated temperature field in a microreactor heat-conduction transient from startup to 40,000 seconds. For this nearly 1.4-million-degree-of-freedom problem, one NVIDIA B200 GPU reduced total runtime from 415.4 seconds on 112 Intel Xeon Platinum 8480+ CPU cores to 64.9 seconds, a speedup of 6.4. Reproduced from [6].

algebraic discretization to each physical process and its associated mathematical operator. The team also made MOOSE's forward-mode automatic differentiation capability available inside Kokkos device kernels, reducing the effort needed to construct consistent Jacobians for nonlinear problems.

A transient heat-conduction model of a microreactor demonstrates the performance potential. The model contains 190,744 elements and nearly 1.4 million degrees of freedom and follows the temperature field through 61 time steps. On this problem, one NVIDIA B200 GPU completed

the calculation in 64.9 seconds, compared with 415.4 seconds on two Intel Xeon Platinum 8480+ processors using all 112 CPU cores—an overall speedup of 6.4. Residual evaluation was 10.1 times faster and Jacobian assembly was 19.2 times faster on the GPU. The linear solve achieved a smaller speedup of 3.9, identifying solver and preconditioner improvements as an important next target.

The second path integrates MFEM, an independently accelerator-aware finite element library, as an alternative backend within MOOSE. Work led by the United Kingdom Atomic Energy

Authority (UKAEA) produced new bidirectional transfers that allow libMesh- and MFEM-based applications to exchange scalar solution fields even when their meshes and parallel decompositions differ. This means existing MOOSE physics can be coupled to MFEM-based physics without rewriting an entire application. Although the transfer operation currently stages data through the CPU, the participating application solves can remain GPU enabled.

Together, the Kokkos and MFEM approaches

provide a practical route to GPU acceleration: applications can move selected, computationally intensive physics to accelerators while continuing to use MOOSE input files, execution control, and multiphysics coupling. Future work will broaden the available finite volume and finite element capabilities, improve solver performance, reduce host-device data movement, and extend testing across additional GPU architectures.

Technical point of contact:

Alexander Lindsay (alexander.lindsay@inl.gov).

References and Additional Reading: 6. N. Choi, A. Lindsay, R. Chowdhury, R. Stogner, and C. Permann, "MOOSE GPU Acceleration," Idaho National Laboratory, Idaho Falls, ID, USA, Tech. Rep. INL/RPT-26-94239, 2026.

LOOKING AHEAD

Upcoming meetings

- NEAMS Annual Review Meeting, February 2–4, 2027, Oak Ridge National Laboratory and virtual.

Upcoming training

- SAM Users Meeting and Training Workshop, November 3–5, 2026, Argonne National Laboratory.

The Nuclear Energy Advanced Modeling and Simulation (NEAMS) program is a U.S. Department of Energy-Office of Nuclear Energy (DOE-NE) program developing advanced modeling and simulation tools and capabilities to accelerate the deployment of advanced nuclear energy technologies, including Light Water Reactors (LWRs), non-LWRs, and advanced fuels. We work with DOE-NE, the U.S. Nuclear Regulatory Commission (NRC), and industry to develop, demonstrate, and deploy usable advanced modeling and simulation capabilities across five technical areas: Fuel Performance, Reactor Physics, Structural Materials and Chemistry, Thermal Fluids, and Multiphysics Applications.

The NEAMS program has major work at Argonne National Laboratory, Idaho National Laboratory, Los Alamos National Laboratory, and Oak Ridge National Laboratory.

Learn more about the NEAMS program at neams.inl.gov