ASC Accomplishments
May 23, 2016

Unified collision operator demonstrated for both radiation transport and PIC-DSMC.
A side-by-side comparison between the DSMC method and the radiation transport method was conducted for photon attenuation in the atmosphere over 2 kilometers in physical distance with a reduction of photon density of six orders of magnitude. Both DSMC and traditional radiation transport agreed with theory to two digits. This indicates that PIC-DSMC operators can be unified with the radiation transport collision operators into a single code base and that physics kernels can remain unique to the actual collision pairs. This simulation example provides an initial validation of the unified collision theory approach that will later be implemented into EMPIRE. (POC: Chris Moore chmoore@sandia.gov) **[Brief: Sandia researchers have shown that two different numerical methods give consistent results for radiation transport in the upper atmosphere.]**

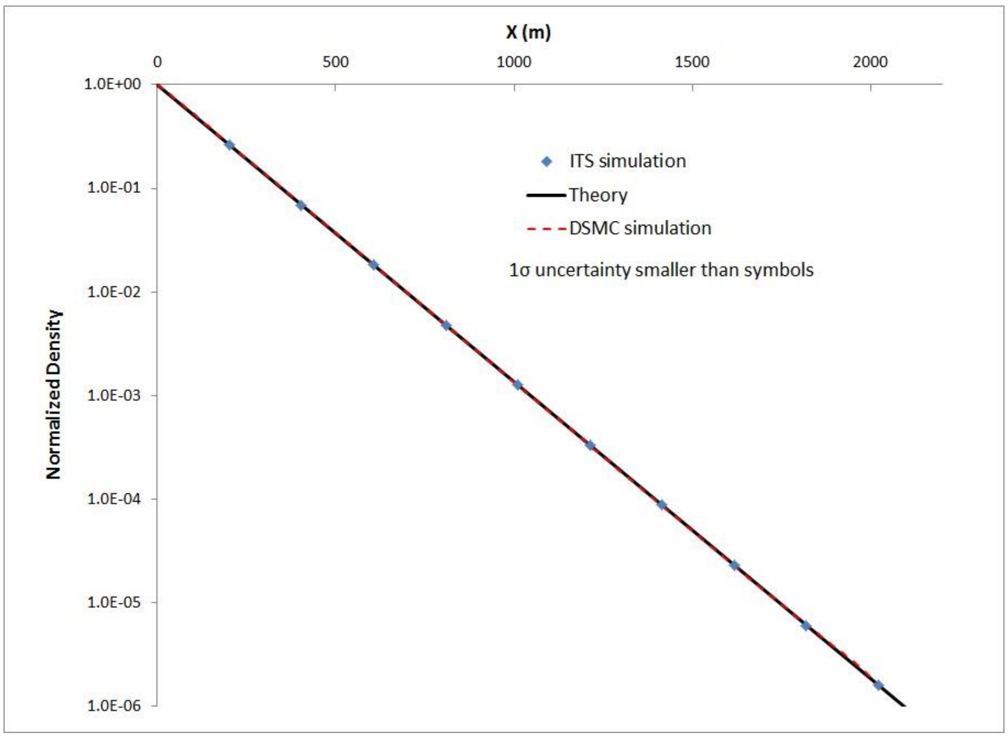

Generalized unstructured low-Mach turbulent flow simulations run during Trinity Open Science Period. Researchers at Sandia National Laboratories have tested the low-Mach code Nalu during the Trinity Open Science period by predicting aerodynamic efficiency of wind turbine blades using high fidelity Wall-Resolved Large Eddy Simulation (WRLES). The open-source application code base Nalu, which leverages the ASC infrastructure including Trilinos and the Sierra Toolkit (STK), is being used during the Trinity Open Science campaign in support of wind energy aerodynamic studies. High fidelity LES simulations for a wind turbine blade section at a high angle of attack (15^{o}) are being conducted on unstructured mesh sizes ranging from 1.6 to 13+ billion elements using core counts in excess of 60,000.

Although the project has specific science objectives to gain a fundamental understanding of complex flow characteristics known as stall cells, the first foundational objective was to ensure that the underlying ASC infrastructure was capable of supporting unstructured fully implicit, low-Mach turbulent wind energy calculations at this novel scale of production computing. Key modules are being tested for NNSA applications, including STK_transfer, mesh-file compression, and algebraic multigrid. Results thus far have demonstrated good solver efficiency (low linear iteration counts) for the difficult pressure Poisson system even at core counts of 64k (Figure 2). (POC: Stefan P. Domino, spdomin@sandia.gov) **[Brief: Researchers at Sandia are deploying the ASC code infrastructure to support high fidelity, low-Mach WRLES in support of Trinity Open Science.]**

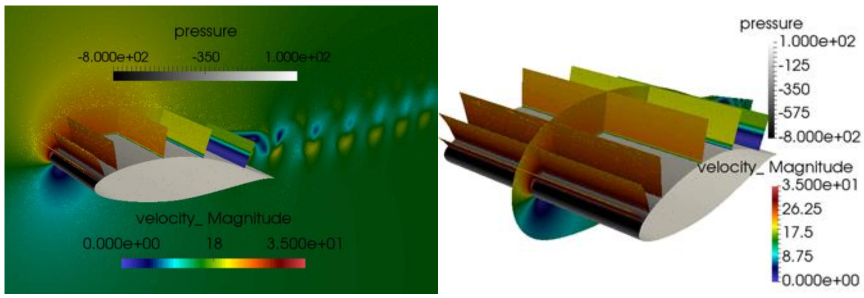

Figure 1: Pressure field from a RANS simulation as an initial condition to the LES simulation (left) and i/o planes (right).

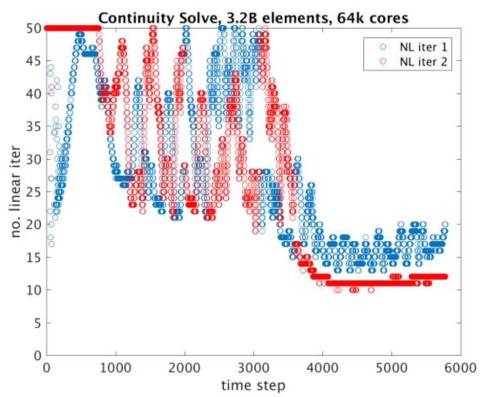

Figure 2. Trilinos/Muelu AMG linear iteration counts.

In-situ visualization on the Mutrino Trinity Platform. Sandia visualization researchers have collaborated with engineers from the Sandia Sierra Toolkit-based Nalu low Mach flow code team to integrate the ParaView Catalyst in-situ visualization library with the Nalu code on the Trinity platform. In-situ visualization allows parallel simulation codes to perform post-processing operations during run time in order to avoid lengthy I/O operations. Nalu simulation-runs with Catalyst image output have been conducted on the Trinity Platform development system Mutrino on up to 80 nodes (2560 cores). These experiments are being used to refine the in-situ visualizations in preparation for larger scale runs on the Trinity machine during the Trinity Phase 1 Open Science Period using 1900 nodes (60800 cores). Ron Oldfield **[Brief: Sandia researchers have enabled in-situ visualization for the Nalu low Mach flow code using ParaView Catalyst on the Mutrino Trinity platform in support of LANL Trinity Phase Science time.]**

ASC Accomplishments
May 29, 2016

Verification and Sensitivity Analysis for Prediction of Process-Induced Distortion. Analysts at Sandia National Laboratories are developing methods to simulate process-induced stresses in simple composite structures, such as flat plates and open-ended cylinders. These stresses form due to differences in the composite materials' coefficients of thermal expansion as well as the shrinkage upon cure exhibited by most thermoset polymer matrix materials. The consideration of these stresses is important when designing composite parts as process-induced stresses of significant magnitude can lead to fracture and failure within the composite. Recently, the developed methods were applied during the mesh convergence studies of three different composite structures, a flat strip, a flat plate, and an open-ended cylinder. The outcome of these mesh convergence studies have indicated the ideal mesh size and model discretization appropriate for future simulations of process induced stresses. Then, sensitivity studies were performed with the convergent meshes for each of the three structures to determine which model parameters are most critical to predictions of process-induced stresses and will therefore require the most rigorous experimental characterization. These critical parameters were the curing temperature and the composite in-plane elastic moduli and coefficients of thermal expansion. (POC: Stacy Nelson, smnelso@sandia.gov) **[Brief: Analysts at Sandia National Laboratories have performed verification and sensitivity analyses to accurately predict manufacturing process-induced stresses in composite structures.]**

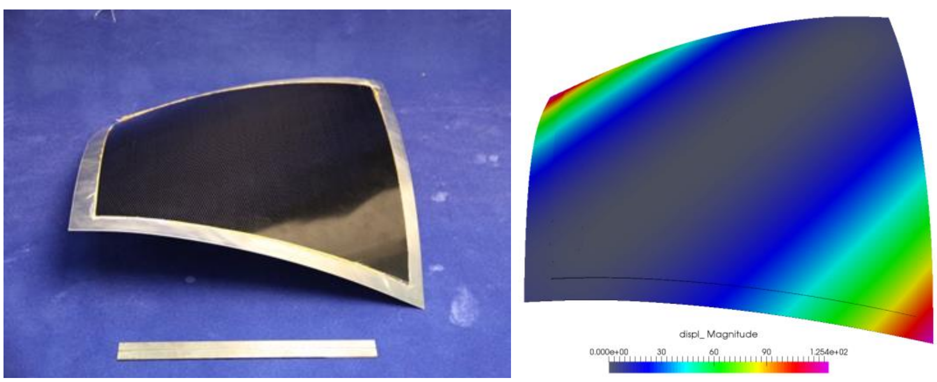

Figure: A composite/aluminum plate as manufactured (left) and as simulated (right). The process-induced stresses manifest physically as the warped shape.

Transient Free-Surface Flows. In current P&EM work at Sandia National Laboratories, researchers have developed and prototyped a preconditioner for the system of equations that are produced by the conformal decomposition finite element methods (CDFEM). Previous work either used expensive general-purpose solvers or modified CDFEM to eliminate the small angle simplex elements. These techniques, however, were not applicable for large scale parallel problems with many, complex interfaces and boundaries. By noting the relationship between CDFEM and hierarchical enrichment, researchers discovered that a relatively simple change of variables allows the system of equations to be recast into a system that is easily preconditioned using simple Jacobi preconditioners. This work has the promise to lead to more robust implementations of CDFEM for complex and dynamic interface problems with complex interfacial multi-physics. (POC: David Noble, drnoble@sandia.gov) **[Brief: Researchers at**

Sandia have discovered a preconditioning technique for the conformal decomposition finite element method, enabling them to provide their customers with a robust method for complex and dynamic moving interface problems including laser welding, organic material decomposition, and foam encapsulation modeling.]

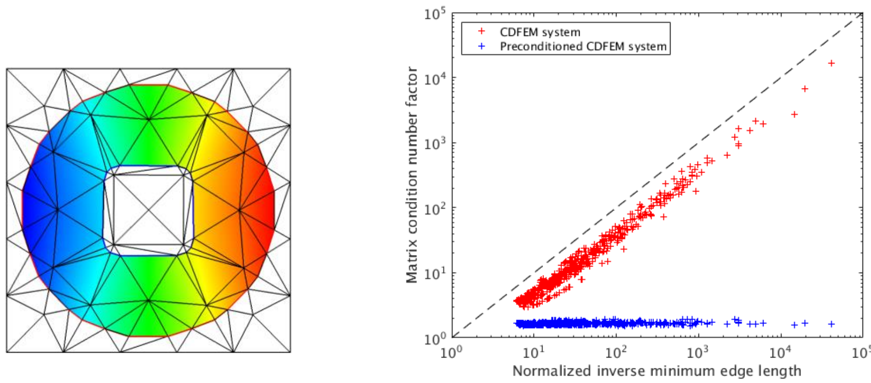

Figures: The plot on the left shows the kinds of decomposed meshes produced by the conformal decomposition finite element method (CDFEM). The small angle triangles in this mesh typically leads to linear system matrices that are difficult to solve (large condition numbers) as shown on the right (in red). A new preconditioner has been prototyped that produces well conditioned matrices regardless of the small angles present in the decomposed elements (shown in blue).

Intel "Dungeon" for Sandia Researchers Preparation for Trinity Phase 2. A "dungeon session" was held at Intel to do an in depth exploration of performance improvements on the upcoming Trinity Phase 2 hardware, Intel Xeon Phi processors. Both on-node parallelism and data movement were analyzed to determine how to extract best performance from the new hardware. This has led to ongoing collaborations with Intel to improve compilers and libraries in support of our code base. (POC: Si Hammond, shammon@sandia.gov) **[Brief: An Intel "Dungeon" Session was held for Sandia Researchers to Prepare ASC codes for Trinity Phase 2.]**

ASC Accomplishments
June 6, 2016

Improved convergence acceleration demonstrated in SCEPTRE. Researchers at Sandia National Laboratories have demonstrated improved convergence acceleration in the deterministic coupled electron/photon radiation transport code SCEPTRE (Sandia's Computational Engine for Particle Transport for Radiation Effects). The traditional source iteration is notoriously slow for diffusive systems (e.g. electron transport). A multi-level method (Transport Synthetic Acceleration [TSA]) has been recently implemented in SCEPTRE which uses a coarse angular solution to accelerate the fine angular problem. Good results have been demonstrated for a cable SGEMP (system-generated electro-magnetic pulse) test problem, which should dramatically improve run-time. Some challenges remain to be addressed. For example, S_N TSA algorithms require Galerkin mapping between moments and discrete spaces to avoid contamination of the correction term (this is now completed). Also, further acceleration of the coarse solves are needed since it is time consuming for large problems or problems with strong heterogeneities. And some additional code restructuring could take advantage of code re-use. (POC: Clif Drumm crdrumm@sandia.gov) **[Brief: Improved convergence accelerated was demonstrated in the SCEPTRE code for a cable SGEMP test problem]**

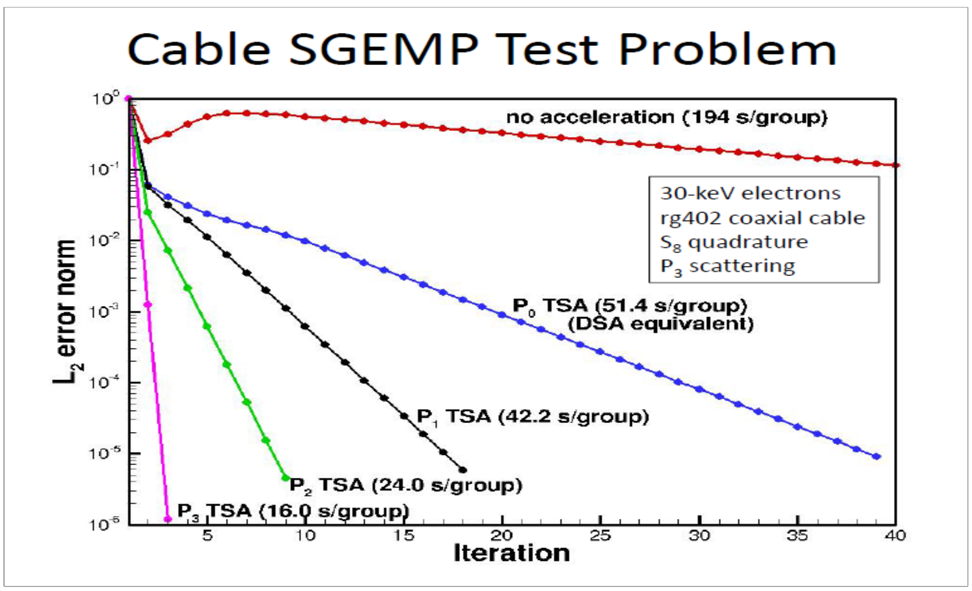

Figure: This plot shows the improvement of iteration convergence with increased order of the Transport Synthetic Acceleration (TSA) recently implemented in SCEPTRE.

Sandia researchers release first version of DARMA portability layer for next generation programming models. Sandia researchers have released the first version the DARMA portability layer for Asynchronous Many-Task (AMT) runtime systems. AMT runtimes are of great interest in the context of next generation platforms because they show promise to mitigate complexity at the runtime system-level. However, a recent ATDM study found that existing AMT runtimes have deficiencies in their APIs and need hardening prior to being ready for production use by ASC codes. The DARMA portability seeks to 1) insulate ASC applications from runtime system and machine architectures, 2) address deficiencies in existing AMT APIs, and 3) provide a formal specification of runtime system requirements. Within Sandia's ATDM program, the portability layer is being co-designed with application and runtime system developers that span the high-performance computing software stack. While the team is focused on meeting ATDM application requirements, their approach facilitates a more general exploration and characterization of runtime design space tradeoffs. These broader community-based activities are an important precursor to establishing community and vendor supported best practices and ultimately standards that will enable AMT technology long term. (POC: Janine Bennett, jcbenne@sandia.gov). **[Brief: Sandia researchers have released the first version the DARMA, portability layer for Asynchronous Many-Task (AMT) runtime systems.]**

Challenge: Addressing the Challenge of Communication-free Parallel Mesh Refinement
Inter-process communication is often the limiting factor for scalable parallel computing. To execute complex numerical simulations, the ATDM project currently adopts offline coarse mesh generation and parallel inline mesh refinement. However, existing mesh refinement methods lack scalability, as they require communication between processors to share local indices of new mesh entities as they get created. Since ATDM's mesh refinement is global and completely deterministic, each processor can predict the refinement operations over shared mesh entities across processors, and consequently, conclude the local indices of the refined mesh entities on neighbor processors without communication. Capitalizing on this observation, we designed a novel parallel mesh refinement approach that has been tailored for the ATDM needs and requires no communication. Implementation and testing of this new approach will begin in the coming weeks. (POC: Mohamed Ebeida msebeid@sandia.gov **) [Challenge Summary: A novel communication-free parallel mesh refinement prototype approach is being designed to address ATDM's challenge of highly scalable mesh refinement.]**

ASC Accomplishments
June 13, 2016

High-order compressible DNS capability. Researchers at Sandia National Laboratories have developed a novel capability in the Sierra Mechanics simulation package for performing direct numerical simulations (DNS) of turbulent flows using high-order unstructured numerical methods. This accomplishment facilitates fundamental studies to evaluate the simulation efficacy of lower fidelity models in both canonical and unstructured configurations. This work leverages previous work to use high-order unstructured methods to perform turbulence simulations using lower fidelity turbulence models. The recent work extends the initial capability to DNS through novel turbulent boundary conditions and a more efficient implementation for this class of problems. This capability is being further extended to perform two-way coupled fluid-structure interaction (FSI) simulations to study panel flutter. (POC: Travis Fisher, tcfishe@sandia.gov) **[Brief: Researchers at Sandia have developed a capability to perform compressible direct numerical simulations of turbulent flows using high-order unstructured numerical methods.]**

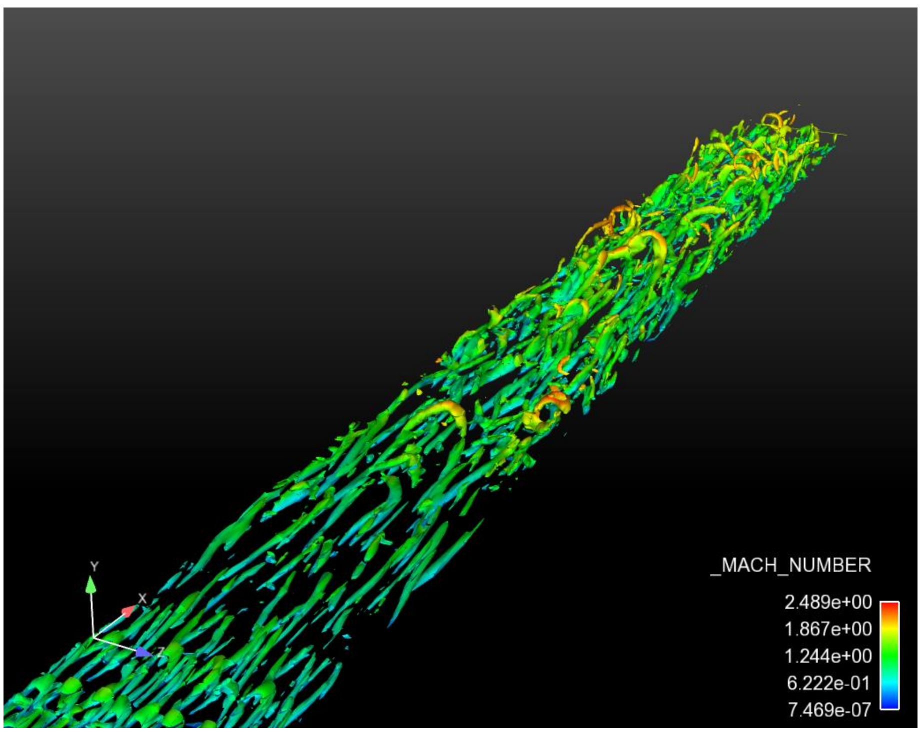

Figure: The breakdown of turbulent structures is visualized for a Mach 2 turbulent flat plate boundary layer. The turbulent structures are colored by Mach number.

Kokkos AMT Capabilities Integrated into AgileComponents Assembly. Sandia's discretization tools team has successfully integrated the experimental Kokkos Asynchronous Many-Task (AMT) capability into the AgileComponents assembly tools. Part of this work was to determine an integration pathway with a minimum disturbance to the current applications. Assembly on a node can now be performed with both task and data parallelism. This will allow us to explore possible efficiency gains from AMT.

Upcoming changes to the Kokkos AMT layer to support CUDA AMT will require more changes in the next quarter. Once finished, a study on various workset configurations will be performed for a representative physics. (POC: Roger Pawlowski, rppawlo@sandia.gov) **[Brief: Sandia's discretization tools team has successfully integrated the experimental Kokkos AMT capability into the AgileComponent assembly tools for on-node hybrid task and data parallel assembly.]**

ASC Accomplishments
June 20, 2016

Improvements to small-signal noise analysis in Xyce. The ADMS (Automatic Device Model Synthesizer) back-end for Xyce has been extended to support noise analysis. Noise analysis is a fundamental feature of most circuit simulators, and the numerical framework for this analysis was added to Xyce last year. However, in addition to the numerical framework, to fully support this capability it is also necessary to implement noise models in all the relevant compact models in Xyce. Until recently, only a handful of devices were supported. This has been greatly mitigated by the ADMS enhancements, which now allow for all the ADMS-generated devices to support noise. As all modern transistor models are implemented in Xyce using this mechanism, this accounts for the majority of models. Models which support noise in Xyce now include: VBIC 1.2, VBIC 1.3 (3-terminal, VBIC 1.3 (4-termina), MEXTRAM 504.11, MEXTRAM 504.11 with self-heating, BSIM6 v6.1.0, BSIM-CMG v 107.0.0 and the PSP. (POC: Eric Keiter,erkeite@sandia.gov) **[Brief: Noise analysis has been extended to support all transistors models in Xyce.]**

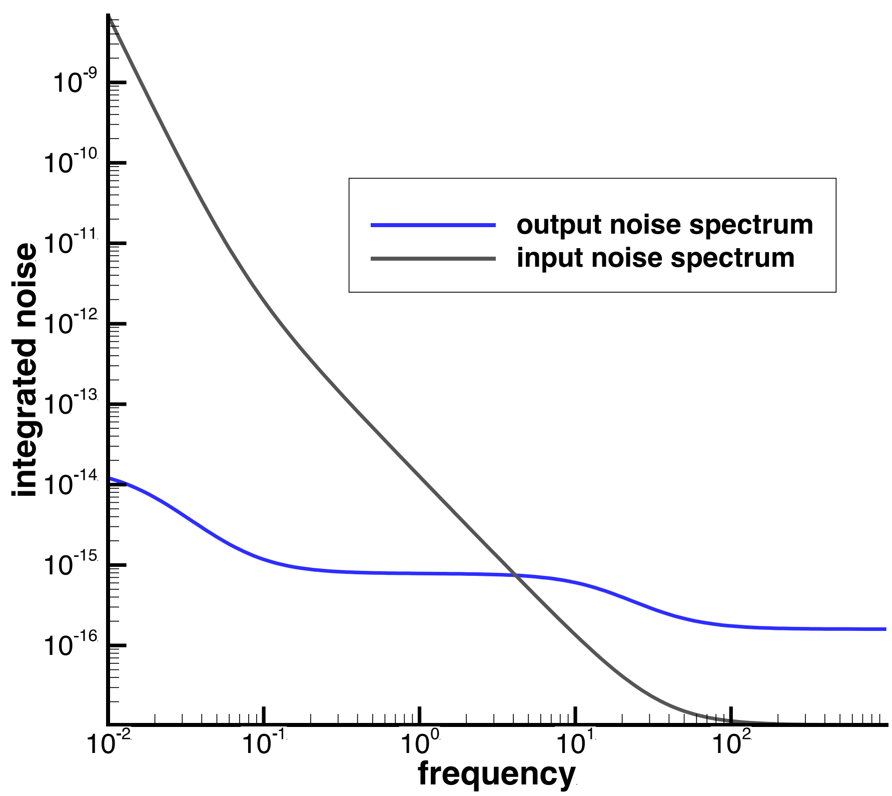

Figure: This plot shows the total output and input noise for a bipolar transistor in a common emitter amplifier circuit. The model used in this calculation is the VBIC 1.3 model. (VBIC = Vertical Bipolar Intercompany Model).

NiMC: Network-Induced Memory Contention.

Researchers at Sandia National Laboratories have introduced the concept of Network-Induced Memory Contention (NiMC) and demonstrated its impact for a variety of current HPC system architectures. It has been shown that NiMC is a concern for both onloaded and offloaded networking hardware, with the onloaded hardware observing the largest performance impact of up to 3X performance degradation for LAMMPS at 8,192 cores when injected RDMA (remote direct memory access) traffic creates contention at endpoints. The Sandia team has offered three solutions to NiMC, each applicable for different intensities of RDMA traffic, that can reduce the overheads associated with memory contention to only 6.4% versus the 300% or more originally observed. (POC: Ron Brightwell, rbbrigh@sandia.gov) **[Brief: Researchers at Sandia have been the first to investigate the impact of RDMA networks on application memory bandwidth on modern systems.]**

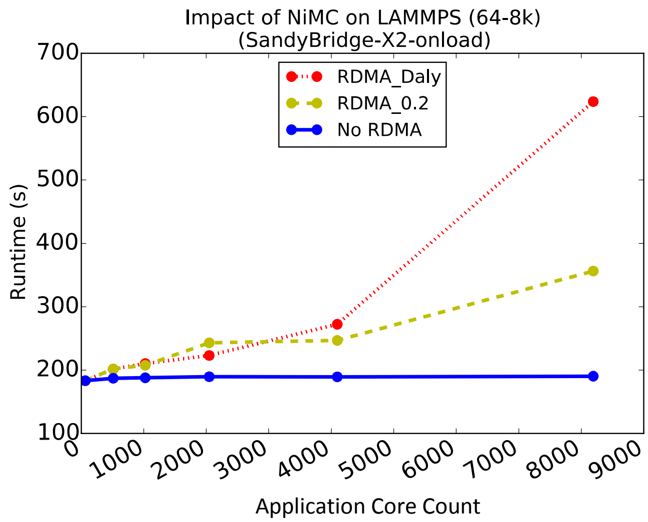

Figure: Impact of Network-Induced Memory Contention on LAMMPS at up to 8,192 cores

ASC Accomplishments
June 27, 2016

The Role of Synchronization in Scheduling Resilience Activities.
Researchers at Sandia National Laboratories have devised a low overhead resilience scheduling scheme that does not have the strict synchronization requirements of the current state-of-the-art. Current coordinated resilience schemes like checkpoint/restart minimize application performance perturbation, but require very expensive synchronization across all nodes. Uncoordinated methods like uncoordinated checkpoint restart, on the other hand, do not require expensive synchronization, but can significantly perturb application performance leading to significant slowdowns. This work demonstrates the meager synchronization requirements needed to keep overheads low when scheduling resilience activities. For example, a global synchronization scheme with a standard deviation of 1 millisecond has nearly the same overheads of a perfectly coordinated scheme, while having the benefit of being easily implemented on current and future extreme-scale systems. Overall, this work is the first in the field to demonstrate how to combine the performance overhead benefits of a coordinated approach with the ease of implementation of an uncoordinated approach. Robert Clay, rlclay@sandia.gov. **[Brief: Sandia researchers have devised a low overhead resilience scheduling scheme that does not have the strict synchronization requirements of the current state-of-the-art.]**

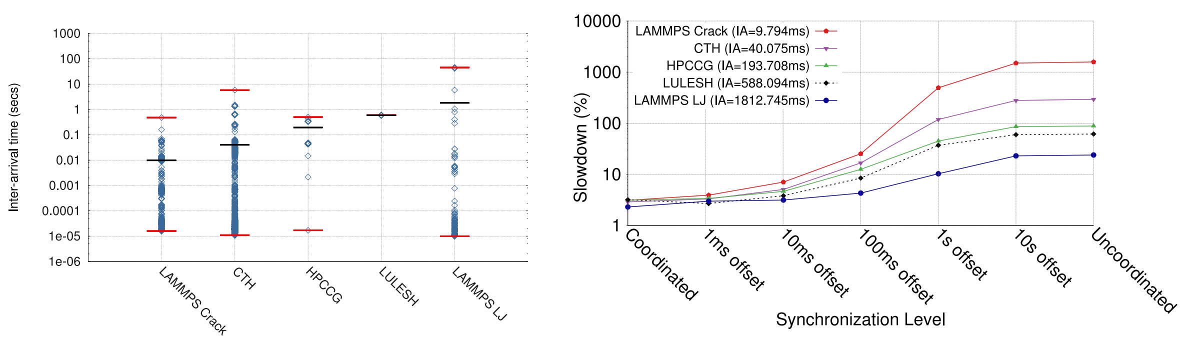

Figure: Measured min, mean, and max collective inter-arrival times for studied applications (left). Application slowdown at 64K nodes for resilience scheme with varying synchronization levels (right)

Improved Reduced Order Models in Sierra/TF. Researchers at Sandia National Laboratories have improved the support for reduced order models in Sierra/TF. This improves the ability of analysts to model pressurization of systems in abnormal thermal environments to support safety analyses. It is now possible to couple lower fidelity models for organic and energetic material decomposition that neglect gas transport with higher fidelity models that include gas transport to compute the pressurization of a full system model. Robustness of the advective bar capability for analyzing performance in normal thermal environments has also been improved by adding new regression testing to address gaps that were identified. Finally, a super-element formulation has been identified for future implementation into Sierra/TF in order to better model component thermal response in full system models at reduced computational cost. (POC: Victor Brunini, vebruni@sandia.gov) **[Brief: Researchers at Sandia have**

developed improved capabilities to couple high and low fidelity models of pressurization in abnormal thermal environments using Sierra.]

Challenge: Enable ATDM's Data Warehouse for broader ASC impact. Sandia's ATDM project is developing common input/output interfaces for our in-memory data management capability (called the Data Warehouse) that is creating opportunities to have near-term impact on other ASC sub-programs. In particular, our implementation of the Sierra IO System (IOSS) is generating interest from the Integrated Codes (IC) and the Computational Systems and Software Environments (CSSE) as an enabler for high-performance application workflows that transfer intermediate application results to downstream workflow components through the network, rather than the file system – significantly improving end-to-end run time. The CSSE Analysis and ATDM Data Management teams are working with ATDM team members to create the first (simple) integrated workflow using the data warehouse as part of an FY16 Level 2 milestone focused on demonstrating analysis capabilities. This early demonstration will provide a good test of our capability and will accelerate our progress toward our FY17 ATDM milestone and our ability to contribute to the Sandia Analysis Workbench. (POC: Craig Ulmer, cdulmer@sandia.gov) [**Challenge Summary**: Common input/output interfaces promise to improved end-to-end runtime performance.]

ASC Accomplishments
July 4, 2016

Sierra Toolkit thread-ready MPI load balance. A new thread-ready MPI load balance capability has been released by the Sierra toolkit and Trilinos Zoltan teams at Sandia National Labs. The teams have collaborated to deploy this capability for us in all NW Sierra engineering mechanics analyses. This new load balance capability supports a 'constrained load balance' algorithm that considers the computational cost of calculating particle physics *and* the computational cost of transferring data between the background fluid grid. (POC: Manoj Bhardwaj, mkbhard@sandia.gov ; Karen Devine, kddevin@sandia.gov) **[Brief: A thread-ready MPI load balance capability has been released by Sierra for use in production NW simulations.]**

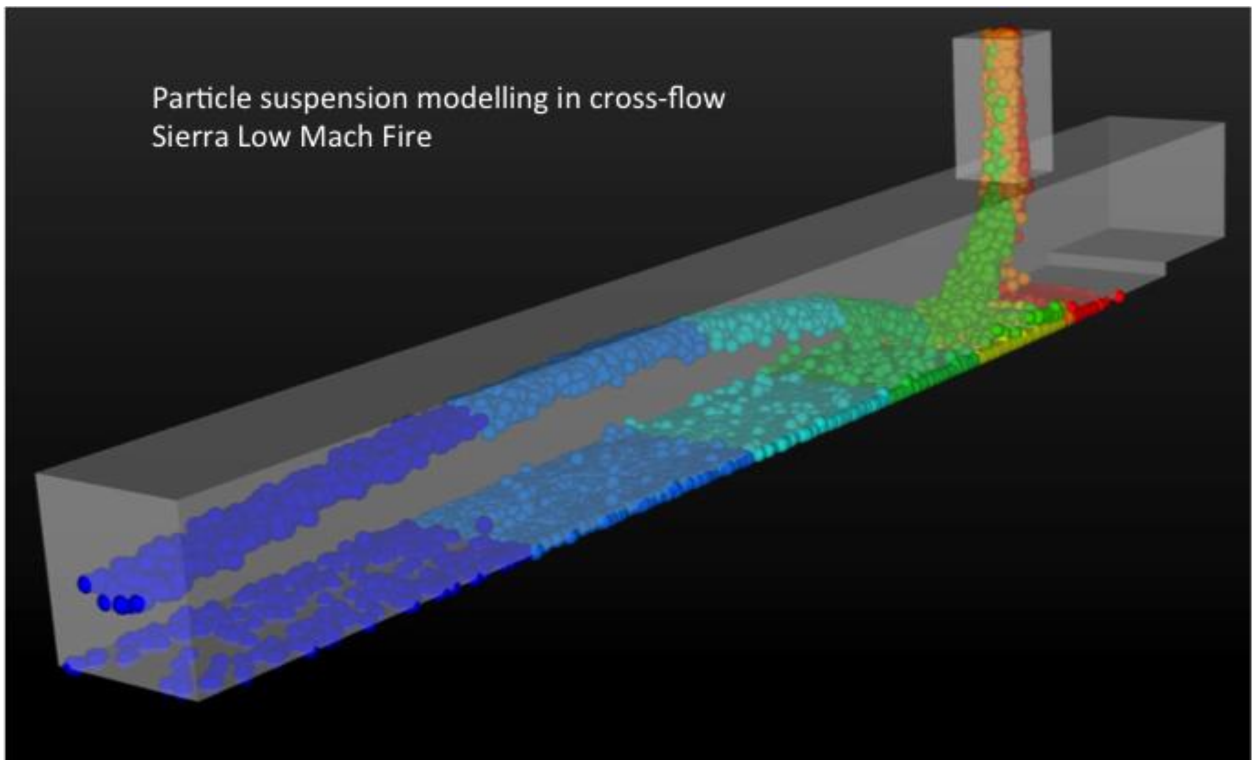

Figure: Sierra load balance capability applied to a multi-physics Sierra low Mach Fire problem. The figure depicts the processor ID that particles are sent to that result in load balanced computational work.

Python Interfaces added to Power API. Sandia and Cray have worked together to add Python interfaces to the HPC Power API. This will enable system administrators and other users who prefer to program in Python to access the capabilities provided by the Power API. The Power API was originally specified as a C API, but Cray and members of the HPC community provided feedback that Python interfaces would be beneficial for their intended use cases. The Python interfaces to the Power API were developed as part of the Trinity program's Advanced Power Management Non-recurring Engineering project. Cray will be deploying a Python version of the Power API on the upcoming Trinity platform. (POC: Jim Laros, jhlaros@sandia.gov) **[Brief: Sandia and Cray have worked together to add Python interfaces to the Power API.]**

Challenge: Bayesian Decision Modeling of Engineering Test Plans. A new model for how engineering programs can assess the importance of testing/experimental and computational modeling/simulation activities using Bayesian Decision Networks (BDN) is being developed. Our goal is to enable engineers and decision makers to better understand their critical activity needs and prioritize work to most efficiently meet a program's requirements. The BDN enables existing engineering knowledge to be quantitatively incorporated through priors (inputs). Design options are modeled as decisions informed by prior information and testing/modeling results. Asserting requirements are met is represented as diagnostic activities dependent on the state of the system. The result is that engineering activity results and expert judgment can be blended to make improved decisions on testing/modeling needs, design, and qualification (POC: J. Templeton, jatempl@sandia.gov) **[Challenge Summary: Researchers at Sandia are developing a Bayesian Decision Network model, which quantitatively incorporates engineering activity results and expert judgment, to assist planning and execution of large engineering programs.]**

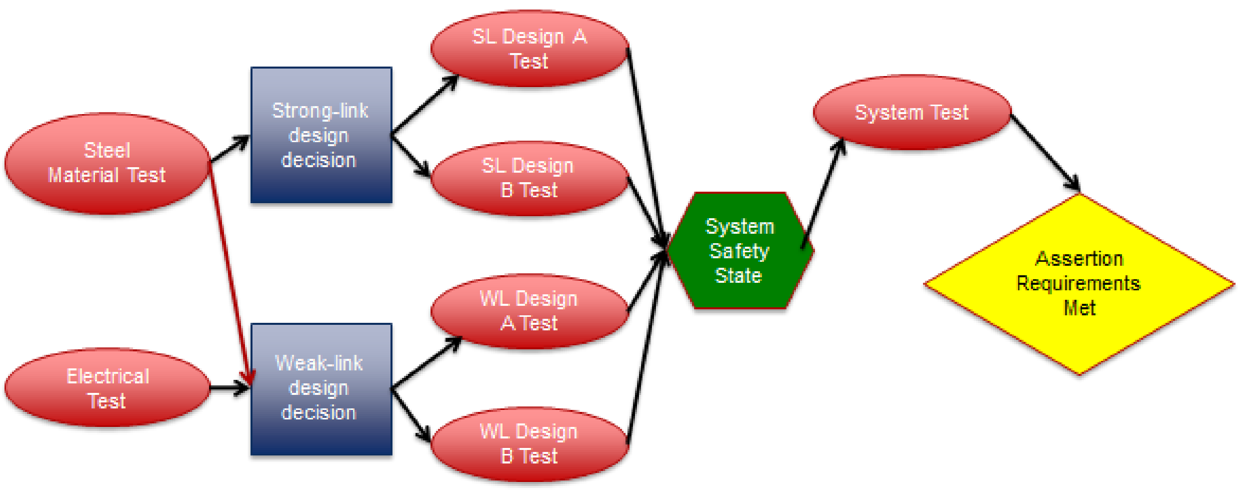

Figure: A simple schematic for representing system safety within a Bayesian decision network.

Red ovals represent test activities which add information and generate posterior distributions from priors. Blue boxes are design decisions selecting between different design options. Green hexagons are included to represent the state of a system, need in order to incorporate downstream diagnostics to assess the state. Finally, yellow diamonds represent the outcome (utility) of the activities, in this case being able to make a credible assertion that the system's safety requirements are met.

ASC Accomplishments
July 11, 2016

Advanced Reduced Order Modeling for Structural Mechanics. Parameterized reduced order models (PROMs) have been successfully applied to reentry vehicle-like geometries, where the thicknesses of sections of the cone are varied independently. Sensitivity information is output from Sierra/SD and used to generate a parameterized model of the eigenvalues and system matrices. The PROM formulation allows for design studies to be performed using only a single model evaluation in SIERRA. By contrast, using conventional techniques, new model geometries, the associated meshes, and evaluations in SIERRA are necessary for every design point considered. In Figure 1, the thickness of the aft rib (where the internal assemblies are attached to the aeroshell) is varied and the change in eigenvalue (which is related to the natural frequency of the structure) is calculated for three different vibrational modes, including two radial and one torsional. The PROM approach is shown to have very high agreement with the results from the high fidelity simulations (SIERRA), but requiring **approximately 1/40th of the computational time**. This novel approach is being demonstrated to significantly improve our ability for design optimization. (POC: Matthew Brake, mrbrake@sandia.gov) **[Brief: Novel reduced order modeling theories are being applied to representative applications at Sandia in order to demonstrate ability to enable design optimization.]**

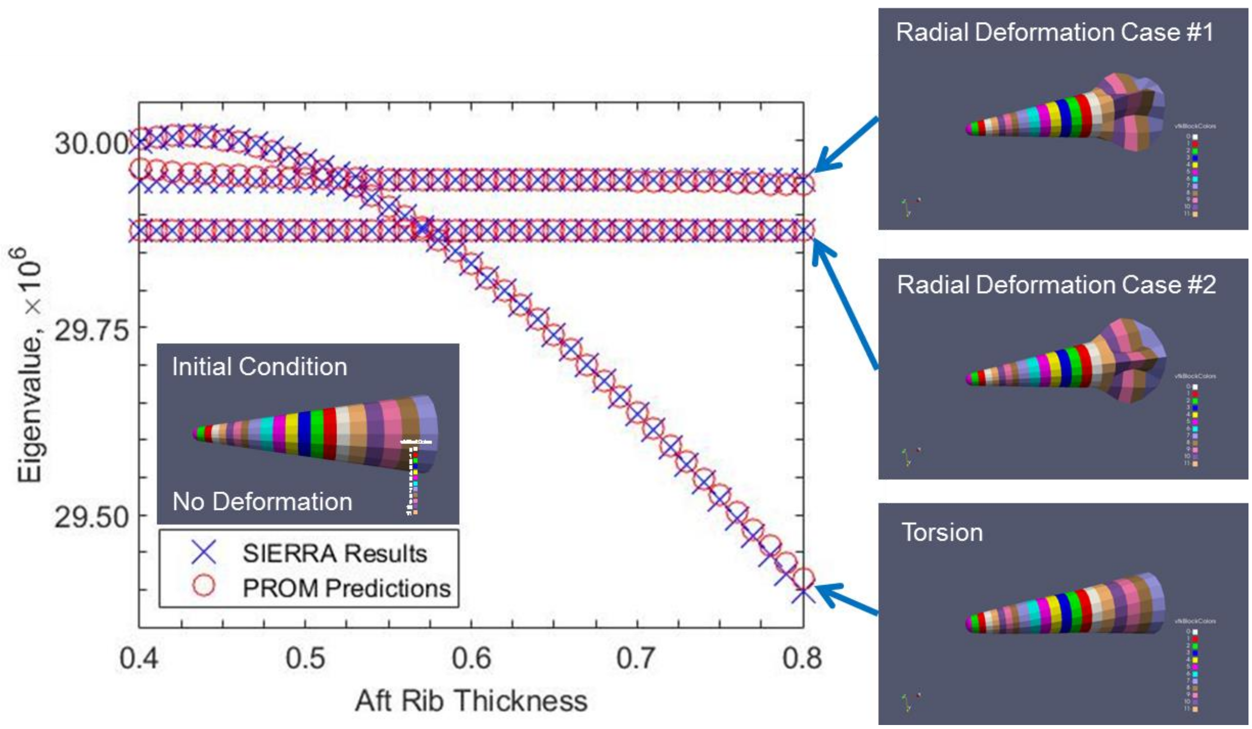

Figure: Eigenvalues of three different vibrational modes as a function of aft rib thickness for a parameter sweep using 41 different models evaluated in SIERRA (**X**), and only one PROM model constructed from a single SIERRA evaluation (**O**). The color bands on the nose cone indicate different regions of parameterized geometry.

High-performance topology optimization-based design for additive manufacturing. Researchers at Sandia National Laboratories are leveraging Sandia's Sierra production physics code, high performance computing, and the Sandia Analysis Workbench to deliver a powerful topology optimization-based design environment. Topology optimization uses computational simulation to iterate on the form of a design given design functional requirements defined by the user. The Sandia Analysis Workbench provides a professional front-end for setting up functional requirements and then launching topology optimization design runs that utilize Sierra physics code on powerful HPC platforms. This combination enables an environment for fast design iterations based on functional requirements. Version 1.0 of this environment, called "Plato", will be released for government use at the end of Q2. This initial release uses Sierra Structural Dynamics in support of doing designs with linear statics structural functional requirements. Future work will concentrate on leveraging more of the Sierra suite of physics code to enable designs with non-linear mechanical functional requirements. (POC: Brett W. Clark, bwclark@sandia.gov) **[Brief: Researchers at Sandia National Laboratories are leveraging Sandia's Sierra production physics code, high performance computing, and the Sandia Analysis Workbench to deliver a powerful topology optimization-based design environment.]**

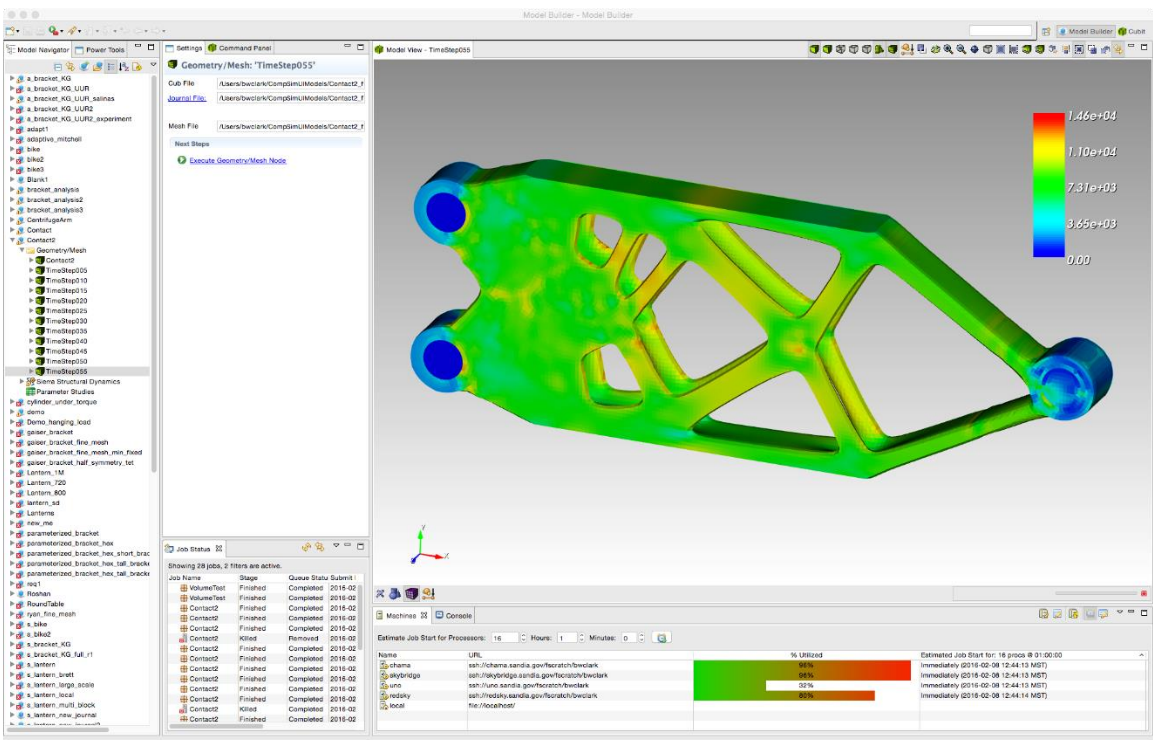

Figure: The "Plato" topology optimization-based design environment showing an optimized design with stresses plotted.

ASC Accomplishments
July 18, 2016

NuGET Free-Field Blast Environment Capability. The NuGET code team at Sandia has completed the implementation of the DNA Nuclear Blast Standard (1 KT) which allows users to calculate the blast overpressure and dynamic pressure from a free air nuclear burst. The new capability will eventually allow users to balance the hostile radiation requirements for NW strategic systems with the axial and lateral G-loading requirements on the re-entry system aeroshell. Consideration of the blast environments has the potential to reduce the radiation qualification test levels (or increase experiment margin) for systems currently in the LEP and ALT processes. (POC: K. Russell DePriest, Org. 01384 krdepri@sandia.gov) **[Brief: The NuGET code team has developed a model based on the DNA Nuclear Blast Standard which has the potential to reduce qualification test levels or increase experimental test margins for strategic systems in the modernizing stockpile.]**

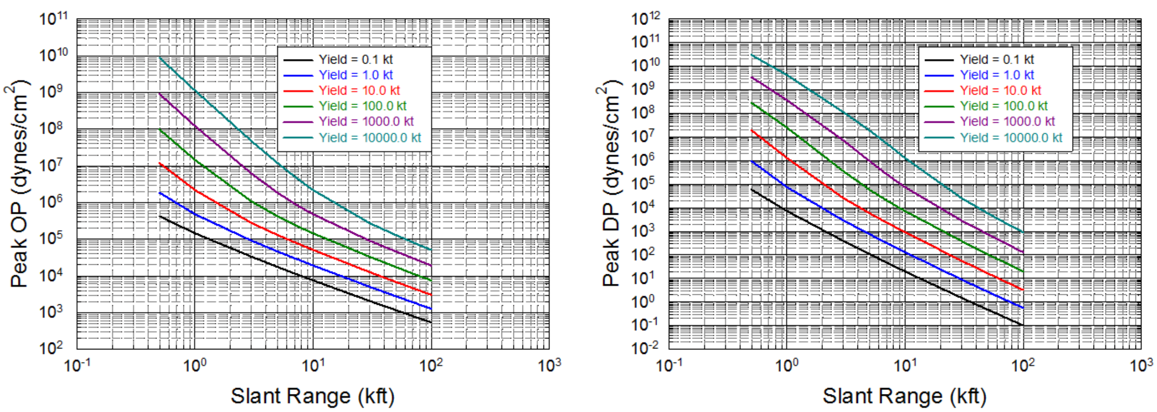

Figure: Peak Overpressure as Function of Slant Range (left) Peak Dynamic Pressure as Function of Slant Range (right)

Major Enhancements Released in CUBIT 15.1. With the release of CUBIT 15.1, on February 29, 2016, the CUBIT Team has delivered major enhancements to three key technologies in CUBIT: Sculpt, Mesh Scaling, and tetrahedral mesh generation. The latest Sculpt contains new smoothing techniques that significantly improve mesh quality. Sculpt also includes new capabilities for generating smooth, all-hexahedral meshes of material microstructures. Mesh Scaling improvements include a new algorithm for more evenly distributed mesh refinement, which gives better results for mesh refinement studies for solution verification. With increasing use of CUBIT to generate tetrahedral meshes, the triangle and tetrahedral meshing in CUBIT have been updated to a new version of the licensed MeshGems meshing libraries. CUBIT is widely used across the Government complex to generate finite element meshes for a broad range of engineering computer simulations. Additional work on Sculpt is still needed to capture sharp geometric features in engineering models. Initial users of Mesh Scaling have identified a number of issues which are being addressed to provide a robust option for solution verification of large meshes.

(POC: Byron Hanks, bwhanks@sandia.gov) **[Brief: CUBIT 15.1 includes major enhancements to Sculpt, Mesh Scaling, and tetrahedral meshing.]**

Challenge: Multiscale uncertainty propagation for materials UQ. Researchers at Sandia National Laboratories are exploring methods for reduced-order multiscale uncertainty propagation. The goal of the work is to develop a tool that provides statistically realistic perturbations to homogenized solutions that are representative of the material's microstructure. A tool to model random fields was applied to generate perturbations based on calibrations to solutions of representative volume elements of elasto-plastic, polycrystalline 304 L stainless steel. Samples of principal strains and stresses were generated for various applied load levels in both uniaxial and biaxial loading. These were then used to estimate engineering quantities, *e.g*., von Mises equivalent stress. The approach demonstrated that statistically accurate samples of principal stress/strain can be generated when a mapping is applied to account for nonlinear dependencies. Further, the model appears to demonstrate low dependence on loading direction and magnitude. This is exciting because it illustrates a path forward; however the post-processing mapping step is undesired. Future work will seek to build these dependencies directly into the model to avoid the necessity of the mapping. (POC: Gustavo Castelluccio, gmcaste@sandia.gov) **[Challenge Summary: Researchers at Sandia are using a random field model to develop a tool for reduced-order multiscale uncertainty propagation.]**

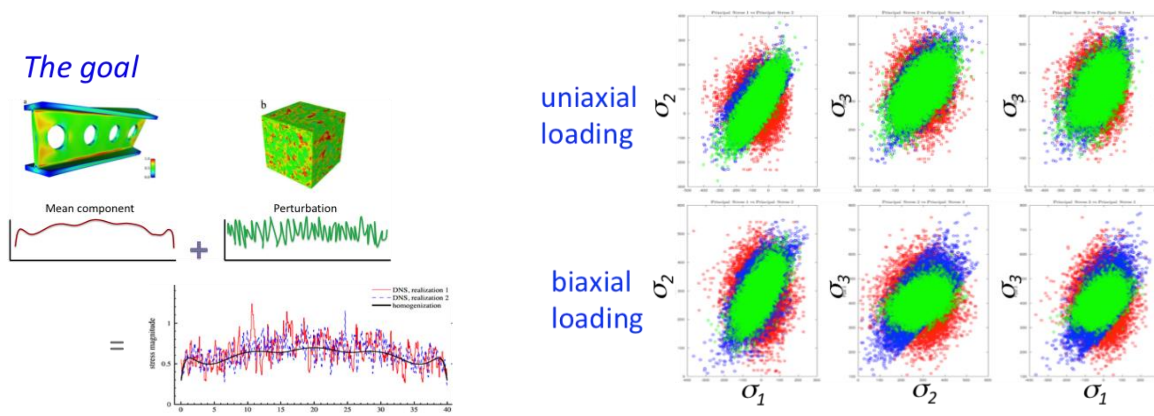

Figure: Schematic of the method (left). Dependencies of the three principal stresses from propagation for two load conditions: calibration data in blue; random samples in red; post-process random samples in green (right).

ASC Accomplishments
July 25, 2016

Improved Radiation Transport in Stochastic Media. P&EM researchers at Sandia National Laboratories have discovered an improved method for performing radiation transport calculations in stochastic (random) materials. Previous approaches can be either expensive (Monte Carlo sampling) or have limited accuracy (atomic mix or Levermore-Pomraning (LP) models). The new "atomic mix closure" combines aspects of atomic mix and LP approaches and is generally more accurate than either method. Such approaches can improve the accuracy and reduce the expense of analyses of radiation effects in certain types of materials in the stockpile. Continuing work will explore generalizations and potential shortcomings of this method. (POC: Shawn Pautz, sdpautz@sandia.gov) **[Brief: Researchers at Sandia have discovered an improved method for performing radiation transport calculations in stochastic materials which reduces the expense and improves the accuracy of such analyses.)**

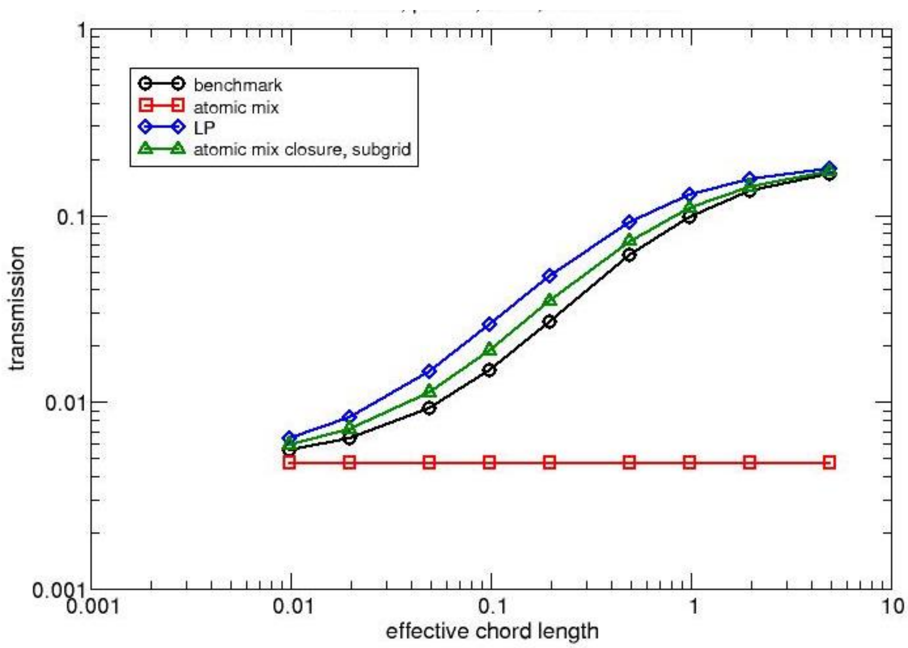

Figure: This plot shows the improved accuracy of the atomic mix closure in comparison to other methods.

Aleph Collision Algorithm Verification: Sandia is developing Aleph as an advanced plasma transport and kinetics simulation code to improve science based design of components. Atomistic collision and energy partitioning algorithms have been rigorously verified against published data for the plasma modeling code Aleph. Using a model problem of heat flux through a rarified gas, we have shown that the Aleph algorithms converge as expected over a range of time and spatial discretizations. Heat flux is a sensitive function of many atomic scale processes including collision rates between atoms, how atoms pick up energy from surfaces and also how energy is distributed between particles after a collision.

Aleph demonstrated linear convergence in both space and time. These studies give our component designer's confidence that Aleph can correctly predict atomic momentum and energy transfer processes for their applications. (POC: Anne M. Grillet, amgrill@sandia.gov) **[Brief: Researchers have rigorously verified collision and energy partitioning algorithms to give our users confidence that Aleph can correctly predict these atomic processes for their applications.]**

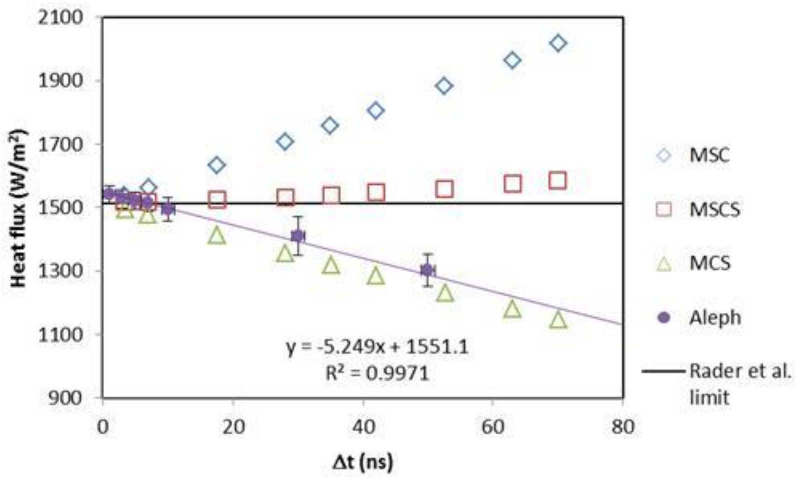

Figure: The calculated heat flux in a rarified gas as a function of temporal discretization showing the expected first order convergence behavior compared to published results by Rader et al. (2006) for several Direct Simulation Monte Carlo algorithms.

Challenge: Generalized Online Fault Tolerance API for MPI Programming. In collaboration with Rutgers University and Intel, Sandia researchers are working toward the specification of a resilience API, **FENIX,** intended for the message passing interface (MPI) programming model. The framework is a merging of separate efforts by Sandia and Rutgers for online application recovery to avoid global checkpoint/restart. The API specification hides the complexity of low-level fault tolerance message passing mechanism (such as MPI-ULFM) to enable quick adaptation of the new recovery scheme. FENIX will support major languages including C, C++ and Fortran. . (POC: Michael Heroux, maherou@sandia.gov) **[Challenge Summary: An SNL collaboration with industry and academia is producing a new, efficient, adaptable application programming interface for resilient message passing programming, usable with all major supercomputing programming languages.]**

ASC Accomplishments
August 1, 2016

System Level Damping Characterization. The challenge of modeling systems with frictional interfaces is that they exhibit significant changes in their stiffness and damping characteristics as the excitation level is increased. This is indicative of a strong nonlinearity. Traditional methods for modeling these systems have relied on linear approaches, which are inadequate. Researchers at Sandia have developed a new suite of tools that adaptively estimate the instantaneous frequency and damping characteristics of a system. As the instantaneous properties of the system change, the toolbox automatically adjusts its windowing parameters to estimate the system properties more accurately. **[Brief: A new toolbox has been developed that adaptively calculates the properties of frictional interfaces in a nonlinear system from experimental data.]**

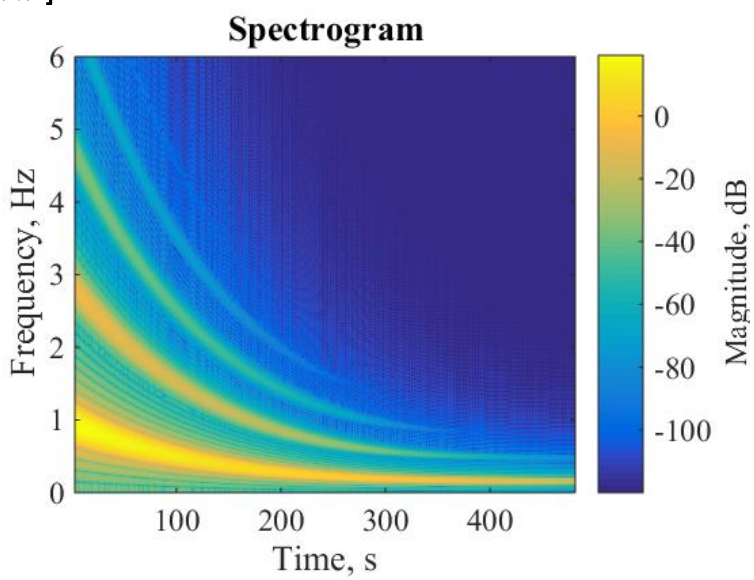

Figure: Example of the extracted information from a ringdown experiment. This data is then processed to determine instantaneous system parameters.

Simulation of Energy Decay in Turbulent Flows. Using 1.57 million cores of Sequoia, DSMC (Sandia's Direct Simulation Monte Carlo code) has recently been able to simulate up to 1 trillion cells and several trillion particles. With a simulation of this size, bringing to bear the whole might of the currently available resources, the energy decay in a three-dimensional flow with homogeneous turbulence at a Reynolds number (Re) of 100 has been successfully simulated (see Figure 1). Since DSMC needs to resolve all length scales from the molecular to the macroscopic turbulent length scales, the newly implemented near-neighbor collision scheme was successfully employed to reduce the discretization requirements, making this calculation possible. (POC: Michael Gallis, magalli@sandia.gov). **[Brief: New collision model has been added to the ASC DSMC code Sparta allowing simulations of energy decay in homogeneous turbulence.]**

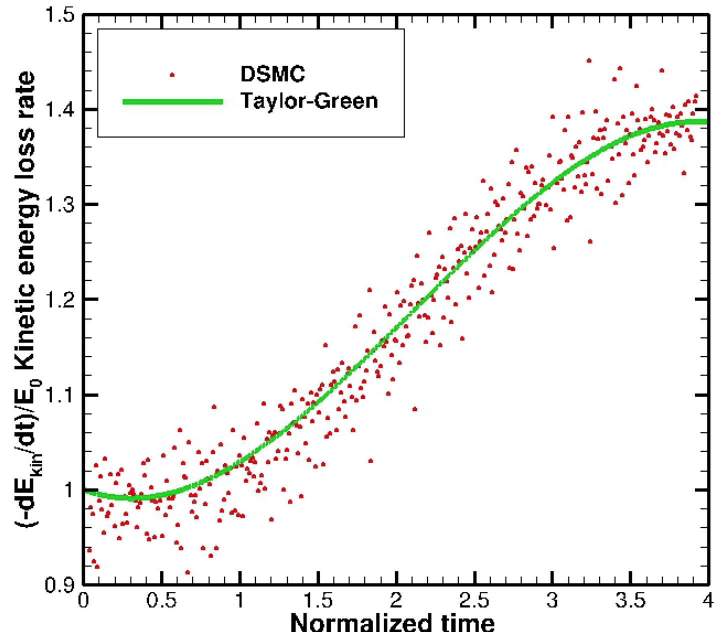

Figure: DSMC simulation of kinetic energy loss in homogeneous turbulence.

Tool Scale Modeling for Additive Manufacturing. Researchers at Sandia National Laboratories have developed computer models that will help to predict porosity and defect formation in metal parts fabricated with the powder-bed fusion process. These simulations provide tool-scale results for the molten metal thermal field which is the required initial condition for metal microstructure simulations. This predictive capability will be invaluable in specifying appropriate process parameters that lead to the fabrication of quality parts using additive manufacturing techniques. (POC: Rick Givler, rcgivle@sandia.gov) **[Brief: Researchers at Sandia have predicted thermal fields for metal additive manufacturing.]**

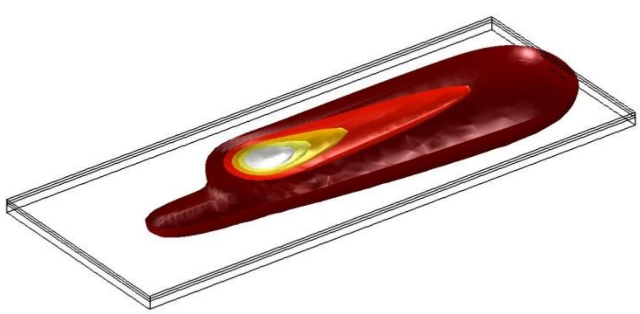

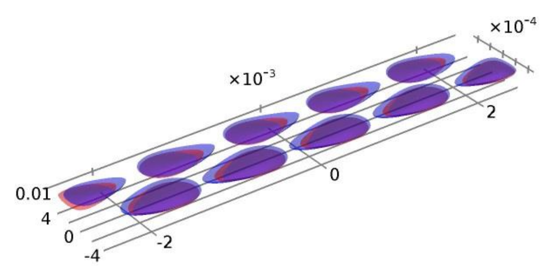

Figure left: Figure: Iso-surfaces of temperature near the top surface of a powder bed during a fusion process - a traversing laser on a layered-bed surface. The laser traverses the same trajectory for successive powder layers.

Figure right: Figure: Sequential snapshots of iso-surfaces of the melt temperature for repeated laser trajectories on successive powder layers. First melt surface is shown in red; second melt surface, after an additional powder layer has been deposited, is shown in blue. Notice that the blue surfaces are slightly larger than the corresponding red surfaces caused by accumulated heating.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.