LLNL-TR-717057

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Carbon Capture Multidisciplinary Simulation Center Trilab Support Team (TST) Fall Meeting 2016 Report

E. W. Draeger

January 9, 2017

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Carbon Capture Multidisciplinary Simulation Center
Trilab Support Team (TST) Fall Meeting 2016 Report

The Trilab support team met with the Carbon Capture MSC project team in Salt Lake City, UT on November 10-11, 2016. TST members present were Eric Phipps (SNL), Ben Bergen (LANL), Greg Burton (LLNL), and Erik Draeger (chair, LLNL). Robert Knaus (Sandia) attended in place of Stefan Domino. Fred Wysocki (LANL) was unable to attend. Rob Hoekstra (Sandia), Jim Costa (Sandia), Fernando Grinstein (LANL), Bob Ferencz (LLNL) and Dan Nikkel (LLNL) were all present from the AST.

The CCMSC is a collaboration of researchers from the University of Utah, Brigham Young University (BYU) and the University of California at Berkeley (UCB), and industrial partners from GE Power (formerly Alstom). Phil Smith is the PI of the project and was instrumental in helping organize and lead the meeting.

Overview and Recent Progress

The theme of this year's meeting was "Predictivity: Now and in the Future". After welcoming remarks, Erik Draeger gave a talk on the NNSA Labs' history of predictive simulation and the new challenges faced by upcoming architecture changes. He described an example where the volume of analysis data produced by a set of inertial confinement fusion (ICF) simulations on the Trinity machine was too large to store or transfer, and the steps needed to reduce it to a manageable size. He also described the software re-engineering plan for LLNL's suite of multiphysics codes and physics packages with a new push toward common components, making collaboration with teams like the CCMSC who already have experience trying to architect complex multiphysics code infrastructure on next-generation architectures all the more important. Phil Smith then gave an overview outlining the goals of the project, namely to accelerate development of new technology in the form of high efficiency carbon capture pulverized coal power generation as well as further optimize existing state of the art designs. He then presented a summary of the Center's top-down uncertainty quantification approach, in which ultimate target predictivity informs uncertainty targets for lower-level components, and gave data on how close all the different components currently are to their targets. Most components still need an approximately two-fold reduction in uncertainty to hit the ultimate predictivity target, but the current accuracy is already rather impressive.

John Schmidt then reported on the recent use of INCITE resources to run long simulations (tens of seconds, hundreds of thousands of time steps) to evaluate USC boiler designs. Several significant improvements were made to the software and infrastructure to enable these runs, including improvements to I/O performance with PIDX, a large scale overhaul of Uintah infrastructure to reduce memory and improve performance, use of a Reverse Monte Carlo Ray Tracing (RMCRT) algorithm, and other performance improvements to reduce the time to solution per time step by over a factor of 2. The two large scale cases planned for Mira are both nearly completed, with a third planned to make use of an additional 100M cpu-hours granted. Titan has presented greater challenges, including a high (>60%) job failure rate at 256k cores, but runs are now underway.

Alan Humphrey described the work performed to move from discrete ordinates to RMCRT on Titan, which required a significant overhaul to Uintah infrastructure. The effort described was substantial and included transitioning the code to the C++11 standard, building performant lock-free data structures, and improving the task graph compilation and communication requirements. Major improvements in memory footprint (~10x), time step (~3x) and overall time to solution (~2.5x) were reported in detail. Ben Isaac then summarized the 8-corner boiler simulations on Mira, including the impact of the previously described optimizations on time to solution. Derek Harris gave an overview of recent work developing a multi-level RMCRT algorithm that gave 5x faster ray tracing and reduced communication times by 500-1000x, with a smaller memory footprint.

The remainder of the meeting was organized around the idea that the project is far enough along and the TST is sufficiently familiar with the high level details from past meetings and reviews that it would be productive to have a TST meeting with presentations and discussions organized to encourage more specificity and greater technical depth. To this end, the meeting followed the theme, with the majority of the time spent in two parallel Physics and Computer Science sessions. The TST split attendance across sessions, with Greg Burton and Robert Knaus attending and presenting talks in the physics session and Eric Phipps and Ben Bergen attending and presenting talks in the computer science session. Erik Draeger also attended the computer science session.

Physics Breakout Session

Sean Smith reported on model sensitivity studies, comparing "top-down" and "bottom-up" studies of model sensitivity. The top-down/bottom-up analysis framework attempts to determine how well current tools propagate the uncertainty in opposing directions within the same model structure, providing a check on individual model contributions. Work continues quantifying model uncertainty, narrowing the parameter range for numerous physical models. They are evaluating models for ash deposition, involving a variety of secondary physical processes, including radiation, probability of deposition, deposition rate, deposit thickness, conductivity, particle stickiness, melting temperature and viscosity. They are also evaluating the models for char oxidation physics, including the contributions of surface reaction rates, mass transfer, pore surface area, pore diffusion rates, annealing, heats of reaction at material interfaces, adsorption chemistry and thermodynamics. Much of this work is being driven by the availability of experimental data and further experimental investigations that may be conducted in the current PSAAP2 project. Ultimately the work will attempt to evaluate all known physical processes/phenomena that impact both the bottom line for the particular model brick, but also the ultimate quantity of interest (QOI) for the project.

Troy Holland then discussed in detail the coal combustion modeling effort conducted since the spring TST meeting. He began by discussing char oxidation, and the model forms for the sub-component physical processes. He indicated that all were complex, and that any one of which could drastically alter the ultimate QOI. These include models for coal plasticity, devolatilization, swelling, thermal annealing and burning rates. He noted that simple heuristic models were insufficient, in part because of the very intense O_2 conditions and non-negligible H_2O and CO_2 gasification. There are a number of more advanced methods, noting Haugen in particular, but the framework of these are often too complex to be easily incorporated in the present global computational model framework. The current model for particle swelling, for example, does not adequately capture the range of bubble sizes, which in turn produces inaccurate estimates of heat and mass transfer. Similarly the thermal annealing models are overly sensitive to coal type, heating rate and peak particle temperature. There remain large areas for obvious improvement, including the reaction kinetics model. Comparisons of the current models optimized for datasets like the Black Thunder Data show a systematic skewness, and remain further problematic given the uncertainty in boundary and initial conditions. The datasets also indicate enormous particle size variation and noise in the measurements. Phil Smith commented that an improved CFD model for these processes must contain information about particle burnout, peak temperature, proper distributions relating to characteristic particle size and p.d.f.s of the inter-particle reactions, given those sizes. He said that we "need a measure of what we are shooting for," and then, at least initially, some way to reduce the degrees of freedom to make the model computationally tractable.

Michael Frenklach of the Berkeley group discussed their continuing work on the bound-to-bound (B2B) protocol for model validation, including the vector consistency measure (VCM). He said that the B2B technique takes the view of uncertainty quantification (UQ) through models constrained by the data, where prediction is established within a range allowed by model form and data constraints. The degree to which a model parameter vector replicates a data set can be measured with the VCM procedure, discussed in some detail as a sensitivity analysis conducted through an iterative process. He then discussed application of the quadratic surrogate model within this framework, which is used to approximate a true solution. The goal is to look for the simplest model that can explain the experimental data. Given this, the discussion focused on combining the various available data to validate the model forms using the B2B formalism. The goal is to look at data for the initial particle size distributions, and other uncertain parameters in the instrumentation models and the physics to establish a likely model bound for each sub-physics process like oxidation reactions and burning models. They did note that updates to the instrumentation models had been particularly effective in narrowing the uncertainty. They then discussed how the B2B consistency analysis and data organization has led to an automated workflow for exploring different model forms. The Berkeley group ended by discussing hybrid statistical deterministic methods for uncertainty quantification, including brute force sampling, random walk sampling, bounding polytope sampling and approximate polytope sampling. In summary the Berkeley group explained their vector consistency measure for evaluating model forms and performance, expanding the validation framework in B2B to include VCM, applying B2B tools to validate char oxidation models for CCMSC and investigating approximate sampling schemes for high-dimensional feasible sets.

Phil Smith then spoke about next steps in improving the char oxidation model to within a 15% uncertainty measure. Saying that "we are not there yet," Phil suggested combining all the data sets (390+) to leverage model development through the Berkeley framework. However, he indicated that the strategy going forward was not entirely clear. He asked whether we should use Troy's model form and simplify it? He said he remained uncertain as to whether those model forms would be computationally tractable. He reminded the group that having started with the final range of acceptable error for the overall project, we could do an inverse problem to determine what the bounds of accuracy should be within each of the necessary model components, even though this would require "throwing a lot of cores" at the problem.

David Lignell then spoke about the BYU effort in soot physics modeling. He first described the physics of the soot-formation process in some detail, including nucleation, growth, oxidation, gasification, coagulation and aggregation. Work since the spring has focused on soot in radiation transport, soot as a mixture fraction sink and as a source of emissivity, with the goal of developing a suite of soot formation models of increasing complexity and accuracy. Their approach uses a Bayesian statistical framework to perform V/UQ studies against experimental data, quantifying modeling errors and performing sensitivity analyses of model parameters and forms. They have recently focused on evaluating model accuracy and sensitivity of models for soot oxidation, soot growth and nucleation rates. Going forward, they will continue to focus on model development and validation, especially extending the soot growth and nucleation rates models of Brown & Fletcher, fitting the model output to previously existing flame burner data. They will also provide guidance for future experimental work that will produce new data that will likely allow insights necessary to develop additional advanced model forms.

Robert Knaus then discussed Sandia NL's predictive wind plant modeling program. He introduced the problem of forecasting the effect of wake turbulence in a large array of wind turbines in a given wind farm operation. He introduced the team, discussed the physics modeling issues, and outlined the computational approach, involving grids of 10 billion mesh elements needed for wall-resolved large eddy simulations of the problem. He also outlined the high-order derivative operator implementation, the use of overset meshes and the issues that the project team has encountered increasing the solver performance at the computational scales required for the problem. He concluded by showing animations illustrating the operation of the sliding mesh around a cube, capturing the resulting wake turbulence in great detail.

Jeremy Thornock then discussed his group's recent efforts quantifying numerical uncertainty of the multi-physics large-eddy simulations. The group performs routine code verification work, confirming that the code gives "the answer to the problem you want", using a variety of tools including regression testing, convergence analysis, resolution studies, and the method of manufactured solutions. The solution validation work has required the group to develop new techniques to characterize and quantify numerical errors relating to model forms and parameter values. Their work has focused on grid resolution-induced errors, errors due to radiation angle, and polydispersity (DQMOM quadrature). This work uses the Richardson extrapolation technique to back out a convergence rate directly from the algebra, which can, in turn, be used to estimate the error in the given measure. They performed a series of simulations of the base case 8 corner unit, and showed animation results from simulations whose resolutions varied by a factor of 64. Their analysis focused on the estimation of midplane temperature and wall heat flux for 5-square meter panels on the surface of the radiation section. Jeremy noted that they were seeing errors in comparison with previously published data in the range of 50%. He said they were as yet uncertain as to the source of these errors, but noted that they are currently evaluating whether they have enough solution samples at a given resolution to be within the asymptotic regime, where statistical measures would be converged. They are in the process of evaluating also whether any of the solutions resolve enough of the turbulent length scales to properly represent a near-fully resolved simulation of the system. Phil Smith noted that these quantities were not the only ones of interest and it was difficult to determine how the measured quantities were related to the quantities of interest defining the performance of the entire coal burner apparatus.

Greg Burton spoke about Livermore's Turbulence Analysis & Simulation Center's (TASC) proposal for the DOE Exascale Computing Program regarding hypersonic scramjet design. He described the overall focus of the project as the simulation of a combined hypersonic boost/glide vehicle in maneuvering flight. Burton noted the many physics-based modeling challenges of the proposed project in conjunction with the computer science aspects of the project directed to efficiently accessing and evaluating the enormous amounts of data generated from such a simulation at exascale. He noted that the project team was informed mid-way through the proposal development process that the DOE Office of Science had already determined that they would not endorse a hypersonics-based ECP proposal, notwithstanding the support of other sections of the DOE for future hypersonics work.

Derek Harris discussed improvements to the Discrete Ordinate radiation solver in the principal physics-based LES code discussed previously. They noted that the traditional form of the discrete ordinates solver had been used for many decades and was a major part of the computational overhead. In trying to reduce the overhead, they had been using HYPRE library to solve for radiation intensities formulated as a linear system, but noted only weak scalability. They ultimately sought to mitigate the burden by solving only every 20 fluid timesteps, by using a gray gas assumption, and by using a new Gauss-Chebyshev quadrature set. A "deep-dive" session at Texas A&M motivated the group to pursue a different route, which has resulted in their adopting an approach based on reverse Monte Carlo ray tracing (RMCRT). They noted that highly parallel nature of the underlying approach, which shifted most communication costs to setting up the initial domain and the recursive nature of the approach, which has the advantages of computational efficiency and the reduction of error due to ray effects. They are also pursuing multi-level RMCRT to further reduce computational cost. They have noted that this approach produces two sources of errors, but that the sampling error reduction tends to offset these errors. They are now in the process of validating the multi-level RMCRT, and are showing excellent results when compared to field data. In the future they will complete addition of a mean adsorption coefficient, Beer-Lambert weighting function for the quadrature, and tabulation and integration of high-resolution spectral data to further validate the implementation. They also plan to develop benchmark cases for coarsening models in RMCRT and to explore radiation diffusion.

Jennifer Spinti next described her experiments in Alt-1500 facility which has allowed the team to generate its own data for model and simulation V&V. The facility has a single burner and is of small volume, making wall effects important. She noted that the use of the facility is a critical part of the Project Plan, because it can respond to the changing needs of the modeling and simulation development effort. Her 2016 effort has focused on getting QOIs relating to air-fired coal combustion in the facility. Obtaining valid data is quite challenging, as one needs to understand where and what to sample and what instrumentation would serve that purpose best. She has been particularly concerned about the reproducibility of individual measurements, and has spent a considerable amount of time redesigning cooling coils, for instance, to obtain better reproducibility of the measurement. She also is looking actively at instrumentation models for thermocouples and similar sensors, to better understand how response of materials within the sensors can affect the quality of the data taken. This is an ongoing effort that will continue into 2017.

Oscar Diaz-Ibarra discussed V/UQ methodology for validation of computer models. This requires working with the QOIs of interest, including incident wall heat flux and wall heat temperature and constructing an input/uncertainty map containing the effects of the various physical models including those discussed earlier in the session like char oxidation, and then performing a sensitivity analysis using the Sandia Toolkit. He presented the outlines of his sensitivity analysis for the wall temperature QOI, and noted that the statistic was most sensitive to the wall-emissivity estimates and the energy deposition, while other parameters like the coal feed rate, swirling parameter and primary inlet temperature were not. The data from this initial effort was used to guide the first experimental campaign in 2015. The group subsequently used a similar analysis to analyze the uncertainty in the experimental measurements.

Computer Science Breakout Session

Martin Berzins presented the evolution of the Uintah framework for predictivity:

- Martin presented a nice summary of their view of the key challenges for ensuring software can be ported to new architectures for 2020 and beyond, and their approach for addressing these challenges:
    - Efficiency at extreme-scale: task-based formulation
    - Programming model to encode tasks: Uintah
    - Run-time system to execute tasks: Uintah
    - Performance portability of task implementation: Kokkos
    - DSL to enhance code developer productivity: Nebo

    - Along with maintaining interoperability of DSL to handle DSL limitations in terms of problem scope
    - Managing resiliency: AMR-based patch duplication
- Martin discussed issues the team is seeing with MPI-thread-multiple working correctly and robustly. This is likely due to lack of extensive use of it in the HPC community, but that would be expected to grow as MPI+X approaches permeate the community.
- Martin also discussed the growing usage of Kokkos in both the Uintah system as well as directly in the application codes. Users appear to prefer it in some cases over the Nebo DSL due to the greater control of loop structure. It also has broader applicability of the DSL, so maintaining interoperability between Nebo and Kokkos will be important.
- Incorporating Kokkos into some relevant Arches kernels has demonstrated significant improvements in performance, mostly due to improved C++ implementation with fewer abstraction layers, less indirection, and better compiler optimization.
- The group has significant experience in gradually porting a code to Kokkos, and it would be important for the group to share their experiences/approach to the broader Kokkos user and developer community.
- The team should consider refactoring Uintah as a stand-alone tasking runtime for users with structured mesh management needs. This would require additional funding and staffing, but could provide a useful tool for the community, especially given the level of maturity of the Uintah software stack. Extensions of the task and data models to support unstructured meshes would also be an interesting direction if resources become available.
- Potential for collaboration: LANL has a BlackBox MG project that provides robust multigrid solvers for structured meshes. David Moulton of LANL (moulton@lanl.gov) is a good contact about this project.
- Valerio Pascucci summarized the status and evolution of PIDX for I/O management
- Substantial improvements in I/O performance were presented that make analysis of very large-scale simulation more feasible, as well as provide the opportunity to save more information than was previously possible. This is a significant advancement over previous capabilities and goes a long way to help the center achieve their analysis objectives.

- The group is extensively exploring compression approaches for I/O as well as the in situ analysis pipeline. They are incorporating external approaches such as ZFP as well as researching other schemes that may be more extensible to multiresolution formats.
- The group is beginning to explore in situ visualization of particle data, which is clearly important for the center, as well as started the development of a new EDSL for interactive data analysis and visualization.
- One suggestion is to implement I/O operations using the task model, as this is one area where there is a potential for increasing concurrency, i.e., use an offload model to arbitrate I/O operations that allows the simulation to proceed rather than blocking until the I/O operation is complete. The structure of the server/client model for accessing remote data seems well-conceived! Image compression, bit representation and resolution reduction for visualization preview are forward looking ideas to reduce time and energy costs for remotely accessing data.

Alan Humphrey discussed recent modifications and improvements to the Uintah infrastructure.

- Alan summarized work to incorporate C++11 features into the code that led to improved performance, reduced substantial amounts of legacy code, and improved developer productivity. This also made it easier to incorporate tools such as Kokkos.
- Uintah has pulled in jemalloc to provide a better memory allocator and reduce fragmentation. This also removed the burden to maintain their own allocator.
- New lock-free data structures were added that improve performance, particularly for high-thread count architectures such as GPUs. These are being used in their production code runs on Titan. One of these data structures (unordered map) is maintained by the Kokkos team, and it may make sense for them to extract these data structures from Uintah and put them in Kokkos (or somewhere else) to make them more publicly available.
- Pushing Uintah to the full Titan machine exposed hardware issues with the nodes going down, which OLCF is in the process of resolving.
- Many optimizations of the Uintah framework for the Titan production runs were presented, including substantial improvements in task-graph compilation. The outcome of this is a strong example of the benefits of

close interaction between the CS and physics teams to make the very large scale runs possible.

- We recommend that the team make sure to keep up-to-date with C++17/20 features (in collaboration with the Kokkos development team, and potentially with the LANL C++ committee representatives). Some of the Kokkos functionality is moving into the standard, e.g., parallel STL algorithms and containers. There are also opportunities for performance improvement through the use of C++11/14/17 features. However, it is clear that the project is aware of this and is making progress in adopting these improved features.
- Programming model interactions: It would be useful to have a meeting between the PSAAP centers on programming models and runtime development. One motivation for this is the issue of data ownership between system-level and node-level runtimes, or between different runtimes within a single level of the hierarchy, e.g. how could one leverage a solver written on Uintah with one written on Legion? (This is probably ill-advised.) Potentially, this type of interaction will take place at the WEST meeting this spring. The utility of such a discussion could expose the need for unmanaged or initialization interfaces in these runtimes to allow management of data that are allocated outside of the runtime, or release of previously managed data.

Eric Phipps discussed the development and use of embedded analysis capabilities at Sandia through Sandia's Advanced Technology Development and Mitigation (ATDM) program. Here "analysis" refers to techniques such as simulation-based optimization, uncertainty quantification, sensitivity analysis, and model calibration, and "embedded" means approaches that require more information from simulations beyond black-box samples. Eric first motivated the need for analysis methods in Sandia's nuclear weapons program, particularly in the context of exascale computing. He then described the need for embedded approaches that provide better performance, scalability, and accuracy than traditional black-box sampling-based approaches. Next Eric summarized how these approaches are being developed in Sandia's ATDM program using code transformation based on automatic differentiation techniques. Then Eric described how this work is realized in software, focusing on how these techniques are integrated with Kokkos for propagating analysis information such as derivatives through Kokkos parallel kernels. Finally, Eric presented recent results of applying these techniques to Sandia's two

ATDM application codes (hypersonic reentry and hostile E&M/plasma) through a recent L2 milestone.

Chuck Hanson and Allen Sanderson: VisIt for Scalable Viz

- One possible extension to the monitoring tool is to consider allowing JIT compilation of user tasks for true computational steering. This would likely be straightforward since the distributed communication infrastructure is in place.

John Holmen presented work on preparing the RMCRT implementation for the Intel MIC architecture.

- This primarily entailed incorporating shared-memory data parallelism into the RMCRT task implementation using Kokkos. Incorporating Kokkos motivated code modifications to reduce indirection/abstraction layers that resulted in substantial performance improvements.
- The use of data parallelism within the RMCRT tasks enables the use of fewer tasks per compute node, larger patch sizes, reduced memory footprint, and reduced task overhead, all resulting in improved performance.
- The use of Kokkos also should eventually allow consolidation of the Intel and GPU implementations.

James Sutherland presented the status and challenges with the Nebo DSL

- James summarized work on the weak scaling and absolute performance of CFD algorithms on Titan, where good weak scaling was observed, but poor performance of the GPU relative to the CPU was observed.
- Capabilities for dense linear algebra were added to Nebo, enabling block-implicit algorithms.
- Work exploring a potential Kokkos backend to Nebo was also presented. This has the potential to reduce the code maintenance and porting burden for Nebo, since it would rely on Kokkos for portability. However poor performance compared to Nebo with its own backends was observed. This performance was also verified by comparing to Kokkos-only implementations of Nebo+Kokkos. The Kokkos team will be at SC, so hopefully they can meet to diagnose what is going wrong.
    - Update: James met with the Kokkos team at SC and discovered data layout issues in their usage of Kokkos, which will hopefully resolve the poor performance observed.

- The current approach of developing raw Kokkos kernels to understand the observed performance variability is deemed good. It would be nice if this analysis could be used to present some basic performance data on the improvement that can be achieved by using the Kokkos data model (layout features) for the types of solvers used by this project. Understanding the inefficiencies in overhead seems to be the core concern, however, so it is recommended that investigations into this issue take precedent.

Ben Bergen then presented an introduction and overview of the FleCSI programming system and reported on the LANL ATDM project FleCSALE. FleCSI is a computer science infrastructure system that allows developers to create mesh and other topological data structures, e.g., hashed octrees. FleCSI also provides control, execution and data models for task-based programming. In particular, FleCSI has backend support for Legion (http://legion.stanford.edu). FleCSI is open-source and is available on github. FleCSALE is an ALE hydrodynamics code that utilizes FleCSI. It currently has support for arbitrary 2D polygonal and 3D polyhedral meshes, and can run in pure Eulerian and pure Lagrangian modes.

Day one concluded with an evening poster session highlighting recent work in both physics and computer scientists. Fourteen posters were presented and as with past years, the committee greatly enjoyed the opportunity to speak to poster authors interactively about their research and go into greater depth. The venue did present challenges in terms of navigating and speaking to all poster authors, however.

Confidential Briefing and Wrap-Up

Day two featured several talks covered under the Center's confidentiality agreement. Specific details will thus be omitted from this report to avoid inadvertent disclosure / unauthorized distribution. Phil Smith, Derek Harris, and Ben Isaac discussed the current status of boiler designs and described some new research challenges that have emerged. It was striking how valuable simulation was in the design process. One inescapable theme was the amount of untapped potential that remains for predictive simulation in this space, for example in the role of troubleshooting design changes or mitigating unforeseen effects once exascale-class resources are widely available to enable the necessary number of runs needed for discovery-level simulation studies.

The next three talks featured graduate students describing results from their PSAAP summer internships at the Labs. John Holmen described the work he did on specialized multigrid methods with Jonathan Hu and Ray Tuminaro at Sandia (CA). Pavol Klacansky described his internship at LLNL working with Timo Bremer and Valerio Pascucci on developing a progressive merge tree that can adaptively change with the mesh. He used SAMRAI to generate test cases, and made a point to note how helpful he found the Computation Hackathon. Josh McConnell worked at Sandia (NM) with John Hewson and Robert Knaus calculating particle statistics and deposition rates for turbulent channel flow. Dav de St. Germain finished the session with a discussion of the project's current HPC resources and plans for how to continue to get the significant resources needed once the INCITE allocation is finished.

Closing Thoughts

The TST continues to be very impressed with the quality of work presented and the rapid progress made by the CCMSC on a highly complex multiphysics problem domain. The Center is gathering increasing expertise and experience in multiple areas likely to be of broad interest to DOE and the greater HPC community, particularly in the areas of programming models and DSLs, I/O, V&V/UQ, and visualization/data analytics. The TST encourages the CCMSC to engage with the Exascale Computing Program wherever possible, and to submit a proposal when the call for academic participation is released in 2017. The TST will coordinate with their ECP contacts to ensure CCMSC has all needed information to submit a strong proposal. The TST is also eager to identify more opportunities for engagement with the Labs. Mass visits by numerous Center members to one or more Labs was deemed likely to be inefficient, as ensuring that pulling this much time from Center staff is worthwhile and that the right people at each Lab are available and willing to participate is difficult. Encouraging participation by both the Labs and CCMSC at PSAAP or DOE workshops, such as the WEST workshop at Stanford in February, 2017 or the Performance Portability Center of Excellence workshop in Glendale, Arizona in April, 2017, was viewed as a more efficient way to make connections and can be followed with targeted visits. The TST will work to facilitate this.

The new format was viewed favorably, as the smaller workshop-style sessions seemed to generate more discussion, although some members did note that they felt like they were missing out by only getting to hear half the talks. One suggestion was to have shorter sessions or more formal breakouts on targeted topics. The poster session was deemed to have been less successful than in past years, as not all posters had authors standing near them and the venue did not allow for all posters to be in the same room. (It must be acknowledged that the venue change is due to scheduling uncertainty created by the TST.) We do encourage the use of the original Guest House venue next time if possible.

For the upcoming AST review, the TST encourages the Center to be sure to include a detailed risk assessment for the remainder of the project, with specific risks and corresponding mitigation strategies explicitly identified. We also suggest including numerical accuracy as a term in the top-down UQ if possible, as the team felt it was important information that would strengthen that story.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.