U.S. DEPARTMENT OF ENERGY SAND2016-7624R
National Nuclear Security Administration

Predictive Science ASC Alliance Program (PSAAP) II

2016 Review of the Carbon Capture Multidisciplinary Science Center (CCMSC) at the University of Utah

A. L. Hungerford and D. R. Montoya (LANL), R. M. Ferencz and A. L. Kuhl (LLNL), R. J. Hoekstra and K. Ruggirello (SNL)

ABSTRACT

The review was conducted on May 9-10, 2016 at the University of Utah. Overall the review team was impressed with the work presented and found that the CCMSC had met or exceeded the Year 2 milestones. Specific details, comments and recommendations are included in this document.

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. SAND NO. 2016-???? TR

Executive Summary

GOAL: This project is developing a numerical simulation capability to accurately predict the two-phase reactive flow in industrial boilers (e.g., 350 MW boiler of Alstom Corp) for electrical power generation. The goal is to use this simulation capability to design future industrial boiler facilities—in effect, a numerical boiler (similar to the numerical wind tunnel of NASA) — thereby streamlining the process, or potentially eliminating the need, for building prototypes to design new facilities. This requires accurately modeling all key physical effects in the boiler combustion chamber.

The review team convened at the University of Utah March 9-10, 2016, to review the Carbon Capture Multidisciplinary Science Center (CCMSC) funded by the 2nd Predictive Science ASC Alliance Program (PSAAP II). Center leadership and researchers made very thorough and informative presentations, accurately portraying their work and successes and candidly discussing their concerns.

As a result of the presentations, the review team identified several areas in which the center is to be especially commended:

- In response to 2015 review recommendations we believe your V&V hierarchy has been improved and its credibility has been strengthened
- The expanded description of particle transformation was very helpful.
- The poster session was very valuable and we commend the students for the preparation and technical acumen.
- Continued expansion of industrial partnerships is a strength for this Center.
- For the second year in a row, you used V&V/UQ to identify new physics, and shifted your program priorities accordingly.
- Your latest emphasis on macro mixing is a natural extension of last year's conclusion.
- The well thought out modeling of instrumentation including uncertainties continues to prove very valuable to the credibility of your validation.
- This center has been the best at procuring cycles on HPC platforms.
- The degree of University support continues to impress, and attests to the trust, faith and commitment that the University places in this project.

At the conclusion of the proceedings, the review team universally expressed that CCMSC has made excellent progress in Year 2, meeting or exceeding all Year 2 goals for milestone predictions. Comments related to the specific areas and recommendations are below.

Science and Engineering Research

In this section, we summarize our understanding of the team's current findings, and offer ideas to be considered as you plan your path forward.

Multi-Phase Physics

The CCMSC's simulation effort is built upon a large eddy simulation (LES) model of the combustion of poly-disperse particles in a turbulent flow field. The LES model equations with appropriate closures are solved for the conservation of mass, species, momentum and energy. The conservation equations are solved using a typical 2^{nd} order finite volume approach on an adaptively refined grid.

In order to model the low Mach number compressible flow in the boilers, the LES conservations equations are solved using a Projection method first proposed by Chorin (Chorin, Numerical solution of the Navier-Stokes Equations, 1968) and extended to second order by Bell et. a.l. (Bell, 1989). In this approach the Naiver-Stokes equations are expanded in Mach number (M). Neglecting terms of order M^2 or higher results in a Low-Mach-number set of conservations equations which are hyperbolic but only contain a single characteristic, the particle velocity (u). The acoustic waves are thus eliminated analytically, and a less restrictive time step criterion can be used based on the u, instead of u+a. This results in stable time steps which can be on the order of 100 times larger than the usual acoustic limit. However, the method now requires the solution of a Poisson equation for the pressure which is currently being solved using the LLNL HYPRE library. Currently the scaling of HYPRE for large problems is a limiting factor for the scalability of ARCHES.

In order to achieve better scalability, the CCMSC is considering pursuing explicit fully compressible approaches. While these methods might have a smaller stable time step it is anticipated that the better scalability and performance of them will result in an improved overall performance. The Center plans to first evaluate artificial compressibility approaches (see (Chorin, A Numerical Method for Solving Incompressible Flow Problems, 1967)), which alter the speed of sound in the simulation to reduce the acoustic stiffness. The risk with this approach is that the that the thermodynamics of the problem are modified and it might affect the results in an unacceptable way. Another alternative the Center plans to investigate is a pseudo-transient continuation method which uses preconditioning within a dual time approach to allow accelerated convergence. The method will require point-wise dense linear solves and avoids any global linear solves, which should result in computationally dense kernels which will be easily scalable.

The particle phase is simulated using the Direct Quadrature Method of Moments (DQMOM) which is based on a transport equation of the probability density function (PDF) of the particle velocity. It uses a moment-transformed quadrature-approximated number density function (NDF) transport equation. The particle phase includes models to capture,

- Convective heat transfer between the gas and particle phase
- Combustion of coal particles
- Particle evolution due to moisture evaporation, de-volatilization, gas-char reactions, and soot and ash formation

The Center is currently looking at the CQMOM approach to overcome some of the deficiencies of the DQMOM, such as the mono-kinetic limitation, wall reflections, and more stable handling of breakage. The new method also allows trajectory crossing and particle rebounding. The Center is investigating new coal reaction models, such as the Reacting Particle and Boundary Layer (RPBL) model to investigate the reaction and diffusion processes in the boundary layer around the particle.

Energy transport due to (thermal) radiation is currently solved using a Discrete Ordinates Method (DOM) and includes the treatment of,

- Radiation scattering
- Radiation absorption on coal particles, soot, and light gases
- Radiation emission from coal particles
- Wall temperature that couples with the heat conduction model of energy flow in the boiler walls and tubes

At the 1st year's review the Center presented a new Reverse Monte Carlo Radiation Transport (RMCRT) method as a replacement for their DOM. The DOM involved multiple, sparse linear solves and comes with challenges for the incorporation of radiation physics such as scattering, as well as challenges for scaling to very large scales. RMCRT offers solutions for many of these deficiencies. During this year's review the Center demonstrated simulation and scaling results which showed that the RMCRT was able to resolve finer scale details more efficiently, and scale to large problem sizes better compared to the DOM. The RMCRT method also takes advantage of the AMR infrastructure present in Uintah to reduce the memory foot print and computational complexity of the method.

Finding:

The usage of low-order methods (2nd order Roe solver with a minmod limiter and 1st order time integrator) in the LES large scale simulations are probably not able to accurately capture the moments of the velocity and other field variables.

Observation:

For an accurate LES, higher-order methods are typically required to accurately resolve the moments of the flow without having to use a restrictively fine mesh resolution due to computational requirements. The minmod limiter is the most dissipative of the limiters.

Recommendations:

- Look at how well the current lower order methods are able to resolve the moments of the velocity.
- Explore using higher-order approaches spatially, and possibly a 2nd order time integrator if needed.
- Utilize a less dissipative limiter than minmod.

Computer Science Research

The Utah PSAAP2 Center continues to model excellent integration between CS, Physics, and V&V/UQ efforts. This is partially driven by the demands of industry driven partners with real world problems to solve and supporting data for analysis.

The Utah PSAAP Center CS effort continues with its categorization of being both evolutionary and revolutionary. It has built on prior recommendations and efforts this past year with exceptional accomplishments. This includes a large effort in refactoring Uintah and integrating Kokkos. The portability challenge is being addressed through a complementary approach utilizing the Nebo DSL and the Kokkos abstraction layer. This has developed a middle ground with varied approaches for exploration and performance assessment. In addition, they have recomposed their data warehouse implementation to support the CPU and GPU individually to support task queues. This was part of the well-integrated work done to support GPU implementations this past year.

There was much effort seen in the evaluation of algorithms and assessment on how to make CS work for them. Assessment and implementation effort to make algorithms more in line with what architectures can effectively compute was well focused. This included compressible CFD, significant implementation work for RMCRT and discrete ordinates method (DOM) using a sweep approach. These are aimed at reducing the Center's dependence on the scalability of external solver solutions (such as HYPRE) and represent a nice body of work to address one of their major concerns from the past review.

The Utah PSAAP Center CS approach is still primarily focused on how to run on the next generation of hardware in the near term (4-5 years) and how to evolve Uintah and associated DSLs to run well on those platforms. This, in itself, is a formidable challenge. There has been broader outreach occurring this past year to further communicate possibilities within the community. Since a key goal of the centers is to capture their experiences with innovative CS capabilities and feed those lessons-learned back to the communities, we feel that this must be reinforced.

Data is becoming one of the key bottlenecks as we move toward exascale. This is an area that is in constant flux with technical transitions expected from industry in the next few years. These technical transitions will likely impact the workflow and pose portability issues based on platform implementation. There are opportunities in this space, but taking advantage of them will require tight coordination with the project's technical strategy.

Finding: The IO aspects of the application appear less integrated with DAV effort

Observation: Work has been done to further PIDX and also to optimize application IO at the node vs. core level but does not seem to fit under a strategic plan.

Recommendations:

1. Assess needs in application IO and focus team efforts to develop a sustainable solution that will be amenable to architecture differences that will impact this area.
2. Integrate solutions to address the application memory constraints and needs as IO layers are blending through to memory.

Integration of all aspects of the computational environment targeting a level of flexibility, balance of resources and computation, and workflow throughput is key to moving toward exascale. The architectures being proposed and implemented, as well as their supporting software environments, are in a broad range and there is much exploration being done to generalize solutions. While generality can be easily over emphasized, striking a balance that sees codes porting to emerging architectures without massive rewrite investments for every architecture is desirable. Careful integration of technologies and feedback to the communities is a key goal.

Finding: Overall integration and workflow assessment needs greater consideration

Observation: A characterization approach has been developed, some analysis is being done, but need to further assess and collect workflow performance metrics across the integrated project and communicate findings with NNSA partners.

Recommendations:
1. Quantify and collect metrics for the seven abstractions that are being used to characterize architectural portability. We would like to see a more strategic approach for addressing scaling and performance characteristics for your integrated applications. NNSA partners would find data collected to be very valuable for co-design.
2. Integration of novel solutions developed by the PSAAP and broader DOE community are key to understanding potential future strategies. Continued integration and assessment with feedback to NNSA partners is key to broader growth in the community.

Finding: Continued integration with the broader community for technical assessment and feedback should be a priority

Observation: Commendable interaction with labs regarding Kokkos and Hypre. Additional interaction needed to impart your knowledge and what is needed for scalable debugging (tools), IO, and other exascale technologies.

Recommendations:
1. Engage Exascale communities in best practices and emerging standards. Look to the broader community for tools for performance analysis and debugging.
2. Broaden interactions with other teams in the PSAAP and NNSA communities that have technical efforts underway that have an impact on project goals.

In general, the Utah PSAAP Center CS effort continues to take a very pragmatic approach. It seems balanced for the point of uncertainty that we are at today and for their focus on a production project integrating physics, UQ, industry efforts and goals. Their programming environment and runtime is further evolving, and there are efforts to further compare and share with other runtime studies. A focused IO strategy, continued work on architecture abstraction, and broader integration analysis will support their continued growth along their roadmap.

Validation and Verification / Uncertainty Quantification

The review presentations and posters provided evidence of many aspects of the VV/UQ efforts incorporated into the Center's research activities. Our comments are derived from these materials as well as our conversations with Center researchers. A formal validation planning document is being recommended again (see ASME V&V Standards Committee documentation), as a useful artifact to provide the review team.

The new validation hierarchy has been improved over what we saw last year. It better captures the intent of the center and more clearly identifies how the large scale prediction is built on a base of more fundamental experiments and calculations. Clearly one of the strengths of the Center continues to be integration of experimental and computational data with a clear focus on quantification of the uncertainties. The continued development and improvement of instrument models is commendable and adds to the credibility of the work.

Finding: The overall validation hierarchy has been improved.

Observation: Continue to develop a clear and comprehensive description of the rollup of "bench-scale" experiments and models up to the "keystone" calculations.

Recommendations: Clearly indicate the inputs and outputs and their associated uncertainties at all levels of the hierarchy. Capture the flow of information as is passes up the hierarchy.

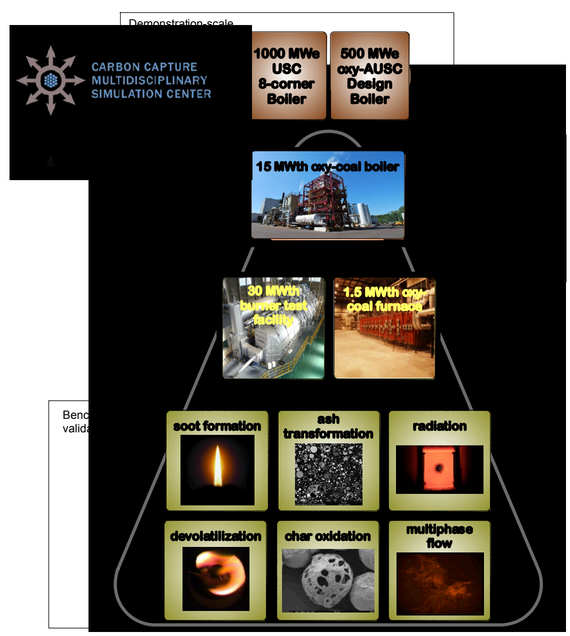

The uncertainty quantification methodology of the Center continues to be primarily based on the bound-to-bound (B2B) approach. Last year, the review team was unclear about the methodology and its efficacy. Considerable time was spent clarifying this methodology in this review and this significantly helped improve the review team's understanding. The B2B methodology continues to be a relatively unknown technique with limited publications and applications. Continued development in this area appears to be warranted with a focus on engagement with the broader community. Again, multiple examples of the UQ analysis impacting the understanding of the models and experiments were demonstrated. This shows how well integrated and valued V&V/UQ are in the Center's efforts.

Finding: Sensitivity analysis has been used quite effectively to redirect the Center's focus.

Observation 1: The Center presented clear evidence that sensitivity/UQ results inform CCMSC research reprioritization by shifting relative modeling focus from devolatilization to wall conductivity.

Observation 2: The B2B methodology continues to be relatively unknown with limited publication and peer review.

Recommendations: In the future we recommend considering the following:

1. Evaluate and publish clear and quantitative comparisons of the B2B methodology with more traditional techniques.
2. Continue to evolve the methodology to incorporate distributional uncertainty information when available.

Again, the Center presented an excellent example showing that (effective wall) thermal conductivity was a large source of uncertainty. It was also reassuring that the UQ method was able to identify faulty gas temperature measurements from the Alstom team. In the future, the review team would like to see more details on how computational errors are determined. What evidence does the Center have for mesh convergence, for instance? Furthermore, we would like additional details on how experimental uncertainties are computed. As the Center pointed out, probe gas temperature measurements are prone to error. In the future, it might be helpful if the Center expands on how 'instrument models' are used to account for these and related uncertainties.

The Center has its sights on a challenging problem requiring major computational resources. The Center's program shows conscious reflection of balancing technical efforts against budgetary constraints. The amount of realizable computational throughput, while acknowledged as a constraint, should be actively planned against with (major) risk mitigations identified.

Finding: UQ analysis has been performed in an effective manor and was even able to identify faulty experimental measurements as well as key sensitivity differences in the large scale and small scale experiments and models.

Observation: Some additional details on experimental and computational uncertainty calculations would be helpful.

Recommendations: In the future, the following additional details would be helpful:

1. Continue to improve the instrument models and elucidate the flow of uncertainties throughout the hierarchy of models and experments.
2. A risk mitigation plan for possibly limited computational availability should be considered.

Additional General Comments and Recommendations

As previously mentioned, the review team finds that this center has met or exceeded all of its Year 2 milestones, and has made excellent progress in all areas. Of all the PSAAP II centers, CCMSC has chosen to work on a problem that is truly Laboratory-like – involving an industrial problem, requiring a design prediction informed by UQ margins and experimentally validated "real world" models (not just realistic ones). This is a very difficult undertaking, and given the global importance of clean-coal combustion could have a significant impact.

The review team makes the following recommendations to further enhance the project's impact to the ASC Program and the robustness of the Center's approach.

1. Perform a "mid-term" update on your roadmap
    a. Clarify the predictions you will run year-to-year.
2. Prepare a flowchart showing how the variety of physics interact within a time step.
3. Continue to explore interaction opportunities with the other PSAAP Centers and labs.
4. Chart and measure your workflow from problem setup through analysis:
    a. What takes the time from asking a question to answering a question?
    b. Strategy for I/O integration.
    c. Data movement and resilience, e.g. SCR, HIO, IOSS.
5. Consider using a higher-order CFD method or document the adequacy of your low-order method.
6. Identify risk mitigation if needed cycles are not available.
    a. Develop a plan for dealing with the loss of INCITE cycles.
7. Engage Exascale communities on emerging standards.
8. Consider having a CCMSC team visit all three labs.
9. Take what you have learned from your GPU algorithms back to your CPU implementations.
10. Consider continuous recording during calculation (e.g. velocity field) to obtain characteristics of turbulent field.
11. Obtain mean profiles across the mixing layer to get a better sense of the flow field.
    a. Do this as a matter of course to collect valuable information.
    b. This can provide basis for comparison between your use of LES and RANS models your industrial partners are using.

12. For a small number of examples, compare and contrast B2B to other methods in a quantitative way and provide a balanced perspective.
13. Prioritize the 8-corner simulations versus the design boiler given the availability of experimental data.
14. Look to the broader community for tools for performance analysis and debugging.
15. We would like to see you become more strategic in addressing scaling and performance characteristics for your integrated application.
    a. Labs would find data you collect to be very valuable for co-design.
    b. What benefits do you see from system software innovations?

Acknowledgements

The Review team would like to thank the members of the ASC Alliance Strategy Team (AST) – R. Voigt (Leidos, Chair), D. Nikkel (LLNL), N. Hoffman (LANL), J. Costa (SNL) – and A. Macaluso (Leidos) for their insights and assistance during the course of the review. The review team would also like to thank the CCMSC personnel and leadership – P. J. Smith (PI), M. Berzins, D. W. Pershing (President, University of Utah) – for their informative presentations and warm hospitality.

Prepared by LLNL under Contract DE-AC52-07NA27344; LANL under Contract DE-AC52-06NA25396; and SNL under Contract DE-AC04-94AL85000.

References

[1] A. J. Chorin, "Numerical solution of the Navier-Stokes Equations," *Math. Comput.,* no. 22, pp. 745-762, 1968.

[2] J. C. H. G. H. Bell, "A second-order projection method for solving the incompressible Navier-Stokes equations," *J. Comp. Phys.,* no. 85, pp. 257-283, 1989.

[3] A. J. Chorin, "A Numerical Method for Solving Incompressible Flow Problems," *J. Comp. Phys.,* vol. 2, 1967.