LAWRENCE LIVERMORE NATIONAL LABORATORY

LLNL-TR-672409

ASC Predictive Science Academic Alliance Program (PSAAP) II Review of the Carbon Capture Multidisciplinary Science Center (CCMSC) at the University of Utah

C. H. Still, R. M. Ferencz, R. J. Hoekstra, A. L. Hungerford, A. L. Kuhl, D. R. Montoya, J. Wagner

June 8, 2015

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

ASC Predictive Science Academic Alliance Program (PSAAP) II

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 J. Wagner (SNL), and C. H. Still (LLNL, chair)

ABSTRACT

The review was conducted on March 31 – April 1, 2015 at the University of Utah. Overall the review team was impressed with the work presented and found that the CCMSC had met or exceeded all of their Year 1 milestones. Specific details, comments and recommendations are included in this document.

LLNL-TR-672409

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 31-April 1, 2015, to review the Carbon Capture Multidisciplinary Science Center (CCMSC) funded by the 2nd Predictive Science Academic 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:

- The degree of University support is very impressive, and attests to the trust, faith and commitment that the University places in this project.
- Engagement with the Defense Program Laboratories has been excellent.
    - The computer science deep dive success set a precedent for the other PSAAP II centers to follow.
    - Sandia engagement with a programming models review has resulted from the CS deep dive.
    - Student internships at the DP Labs have been very successful; the students have been excellent and their experiences have beneficial to all of the parties (student, Center, and Laboratory).
- The multidisciplinary integration across CCMSC is impressive, with all of the disparate skills becoming fully interwoven into one cohesive team.
- The inclusion of an industrial partner (Alstom) is a valuable component of the project, and will provide a dimension not present in any other PSAAP II center.
    - Engaging with Alstom on data uncertainties, or measured uncertainties, via continuous dialogue has benefitted both parties.
- The combination evolutionary/revolutionary computer science approach taken by this center is working well, and in particular enables agility in future years.
- The visualization capability is very high quality, and is making significant contributions to the project.
- Creating a three-site (University of Utah, Brigham Young University and the University of California – Berkeley) course on Validation and Uncertainty

Quantification will ensure knowledge capture and transfer of the methods and best practices established by this project.
- The idea of using crowd-sourcing for V&V data is innovative.
- The poster session, in which center associates and students presented their work, provided an excellent opportunity for in-depth interaction by the review team.

At the conclusion of the proceedings, the review team universally expressed that CCMSC has made excellent progress in Year 1, meeting or exceeding all Year 1 goals for milestone predictions. Comments related to the specific areas and recommendations are below.

Science and Engineering Research

Information regarding the progress and accomplishments from the CCMSC team members was organized into three broad categories: Computer Science (CS), Validation/Uncertainty Quantification (VUQ) and Science & Engineering (S&E). This section details a summary of the review panel comments regarding the latter category. The challenges this center faces from a CS and VUQ perspective share many similarities with the high performance computing challenges the review panel experts grapple with in their daily work. This is largely because those challenges arise from the grand-scale, multi-physics nature of the application, rather than the details of the physical conditions being probed. The S&E application concerns, however, are a bit further from the daily experience of the review panel. A foundational understanding in the fields of two-phase reactive flow and radiation transport is certainly shared, though what constitutes a best numerical approach depends sensitively on the detailed conditions in a coal boiler environment. The review panel recognizes this and commends the CCMSC team on a job very well done at bridging this gap through their overview of both the industry-wide challenges and unique numerical concerns arising from the conditions in a coal boiler environment.
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
Their simulation effort is based on large eddy simulation (LES) models of combustion of poly-disperse particles in turbulent flow. It solves the LES model equations of the gas phase[1]
- Conservation of mass
- Conservation of species
- Conservation of momentum
- Conservation of energy

The particle phase is simulated using the Direct Quadrature Method of Moments (DQMOM) and includes models for
- Heat transfer between gas and particles
- Combustion of coal particles

- Particle reactions: moisture evaporation, de-volatilization, gas-char reactions, soot formation, ash formation

Energy transfer due to (thermal) radiation is calculated currently using a Discrete Ordinates Method (DOM) and includes treatment for

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

The LES model of the turbulent combustion field of the gas is solved with a Projection Method first proposed by Chorin[4] and extended to second order by Bell et al[5]. In this method the Navier-Stokes equations are expanded in Mach number (M). Neglecting terms of order M^2 and higher, one can derive the low-Mach-number version of the conservation laws. The equations are hyperbolic but with a single characteristic: the particle velocity; in this way, sound waves are eliminated analytically. This allows one to explicitly follow the velocity field, but use time steps based on u rather than $u+a$; this speeds up the computations by a factor or 100 or more. According to the notes provided to us by Professor Thornock[1], their LES hydro algorithm, called ARCHES, is based on the Projection Method. It requires the solution of a Poisson equation for pressure; we believe they use the LLNL HYPRE library to solve the Poisson equation. So, it would be difficult to speed up the hydro algorithm (unless one could speed up the HYPRE solver).

The DQMOM method is used to advance the particle phase. It is based on a transport equation for the probability density function (PDF) for the particle velocity. It uses a moment-transformed quadrature-approximated NDF (number density function) transport equation. The method seems quite useful for this application; see Pedel *et al.*[2, 3] for more details. The DQMOM method will be investigated in future PSAAP II reviews.

The CCMSC team placed significant emphasis on results from a set of sensitivity studies. These studies explored the impact on boiler performance from the devolatilization model and from the heat conduction model used to assess energy flow out of the boiler walls. Contrary to their expectations, the devolatilization model was not a strong lever arm on the integrated results, while the impact from changing the conduction coefficient in the wall conduction model was significant. The team believes that an increased focus on understanding wall deposits from the combustion chamber will be an important step in bracketing a realistic treatment for the wall conduction coefficient. The panel recognizes that this is an important finding and expects future PSAAP II reviews will showcase progress towards an improved deposition model, both in the properties of the deposition and the radiative response to those properties. While the sensitivity to the devolatilization model was found to be smaller than expected, the body of work proposed to advance this modeling to a modified 2-step process seems to be a solid step in the direction of a more defensible capability.

Significant focus was also placed on the successes achieved in the development of a new radiation transport package, the Reverse Monte Carlo Radiation Transport method

(RMCRT). The RMCRT method[7, 8] was developed as a replacement for their current Discrete Ordinates method that models radiation transport in the combustion chamber. The method seems quite promising and offers distinct advantages over the DOM. Specifically, DOM involves multiple, sparse linear solves and comes with challenges for the incorporation of radiation physics such as scattering, as well as challenges in the use of parallel computers at very large scales. RMCRT offers ready solutions for handling the radiation scattering physics, as the microscale treatment of physics in Monte Carlo methods naturally allow for detailed scattering physics to be included. Year 1 code development by the CCMSC team identified and demonstrated how the RMCRT method could be employed in a scalable way for the coal boiler application. They have taken advantage of the Adaptive Mesh Refinement (AMR) infrastructure in Uintah to reduce the radiation mesh memory footprint. This reduced memory allows for a replication of the entire radiation domain to each node. Next steps for the team will need to include demonstrations that this AMR implementation has the accuracy needed for the radiation physics of their application.

The AMR is also used to follow turbulent mixing structures on the computational grid. Since one has the complete solution on each grid level, one can calculate flow average (0^{th}, 2^{nd} and higher moments) for each level—to demonstrate solution convergence with grid refinement. The CCMSC project should consider greater use of AMR—for example, one can allocate grid patches to flame sheets. This can save considerable amount of grid points—thereby reducing the computational time and storage requirements.

Finding: The AMR infrastructure provides a significant memory advantage where it can be employed, and in general will facilitate the overall scalability of the computational method.

Observation: If sufficiently accurate, the AMR infrastructure may be advantageous more broadly than just in the radiation physics.

Recommendations:
1. Consider directing some effort to demonstrating the AMR implementation meets the accuracy needs for the radiation physics of the application.
2. Consider greater use of AMR within the turbulent mixing calculation.

Computer Science Research

The Utah PSAAP II Center displays one of the tightest integrations between the CS and Physics efforts of any of the PSAAP II Centers. They work together to the point where it seems much more of a large team effort than singular domain goals. This is partially due to having an industry partner, Alstom, and an industrial goal to model carbon capture technology for pulverized coal power generation. This combination also puts them in a position to develop a capability that does not currently exist and provides a mission that is well understood and envisioned.

The Utah PSAAP Center CS effort was categorized as being both evolutionary and revolutionary. Many of the technologies that they have integrated into the CS solution could be considered revolutionary, including the Asynchronous DAG-based parallel runtime, and the key-value store data warehouse.

Other areas in their approach are based on more evolutionary technology developments than you may see in other parts of the community, including 2-side MPI inter-node communication and limited dynamism and granularity of the tasking layer.

Their approach is very 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 does introduce churn in that there will be continuous updating of production based software and environments to take advantage of the framework additions and also optimizations developed for a given architecture. These are done at the runtime layer and within the DSLs (these are different teams). It is not clear if there was a clear strategy regarding the level of integration at the different layers. There is clear evidence of progress with the runtime system regarding scalability and support for GPU accelerated architectures.

Finding: Interface points to the larger system environment are not being pursued.

Observation: The march toward exascale computing will require integration at all levels that will require community consensus and interaction.

Recommendations:

1. Suggest more effort be directed to exploring new technologies (beyond the “next platform”) and participate in community standardization efforts.
2. Tools are lacking for asynchronous approaches. CCMSC should engage with lab tool development efforts and leverage or articulate requirements, e.g. Open|SpeedShop, and other efforts to identify gaps and assess needs.
3. Engage communities in best practices and emerging standards.

There has been an excellent advance in supporting GPUs with the Nebo DSL. The approach with the DSL implementation is to continue to manage fine-grained parallelism, versus pushing more to the runtime system. This requires more effort in porting DSLs to new architectures to ensure performance with effort going to both DSLs and the runtime to extract performance. Many-core architectures have not been focused on because of initial performance issues with the architecture; we hope that there will be a refocus with new releases. Regarding integration of external solutions we applaud the willingness to look outside CCMSC for a solution to fill possible gaps with respect to performance portability, e.g. Kokkos

Finding: Architecture changes require a heavy investment to both runtime and DSL layers

Observation: Architecture integration points are divided between DSLs and runtime environment.

Recommendations:

1. Development of an overarching design/architecture for these layers with input and lessons learned from the community will have value.
2. Start investing in resilience.

Data Management is an area with much opportunity. The data analysis infrastructure is forward thinking and has excellent analysis capability. The I/O approach is of concern with the sheer number of files seeming unscalable in the long-term. A novel approach is the data warehouse implementation for tracking and bookkeeping status and scheduling.

Finding: Data Management infrastructure strategy not complete

Observation: Potential points that will impact increased scaling; number of output files and data warehouse constraints

Recommendations:

1. Consider rethinking I/O and data management from the simulation perspective.
    a. Can number of files and their size be reduced?
    b. It was not clear how much of the very large dataset is needed for direct comparison and analysis at the scientific comparison level.
    c. Current file size and numbers are not sustainable for archiving.
2. Data warehouse is similar to LANL MDHIM, LLNL CSToolkit and Sandia Kelpie.
3. This is a potential area for collaboration and possible topic for cross-cutting Deep Dive.

Algorithms

RMCRT has high potential to be scalable to exascale, assuming domain replication and/or multi-level approaches remain feasible. Potential efficiency and applicability of RMCRT is much better than linear solves in DO, as long as scalability can be maintained. Some form of hierarchical or multi-level approach is probably necessary.

Observation: Linear solvers are a big issue that the community as a whole needs to address.

1. Consider investing in this area because others may not deliver the solvers you need.
2. Continue dialog with this community with help from our labs.

A discussion item focused on whether the CS effort was doing enough to build for architectures and environments that were more than a few years out. The Utah CS response was that the uncertainty of architectures and environment more than 5 years out was a bounding point and it was more beneficial to target and optimize for known quantities in that near horizon to support their physics goals. This does work for them and is a good product goal related to what they are doing with Alstom – they have the appropriate mix for the uncertainty that is currently out there. We do recommend that they do more outreach with the runtime community to further establish the requirements for cross system interfaces, which will be more important on future systems and cannot be done alone.

The one area that seemed understated and possibly lacking was data management from the application side (the DAV side was well covered). It was noted that the volume of files was probably not sustainable (one file per processor), and they expected archiving problems. The addition of a Data Warehouse component, analogous to a key value store approach to the scheduling is novel and provides a starting point to longer-term data management solutions. This was similar in concept to key value store efforts at the labs, such as MDHIM (LANL), LLNL CSToolkit and Sandia Kelpie. This is a potential area of collaboration with the labs with extension to storage strategies.

In general the CS effort is taking a very pragmatic approach; it seems balanced for the state of uncertainty that exists today and their focus on a production project integrating physics, UQ, industry efforts and goals. Their programming environment and runtime is still somewhat an integrated system but more because there is no need to cleanly separate it at this point and the center charter is not a pure CS effort. Given that as a core goal and charter, which it currently isn't, they could broaden the impact of the programming model and runtime work that they are doing by contributing to the process to converge on a broader community direction for standards in these areas. Even with a limited focus on community convergence, as tighter integration of system resources and interfaces become a necessity, it may drive involvement in developing a broader runtime interface with the larger community and/or a greater ability to "mix and match" internal and external technologies through published APIs. This is an area that they should become active in developing.

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. If a formal Validation planning document has been completed, such as recommended by the ASME V&V Standards Committee, that would be a useful artifact to provide the review team prior to next year's review. Otherwise, we will draw our own inferences from sources such as the V&V Hierarchy.

V&V Hierarchy & Plans

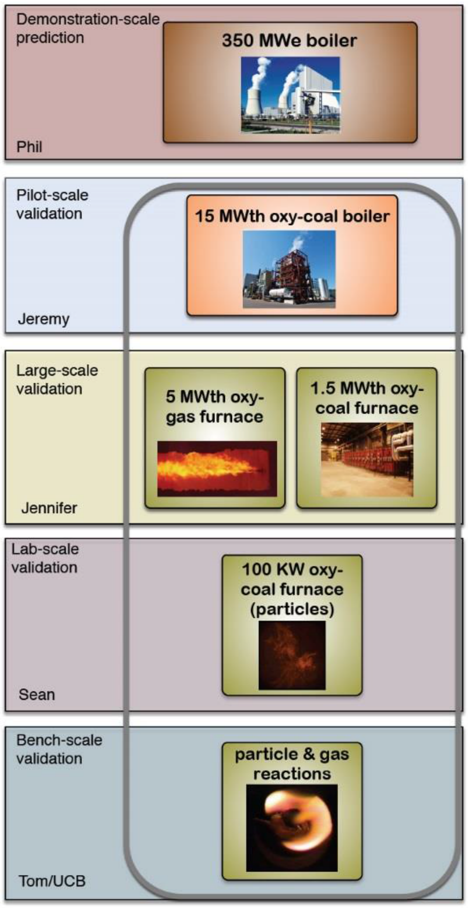

The team showed a validation hierarchy as its consistent focus. We reprise it here, as it has non-traditional aspects that are worth noting. Many levels of the hierarchy consist of integral tests at unique combustion facilities. This contrasts with the notional pyramidal hierarchy, where low levels typically isolate single-physics or a limited combination of physical phenomena. Center leadership could readily articulate that their hierarchy, although comprising many integral tests, and using closely related Quantities of Interest, manifests the dominant physics in different proportions owing to changing physical domain sizes. Conversations offered some indication that there are further bench-scale efforts that also support the overall investigation and validation. It would be good to accord these at least a brief acknowledgment or overview. We applaud the Center for identifying a hierarchy that is appropriate to their available experimental evidence. Validation should not be formulaic, but rather challenge researchers to think through an approach to best assemble a body of evidence characterizing the accuracy of the computational modeling and simulation. The validation data they intend to rely upon represent an interesting combination of 'home grown', industrial partner, and community sources.

At the review, there was a great focus placed on the simulations of the Alstom pilot-scale boiler. Wall heat-flux and a wide sampling of gas temperatures served as the primary validation here. On one hand, it is great that the primary validation case was based on the most challenging and complex experiment currently available. The fact that the Center was able to obtain many measurements from an industry partner, who may not have as much of an appreciation for science, makes the data even more valuable. On the other hand, given the size and experimental challenges associated with the 15 MW boiler, the data from such experiments will be inherently limited. The validation hierarchy clearly seems to address this concern with the lower-level experiments. However, the review team would like some more details as to how these experiments provide validation data and/or physical understanding to inform models. What specifically is to be learned from lower scale experiments?

Finding: The Center utilizes a validation hierarchy appropriate to their applications.

Observation: A succinct contrast with the traditional pyramid validation hierarchy can help provide context to people unfamiliar with the Center's activities.

Recommendations: In the future we recommend briefly outlining for each of the experiments the following:

1. What experiments or data analysis were performed and what is planned?
2. What quantities of interest were measured and how were the uncertainties for these determined?
3. How do these measurements feed into the overarching goals of the project? Since radiative heat transfer does not scale, for example, what other valuable information is being gleaned? What about wall-conductivity / conditions?
4. What type of data would be really useful for the modelers (e.g., non-intrusive temperature species measurements), but seems out-of-reach/ out-of-scope for the experimental team?

Uncertainty Quantification Methodology

The Center presented their uncertainty quantification method in a straightforward manner. In short, sets of model form and measurement uncertainties are bounded. The simulations were then run over a range of conditions corresponding to the model form uncertainties. The simulation was then considered valid over regions of model parameter space where the returned values for quantities of interest such as gas temperature and wall heat-flux overlapped the range of experimental data. By this method, the Center was able to demonstrate that the simulations were surprisingly insensitive to the particle burn model. The thermal conductivity of the wall, however, turned out to be quite important and will get much more attention in the future. *This was a refreshing example of how uncertainty quantification can be used to inform reprioritization by shifting relative focus from devolatilization to wall conductivity.*

The Center wishes to simulate and characterize the steady-state performance of advanced coal boilers. We would like more information on solution quantities that are being tracked to assess whether simulations from 'start-up' have been run long enough to establish sufficiently representative steady conditions.

At this stage of the project, it appears the Center is focusing on uncertainty of the input versus uncertainty of the system performance. This seems entirely appropriate at this early stage, though it would also be interesting to see how uncertainties affect the heat transfer into the water, which is the bottom-line performance metric for the project. Finally, at this early stage it is perhaps appropriate to use simple uncertainty bounds, though in the future, it might be beneficial to investigate results where uncertainty is propagated using statistical distributions (tails can be very important.)

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

Observation: At this time the center seems more focused more on input uncertainties rather than the output.

Recommendations: In the future we recommend considering the following:

1. What quantities are being tracked to ensure a steady state has been achieved?
2. Can the center compare a case or two where uncertainty is propagated using statistical distributions to their current approach of using bounds?

Sources and Magnitudes of Uncertainties/Errors

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.

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. How are experimental uncertainties calculated? What does the instrument model for the gas temperature data look like, for instance?
2. A few examples of grid convergence would be beneficial.
3. 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 1 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. (It should be noted that global impact is not required for a project to be successful as a PSAAP II center.)

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. Identify risk mitigation if the needed computing cycles are not available.
    a. Priority should be given first to the cycles required to meet the goals and deliverables associated with demonstrating the success of the CS approach, and performing the annually required integrated simulation.
    b. The size of the UQ margins may need to be adjusted, or the approach modified, if there are insufficient cycles available to meet the estimated need.
2. Assess whether evolving priorities require redistribution of CCMSC funds as the project moves into Y2.
    a. Does the distribution of resources fit the milestones and goals for the coming year?
3. Engage the exascale communities on emerging standards.
    a. CCMSC is building up experience in the areas of integrated multiphysics and large-scale computing that provide high value to efforts on emerging standards, and may very well be the difference between the selection of a successful approach and settling on one that is ultimately unworkable.
4. Consider investing in linear solvers because others may not deliver the solutions you need.
5. Continue to examine the differences between GPU approaches and manycore approaches.
6. Consider directing more CS effort to exploring new tchnologies (beyond the "next platform").
    a. Technologies present in the "next platform" may set the path for future follow-on architectures, but this seems unlikely given the complexity of issues and uncertainty in the currently-emerging hardware technology markets.
7. CCMSC should engage with Laboratory tool development efforts, e.g. Open | SpeedShop.
8. Start investing (at some level) in resilience to mitigate possible issues with next-generation platforms.
9. In future presentations, please provide more details about how experimental uncertainties are assessed (e.g., instrument models).

10. Some focus should be placed on how to transfer properties of the deposition to the wall heating model.
11. Explore using adaptivity to better resolve reaction flow regions and t0 de-refine in other regions.
12. Consider performing bench-top laboratory experiments to study wall deposition.
13. Upload the posters onto the CCMSC Wiki under the meeting heading.
14. Place links on the CCMSC Wiki to recent project-related papers.

Acknowledgements

The Review team would like to thank the members of the ASC Alliance Strategy Team (AST) – R. G. Voigt (Leidos), D. J Nikkel (LLNL), N. M. Hoffman (LANL), T. D. Blacker (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. Jeremy N. Thornock, *Current ARCHES LES Algorithm without DQMOM*, 2015 (private communication).
2. J. Pedel, J. N. Thornock, S. T. Smith, P. J. Smith, Large eddy simulation of polydisperse particles in turbulent coaxial jets using the direct quadrature method of moments, *Int. Journal of Multiphase Flow* **63** (2014) 23-38.
3. J. Pedel, J. N. Thornock, P. J. Smith, Ignition of co-axial turbulent diffusion oxy-coal jet flames: Experiments and simulations collaboration, *Combustion and Flame* **160** (2013) 1112-1128.
4. A. J. Chorin, Numerical solution of the Navier-Stokes equations, *Math. Comput*. **22** (1968) 745-762.
5. J. Bell, P. Colella, H. Glaz, A second-order projection method for solving the incompressible Navier-Stokes equations, *J. Comp Phys.* **85** (1989) 257-283.
6. R. Fan, D, L. Marchisio, R. O. Fox, Application of the direct quadrature method of moments to polydisperse gas–solid fluidized beds, *Powder Technology* **139** (2004) 7-20.

7. Xiaojin Sun, *Reverse Monte Carlo Ray-Tracing for Radiative Heat Transfer in Combustion Systems*, dissertation, U. Utah (2009), 173 pp.

8. A. Humphrey, T. Harman, M. Berzins, P. Smith "A Scalable Algorithm for Radiative Heat Transfer Using Reverse Monte Carlo Ray Tracing" *The International Supercomputing Conference* 2015 (ISC'15).

9. J. Bell, M. Berger, J. Saltzman, M. Welcome, Three-dimensional Adaptive Mesh Refinement for hyperbolic conservation laws, *SIAM J. Sci. Stat. Comput.* **15** (1994) 127-138.