Annual Technical Report

DE-NA0002375

University of Utah

The UQ-Predictive Multidisciplinary Simulation Center for High Efficiency Electric Power Generation with Carbon Capture

Carbon Capture Multidisciplinary Simulation Center

Philip J. Smith, Principal Investigator

philip.smith@utah.edu 801-585-3129

March 1, 2016 through February 28, 2017

May 31, 2017

Annual Technical Report

DE-NA0002375

March 1, 2016 through February 28, 2017

Carbon Capture Multidisciplinary Simulation Center

University of Utah – Institute for Clean and Secure Energy

The Carbon Capture Multidisciplinary Simulation Center (CCMSC) exists to demonstrate positive societal impact of extreme computing by accelerating deployment of low-cost, low-carbon energy solution for power generation: Advanced Ultra Super Critical (AUSC) Oxy-coal Technology. The overall strategy includes collaboration with industrial partners and interdisciplinary focus on development of technology. Three teams contribute to the overarching predictive design: the computer science team, the physics team and the validation/ UQ team. The center is partnered with General Electric Power.

Outreach and Education

The V/UQ faculty of the University of Utah and the University of California-Berkeley taught V/ UQ course for the second time in the fall of 2016. The course is targeted for graduate students and computer simulation and engineering practitioners. Last year the students enrolled in the course were mostly students from our PSAAP center. While this year, the majority of the students were from other disciplines and research activities. There were a large number of students from nuclear engineering research areas. This participation broadened the applications used as examples throughout the semester.

The course, entitled "Modeling/Validation UQ," covered (1) Introduction and Motivation for V/ UQ, (2) Semi-Definite Programming, (3) Surrogate Modeling, (4) Experimental Uncertainty, (5) Dimensionality Reduction, (6) Kennedy O'Hagen Analysis, (7) MCMC Sampling, (8) Bounds-to-Bounds Analysis and (9) Practical Workflow. Lessons learned in this first course were addressed in designs for the fall 2016 course.

Next year, the course will not follow the academic calendar, rather a new approach will be implemented to reach a broader and more technically diverse audience. A Modeling Validation through Uncertainty Quantification Short Course will be held in Park City, Utah in February 2018. The intensive short course is being advertised to national lab personnel and all of the PSAAP II centers. University students will be invited to participate as well. We hope to address a technically diverse audience on an application level, thus bringing our research developments and understanding to a wide audience.

Key meetings and conferences attended by members of CCMSC include: Combustion Institute (Western States Section) Meeting, Clearwater Coal Conference, WEST Workshops, PSAAP IPDPS Meeting in Chicago, GPU 2016 Conference, Radiation Deep Dive Workshop (hosted by Texas

A&M University) and Multiphase Flow Deep Dive (hosted by University of Florida), American Flame Research Committee Conference, Kokkos Technical Review that was held in Albuquerque, IEEE VIS Conference, and SIAM CSE Conference.

To facilitate more integration with CCMSC, and after group deliberation of communication tools available, a core subset of the Center has adopted #Slack as a team communication and collaboration tool. As technical problems are discussed and progress is made between individuals on a more fine-grain scale, #Slack has allowed others in the Center to be made aware of both the issues and accomplishments. It has particularly helped the remote collaborations. Coordination with collaborators at BYU has been quite active, with as-needed discussions between campuses occurring nearly weekly. The collaboration with UC-Berkeley, in particular, has been tightened by the use of #Slack.

Four University of Utah graduate students completed internships during year-three: namely, Teri Draper (1/23 – 4/16/16) with Scott Skeen at Sandia National Lab – Livermore; Oscar Diaz-Ibarra (2/13 – 4/23/16) with Chris Shaddix at Sandia National Lab – Livermore; John Holmen (5/16 – 8/12/16) with Jonathan Hu and Ray Tuminaro at Sandia National Lab - Livermore; and Joshua McConnell (8/22 – 10/8/16) with Stefan Domino and Paul Krosier at Sandia National Lab – Albuquerque. Additionally, Alex Josephson (BYU) completed his internship at Los Alamos National Lab with Rodman Linn.

Professor Tom Fletcher made a research visit to Sandia National Laboratory in Albuquerque to discuss fire research with Alex Brown during August 2016. He also visited Joel Kress, Troy Holland (student) and Rod Linn at Los Alamos. Professor Lignell has planned a research visit with John Hewson at Sandia National Lab and Rod Linn at Los Alamos in July 2017.

Computer Science:

Runtime System and Infrastructure and I/O

The Runtime System and Infrastructure efforts for year 3 were focused on integrating Kokkos and demonstrating scaling through by using the computational resources provided by a new large INCITE Award.

The major accomplishments of this year is the development within the Uintah infrastructure for its asynchronous, many-task runtime system to efficiently handle globally coupled problems involving radiation.

The initial and principal improvement regarding development within Uintah has been the adoption of the C++11 standard, upon which nearly all other infrastructure improvements have been derived. With this adoption, we have removed thousands of lines of code previously used to provide functionality that now exists in the standard library, specifically from the Pthread library, for synchronization primitives, atomics and other concurrency offerings. This has also allowed for the development of novel lock-free data structures to handle MPI communication records and replace previous task queues. With these changes, we now have a portable approach to multi-threading within Uintah without having to maintain complex and error-prone code.

In our production case, the radiation calculation is performed roughly every 20 timesteps. With the automated MPI message generation Uintah provides, our previous strategy to avoid the heavy communication incurred by the global halo requirement in the RMCRT algorithm, has

been to simply recompile the task-graph on radiation timesteps. This is not a tenable solution in terms of the cost of recompilation. Profiling our code revealed that Uintah's task-graph compilation algorithm was performing an extra O(N2) search for patches on the fine, CFD mesh when constructing lists of neighboring patches for local halo exchange. The cost of this operation grew despite the number of patches per node remaining constant, resulting in untenable task-graph compilation times at large scale with high patch counts

For each compute node, Uintah generates a local task graph for tasks residing on patches owned by that node, and the resulting data dependencies for automated MPI message generation. The Uintah load balancer then creates a "processor neighborhood" for halo exchange. Previously a naive approach was used, so that the task with the maximum amount of halo layers designated the halo length for the neighborhood. For the target production problem, this naive approach was no longer viable for three reasons, 1) a large halo number due to the global nature of radiation calculations, 2) non-uniform halo requirements across AMR mesh levels, and 3) applying this large halo number to almost 1000 tasks per timestep. An approach was implemented which looks at halo lengths on a per variable and per mesh-level basis. Our solution has been to split tasks into two halo neighborhoods (one for local halo exchange and another for the potentially global halo requirement). This has resulted in reducing task-graph compile times from 4.5 hours to roughly 20 minutes at 128K cores for a 2-level mesh problem. Work is currently underway to reduce this compile time down to about 2-3 minutes.

Uintah threaded scheduler works by having all CPU threads first checking which task is available to process, and then performing work for that task. Part of the infrastructure duties involve MPI halo sends and receives. To ensure no two CPU threads performed the same send or receive, these were managed in a mutex-protected priority queue. When we moved to larger problems with hundreds of tasks and hundreds of thousands of cores, these locks began to noticeably effect performance. A novel lock free pool was developed and implemented to eliminate contention related to acquiring this mutex. As the name suggests, it now allows all CPU threads to concurrently processes MPI sends and receives in a fully lock and contention free manner.

Initially the target production problem was not able to fit into the available memory per node on Titan. As part of wider memory reduction efforts, we discovered that the MPI packed buffers had a quick memory buildup and slow release, resulting in substantial memory overhead. We subsequently modified the infrastructure to avoid packed buffers and instead treated multiple items in a message as simply offsets of an array. This resulted in a significant reduction in memory usage.

The RMCRT algorithm functions by randomly selecting rays to process. The Latin Hyper Cube sampling algorithm allows for fewer rays to be selected for the same accuracy. Previously this was implemented on the CPU, but not fully on the GPU. The greatest challenge here was not duplicating logic for GPU code, but rather having Uintah providing proper random number seeding, so that each GPU thread can have its own unique seed. A solution was adopted so that potentially millions of unique seeds can be generated for potentially millions of GPU threads. The infrastructure performs most of the work, avoiding burdening task developers with this responsibility. Moving forward Kokkos will be an ideal solution, as it both allows for a single portable codebase and also a single random number generation engine, instead of maintaining separate GPU and CPU implementations.

In summary, these and other infrastructure improvements have yielded the improvements shown in Table 1 below.

Operation	Before	After
Nodal Memory footprint (GB)	21	2.5
Timestep (sec)	10-11	~3
Radiation Solve (sec)	60+	50-55
Pressure Solve (sec)	2.5 - 5	1 – 2.3
Task wait comm time (sec)	3.25	0.065
Taskgraph compile	4.5 hours	20 minutes*

Table 1: Summary of infrastructure improvements and optimizations

Visualization and I/O

During the past year, we focused on more tightly integrating Uintah for in-situ visualization and analysis. This integration allows for better coupling of Uintah to VisIt so that more of Uintah is exposed to the user at runtime. An example of this integration work was to rewrite the Uintah simulation controller so that Uintah and VisIt can work more cohesively with initial and re-start data thus allowing for exploration of the initial or re-start conditions before advancing to the first time-step. The re-writting of the simulation controller also lead to improvements to run time performance measures which are especially important as new components such as the PIDX I/O are incorporated. All of this integration has cumulated into an in-situ simulation dashboard.

The in-situ simulation dashboard focuses on three areas; visual debugging, run-time performance and simulation monitoring. Application scientists are using the in-situ tools to fine tune parameters on smaller test cases before running larger cases on DOE machines. Being able to obtain real-time feedback to parameters changes is helping accelerate the science. At the same time, the in-situ tools are being used to explore new runtime performance measures that are patch rather than rank base (there can be one or more patches per rank), which may help in better load balancing and task management. As threading becomes more prevalent, we will look at node based runtime performance measures as well.

We continued to support of VisIt on various platforms while providing support to large runs, such as those that were part of the INCITE awards. INCITE had unique data requirements in terms of structure and size. The work was presented at the SIAM CSE17 Workflow Mini-symposium. Also, we continued working on a method for compression on the GPU with ZFP. The CUDA encoder and decoder is up to order of magnitude faster than the CPU version. We have tested this against several CCMSC datasets and a paper is currently being revised for submission.

We have been continuously working on improving Volume Rendering in VisIt. Part of this effort was geared towards speeding up parallel volume rendering on distributed memory machines. The current volume rendering in VisIt (Ray casting: compositing) uses a sort-middle approach to parallel rendering. We have completed working on providing a sort-last parallel rendering solution, Ray casting: SLIVR, which should offer the same quality as Ray casting: composting - Trilinear (that we added to VisIt some years back) since both use trilinear interpolation for sampling, but should be faster. Some initial code for Ray casting: SLIVR had been committed to VisIt and recently a bug, which caused gaps to be visible, has been fixed and committed to the VisIt repository. It is available in the current release.

We completed work on tightly integrating the output of Ray casting: SLIVR with the VisIt interface. While initially the rendered output of Ray casting, SLIVR would overwrite the bounding box and other plot integrations. Now they can live together as shown in the figure below. The effort also involves some major refactoring of our code and developing a faster compositing algorithm. This has been committed to the VisIt trunk and is available in the current release. We have initiated investigation of integrating the Intel OSPRay renderer into VisIt for use on KNL HPC platforms.

Heterogeneous Computing, Performance & Scalability of CFD in the DSL

In addition to improving the robustness and range of applicability of the low-Mach algorithm used in Wasatch, the Domain Specific Language (DSL) prototype, we have performed scaling studies on the low-Mach and compressible algorithms used within Wasatch on Mira and Titan.

On Mira, the low Mach algorithm shows good weak scaling up to 524,188 cores, with scalability limited by the linear solver at smaller patch sizes. For the compressible algorithm on Mira, good weak scaling was observed using patch sizes as small as 64x64x128 for core counts up to 524,188. Poor scaling at smaller patch sizes was due to an MPI reduction of the time-step. The scalability will improve as the reaction modeling is implemented since this will significantly increase arithmetic intensity.

We have demonstrated performance and scalability of full CFD on GPUs, scaling up to 18,000 GPUs on Titan. These calculations, running a compressible CFD algorithm, are run completely on GPU. We have also shown that running a basic low-mach algorithm with only the pressure poisson solve on the CPU and everything else on the GPU does not provide a speedup due to data transfer latencies. However, this may be less problematic once full reacting flow calculations are incorporated since that increases arithmetic intensity. The scaling studies on Titan using CPUs and GPUs to compare the low-Mach and compressible algorithms identified critical gaps in the Uintah infrastructure that exist for GPU execution of tasks typical in CFD algorithms. We also characterized the impact that a CPU-based linear solver has on performance for low-mach algorithms that are otherwise fully deployed on GPU. This is motivating exploration of alternative algorithms (point-implicit) that do not rely on global linear solves and are, therefore, more amenable to GPU implementation. We have also added artificial compressibility capabilities to Wasatch to increase the stable time-step for the compressible algorithm.

We have demonstrated particle-cell interpolants that provide significant speedups on GPU, and these are supported within Wasatch. Additionally, we have performed extensive testing and performance profiling of Kokkos integration in Nebo, showing performance bottlenecks and

gains in various use cases relative to native Nebo implementations on both GPU and multicore systems. We have implemented GPU support for TabProps (tabulation of gas phase properties and chemistry calculations) and RadProps (tabulation of radiation properties for grey gas mixtures), libraries for are now available, providing very significant speedups (10x-80x) over the CPU counterparts.

Physics: LES, Multiphase Flow, Particle Combustion, Radiation

Large Eddy Simulation Integration

Computational efficiency for the production, CPU-only production Arches code was increased significantly over the past year. Much of the efficiency gains were obtained by identifying algorithmic bottlenecks and abstractions that did not scale to the large numbers of cores required for the INCITE production cases (~256K cores). This work resulted in a 63% efficiency gain over the previous code at large scale. This work was crucial in enabling the prediction calculations of the 8-corner unit.

The longer term development path for Arches, as decided on the outset of Year 3, is the incorporation of Kokkos coding constructs, abstracted as Uintah::parallel_for function calls, to provide hardware portability and execution performance. In Year 3, the adoption of Kokkos into Arches continued, with the year ending with a demonstration of a single phase fluid solve with spatial and temporal stencil operations constructed within the Uintah::parallel_for construct with a boundary condition implementation. The pressure linear system (i.e., pressure projection) was also constructed within the Uintah::parallel_for, but the solution procedure was transferred to the HYPRE library to obtain the solution to the global solve. The demonstration represented core CFD elements to enable a full physics solution that fully incorporates a Kokkos back-end. Code verification of these core components is ongoing and projected to finish within the first part of Year 4. Additionally, all the coal physics were wrapped with the Uintah::parallel_for and used for all production runs. Finally, a strategy for performing tabulated property look-up during runtime was developed using Kokkos objects.

Besides the adoption of the Kokkos back-end, Arches adopted a lightweight abstraction layer to sit between the physics components and the UCF. This abstraction layer was implemented in Year 1. Initially, the motivation of the Arches Task abstraction was to reduce UCF boilerplate for physics developers and aid in the addition of new physics within Arches LES algorithm. The interface evolved over Year 3, increasing its robustness and including the ability to combine once separate Uintah tasks into a super-task-set, all the while maintaining a fine granularity on the physics and their implementation for developers. This allows for two advantages moving forward; 1) Lightweight tasks can combine with heavier task to enable better hardware utilization and 2) The work combining for common sets of physics will allow for implicit solves. The later will allow for tighter physical coupling in the case of the coal combustion modeling (devolatilization, char-oxidation, etc) while the former will offer better use of computational resources.

The LES integration team also contributed to the physics fidelity of the program including the incorporation of new turbulence closure models (Sigma Model, NLES) as alternatives to the stock dynamic Smagorinsky closure model. A method for measuring the “LES quality

index" (LESQI) for production runs was implemented, which includes an estimation of the percent-resolved-turbulent-energy locally. This LESQI provides a more quantitative number for determining if the mesh resolution is sufficiently fine. Solution verification methods were also explored for the full-physics production cases to identify errors resulting from numerical uncertainty. Recommendations have been made, but due to the computational cost of the methods arising for grid-resolution/Richardson extrapolation techniques, the solution verification techniques are still being investigated. Finally, an analysis was performed to test the hypothesis that subgrid turbulence/particle interactions may be neglected in the coal problem with a sufficiently fine mesh and a sufficiently high Stokes number particles. The initial analysis performed seems to support the hypothesis but later analysis, which included more physical effects, raised some doubt. This work is continuing into Year 4.

Coal / Ash Deposition

It was shown at the end of our year 2 in top-down sensitivity analysis that ash deposition on the walls of the boiler has significant influence on and coupling from the combustion process. Being a boiler, it should be no surprise that the primary concern is the heat transfer effects — and since radiation dominates, modeling is focused on the surface temperature and emissivity. Determination of these values is itself a multiphysics problem requiring the deposition rate, a quasi-steady multilayered one-dimensional wall heat-transfer model, a sintering model, and an emissivity model. While the one-dimensional wall heat-transfer model had been developed and refined in previous years, significant progress has been made this year in the creation of a deposition model, a model for ash thermal conductivity (used in the one-dimensional heat-transfer model), a sintering model, and an ash-emissivity model (based on sintering).

The deposition rate depends on the wall temperature & viscosity as well as the flux of particles to the wall along with their properties. The Arches deposition model combines the well-established approaches of Brink et al.[1] & Walsh et al.[2] by using a sticking probability with the thermodynamic properties of the ash. One interesting requirement is that the viscosity of the incoming particles has been shown in the literature to be better modeled using the maximum temperature, historically experienced by the particle rather than its current temperature. As such, a new particle property was added to the code for transporting maximum particle temperature.

The model for the deposit thermal conductivity is not entirely novel since there are sufficiently accurate models in the literature for our specific situation.The approach starts with a solid thermal conductivity (which can be measured directly) and an initial porosity. For long sintering times (which the bulk of the wall experiences), the porosity is effectively a step function in temperature — dropping to zero porosity at a specified value.

---

[1] A. Brink, D. Lindberg, M. Hupa, M. E. de Tejada, M. Paneru, J. Maier, G. Scheffknecht, A. Pranzitelli, and M. Pourkashanian, A temperature-history based model for the sticking probability of impacting pulverized coal ash particles, Fuel Processing Tech., 141 (2016) 210–215.

[2] P. M.Walsh, A. N.Sayre, D. O.Loehden, L. S.Monroe, J. M.Beer, and A. F.Sarofim, Deposition of bituminous coal ash on an isolated heat exchanger tube effects of coal properties on deposit growth, Prog. Energy Combust. Sci., 16 (1990) 327–346.

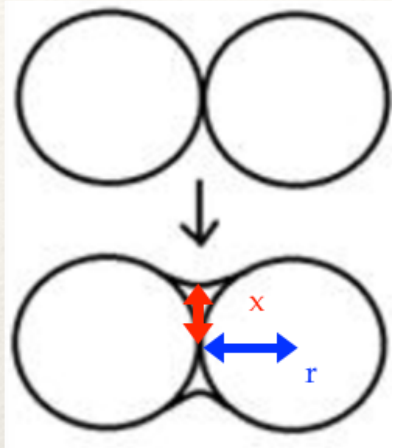

Figure 1: Schematic of deposition and sintering process.

The surface sintering model is novel, but follows the recommendation in the literature by TF Wall[3] (a well-established ash deposit expert in pulverized coal). Additionally, certain aspects of the model rely on approaches proven in ceramics literature. The core of the approach has, in this last year, relied on the Frenkel sintering model.[4]

However, alternative sintering models such as Pokluda et al.[5] have been considered. The core sintering model relies, in turn, on models for the surface tension, ash viscosity and sintering timescale. The surface tension is provided from the ceramics literature. The current choice for the viscosity is the modified Urbain[6] model. And since a simple radiation calculation shows that only the top layer of particles influence the value of emissivity, the sintering timescale is modeled as the deposition particle diameter divided by the deposition rate.

The emissivity model, in conception, is based on Mie theory. In practice only two limiting cases are needed: the Fresnel relation, and the Bohren[7] two-stream model. Mie theory, however, is sensitive to the optical constants which are particularly tricky for complex coal ashes. Significant work on ash optical properties has been reported by Ebert & Goodwin[8], and the results of these studies are used directly in the Arches emissivity model.

Combining these models and comparing to data from Boow & Goard[9], the Arches models for sintering & emissivity behave as shown in Figure 2.

---

[3] Wall TF, Bhattacharya SP, Zhang DK, Gupta RP, He X, "The Properties and Thermal Effects of Ash Deposits in Coal-Fired Furnaces," Prog. Energy Combust. Sci., 19 487-504 (1993).

[4] Frenkel J, "Viscous flow of crystalline bodies under the action of surface tension," J. Phys., 4: 385–431. (1945).

[5] Pokluda O, Bellehumeur CT, Vlachopoulos J, "Modification of Frenkel's model for sintering," AIChE J., 43-12 (1997), doi:10.1002/aic.690431213.

[6] G. Urbain, Viscosity estimation of slags, steel research international 58 (1987) 111–116.

[7] Bohren CF, "Multiple scattering of light and some of its observable consequences," American J. Phys., 55, 524 (1987); doi: 10.1119/1.15109.

[8] Ebert JL, "Infrared Optical Properties of Coal Slag at High Temperatures," Ph.D. Dissertation, Stanford University, August 1994.
Goodwin DG, "Infrared Optical Constants of Coal," Ph.D. Dissertation, Stanford University, July 1986.

[9] Ebert JL, "Infrared Optical Properties of Coal Slag at High Temperatures," Ph.D. Dissertation, Stanford University, August 1994.
Goodwin DG, "Infrared Optical Constants of Coal," Ph.D. Dissertation, Stanford University, July 1986.

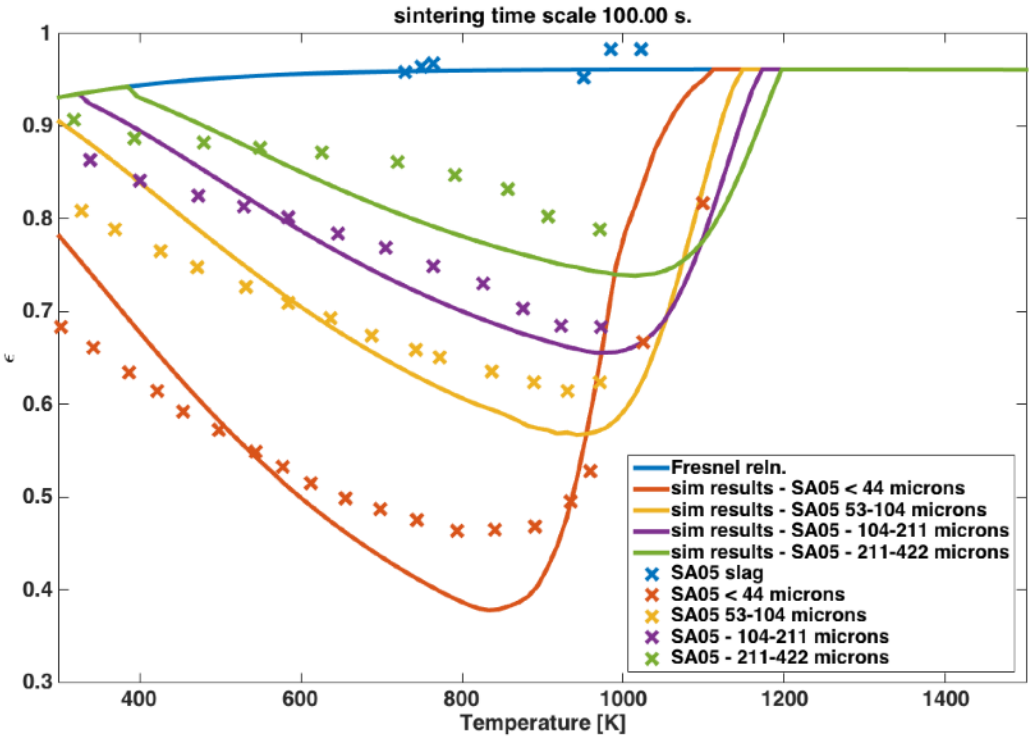

Figure 2: Comparison of simulation results (solid lines) and measurements (symbols).

While these results are quite promising, the experimental data did not include a sintering timescale. We are continuing to scour the literature for better validation data to quantifiably reduce our uncertainty.

Coal Devolatilization

The analysis of the performance of simple, coal devolatilization models at different heating rates vs. the CPD model predictions was completed and published. The fitting procedure can be used to determine devolatilization rates for any coal, and can also be used to generate a simple model for tar release rates. Since the CPD model describes pressure effects, this procedure could also be used to generate coefficients at pressure. In addition, lab experiments were performed in the BYU flat-flame burner on Sufco coal devolatilization to help in the VUQ analysis for that coal.

One of the current assumptions in the ARCHES simulation is that any gas generated from the coal has the same composition, so that one coal gas mixture fraction can be used. It is well known that the volatiles are enriched in hydrogen and the char is enriched in carbon. We have been gathering data from the literature on elemental compositions of char and tar. These data will be used to develop a correlation that will describe these elemental compositions as a function of coal type and pyrolysis conditions. The data fitting for this correlation will use VUQ, possibly in conjunction with Dr. Habib at Sandia Livermore as an internship for Andrew Richards.

Coal Char Oxidation and Gasification

The previous sensitivity analysis on the advanced char combustion model (CCK/oxy) showed that annealing was the most sensitive rate besides the surface oxidation and gasification kinetics. The annealing model used in the best previous char oxidation model was based on just a few data sets that were available in the 1980's. An extensive literature review was conducted, and many more data sets on annealing behavior were found. Troy Holland used the Bayseian VUQ codes at Los Alamos National Lab to develop an improved char annealing rate model, and incorporated this model into the CCK/oxy code. The CCK/oxy code was combined with the CPD model (Fletcher et al., 1992) and a new swelling model (Shurtz et al., 2011). The resulting combined particle reaction model was used to explain the char oxidation data of Shaddix under oxy-fuel conditions (Shaddix and Molina, 2009; Geier et al., 2012). Particle diameter was found to be critical to explaining these data, along with the kinetics and the annealing process. Figure 3 shows the comparison of the observed and predicted particle temperatures for four coals.

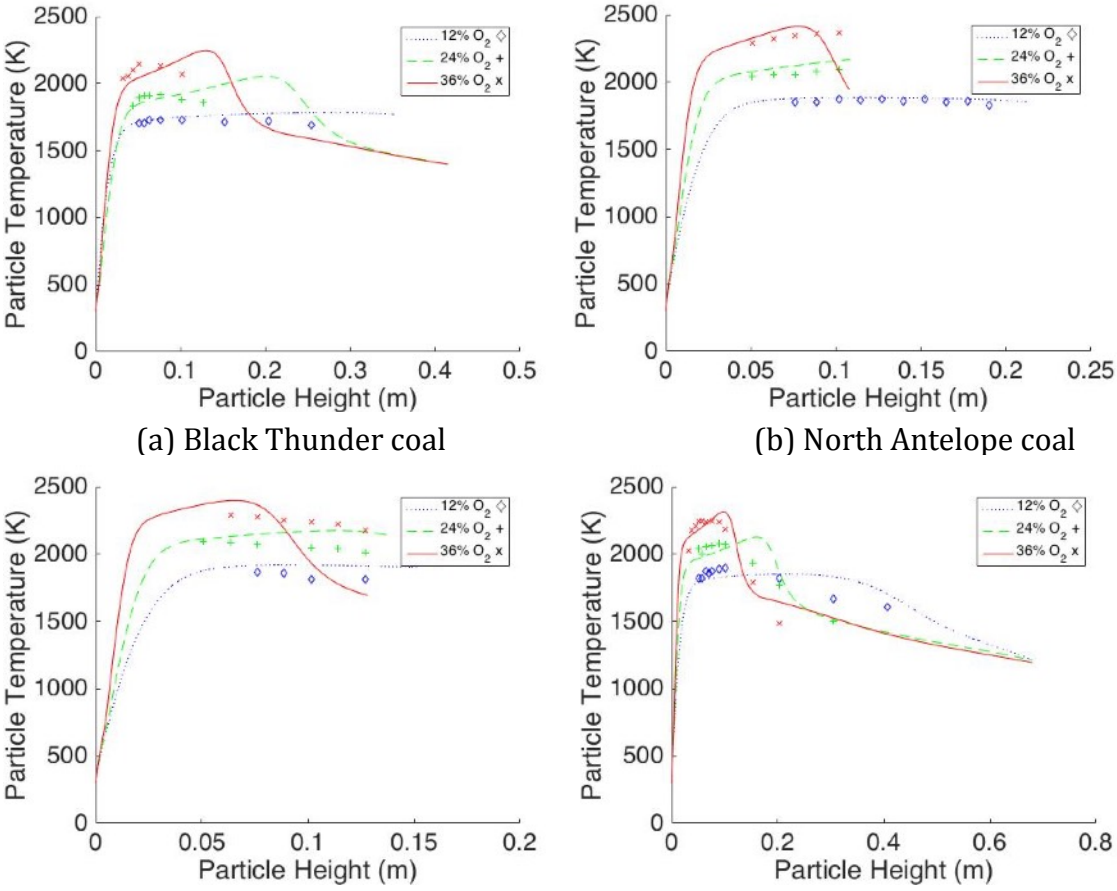

Figure 3. Comparison of CCK/oxy model calculations with coal data from Shaddix and coworkers (Shaddix and Molina, 2009; Geier et al., 2012) using the measured particle diameters.

Adopting a multiscale modeling approach, we have employed the one-dimensional turbulence model to identify model pairings that provide accurate descriptions of flame standoff following devolatilization as well as char oxidation/gasification. An assessment assessment of the predictive capability of the high-fidelity Chemical Percolation Devolatilization (CPD) and a simple two-step (TS) devolatilization models and the high-fidelity Char Conversion Kinetics (CCK) and a simpler nth-order Langmuir-Hinshelwood (LH) char chemistry models was conducted. Sensitivity to furnace temperature and initial particle diameter was investigated in the study. Results indicated that the predictions for particle temperature and particle mass evolution are strongly dependent on the initial particle diameter for both CCK and LH but are weakly dependent on the initial furnace temperature for both CCK and LH. The choice of devolatilization model has a strong impact on particle mass histories but only somewhat impacts accuracy of particle temperature calculations. The chosen devolatilization model has a noticeable effect on the CO chemistry in the gas phase, even after the devolatilization process has ended. Furthermore, a significant overlap in char oxidation and devolatilization is predicted by all char and devolatilization model combinations considered. In general, CCK yields more accurate predictions of the particle mass than LH, although CCK fails to accurately predict DAF particle mass at elevated oxygen concentrations for low-carbon coals. This is most likely because correlations for kinetic parameters relied upon by CCK are based on data that are sparse for low-carbon coals. Neither CCK nor LH perform well at predicting particle temperature through char oxidation. We have also implemented the spherical harmonics, discrete ordinates and reverse Monte-Carlo ray tracing radiation algorithms within Wasatch.

Soot in Coal Flames

The soot formation and oxidation model of Brown and Fletcher (1998) was previously implemented into the Arches code. A new soot gasification model was derived with parameters tuned to data gathered from six different published experiments in the literature. The new soot gasification model includes gasification by both CO2 and H2O, which is important in oxy-coal combustion. Conservation equations for soot and tar were coded into the pressure solvers used by the input files in the Arches software. Parameters for both oxidation and gasification models were tuned using Bayesian methods, allowing for better quantification of uncertainty. Probability density functions were derived for each model parameter. Use of the soot gasification model eliminates the small amounts of soot predicted by the base soot model, but not experimentally-observed, in fuel-lean regions in oxy-coal simulations of lab-scale furnaces. The soot model of formation from coal tar, agglomeration, oxidation, and gasification was coupled with the radiation model implemented in the Arches code.

A more detailed soot model has been derived which includes: soot nucleation from tar and light gases, surface growth of soot particles using the HACA (hydrogen abstraction carbon addition) mechanism, and soot consumption. The new detailed model uses MoMIC for soot predictions and a sectional method for soot precursor (tar and PAH) predictions. The model was implemented in Python for further validation testing and development. We have been analyzing existing BYU soot measurements from coal pyrolysis experiments, including six coal types of different ranks at three temperature profiles using a flat flame burner. We have now reproduced the original curve fits with the previously-published empirical model. This data set is being used for validation of the new physics-based coal model.

We developed and implemented in Arches a new method for accurately treating stiff chemistry within the explicit solution advancement approach. This method uses analytic solutions to determine rates for use in the explicit solver. These rates allow exact solution for any step size. This was implemented for the soot formation rate from tar, which was limiting. We also demonstrated the method for the full nonlinear coupled soot reaction system. This was done using analytic solution of the linearized equations. In that case, the stiffness and stepsize constraints were significantly reduced, though not eliminated, while remaining much more accurate than the explicit counterpart.

Radiation

Computation of the gas phase radiation properties in Arches was reviewed. A hybrid approach of WSGG and Hottel-Sarofim charts was the primary method in use. An improved approach using the HITEMP spectral database was formulated. Tools for the computation of mean or effective absorption coefficients has been developed.

To improve the performance of the radiation linear-solve for the INCITE runs, new radiation solver options were added for use with discrete-ordinates. This included using a more efficient use of the Hypre library and the shedding of some legacy Fortran code.

RMCRT validation runs have been conducted, and show good agreement with experimental data for the radiative heat fluxes to the wall and gas temperatures. The Legendre-Chebyshev spherical quadrature set has been added for use in discrete ordinates and shows a ~40% increase in accuracy relative to the level-symmetric quadrature set. Reformulating the discrete ordinates equation in a second derivative form was investigated but showed similar performance and accuracy to our current discrete ordinates method.

The discrete-ordinate branch sweeping branch has been developed further and merged into the Uintah trunk. The method has shown excellent performance exhibiting 100x to 3x faster performance than the linear solver counterpart, depending on optical thickness of the system, the effectiveness of the initial guess for the linear solver, and scale of the problem. Detailed scaling data has not been generated at this point.

During year-3, a primary goal was to utilize the Reverse Monte Carlo Ray Tracing (RMCRT) model in a production size simulation. This required a significant amount of development, especially with regards to the GPU infrastructure. We have spent years developing RMCRT for the CPU and GPU, testing the framework and algorithm with simplistic scenarios that only ran for a few time-steps. Once we started running realistic production calculations, with a large number of time integration steps, we quickly encountered memory leaks/fragmentation issues in addition to excessive and expensive communication on non-radiation calculation time-steps. The RMCRT model requires all-to-all communication of the radiative properties, which is the single biggest drawback of the algorithm. These problems were addressed, but it took a major effort by the computing team for a large fraction of the year.

Using the INCITE award on Titan (71M SUs), we performed a suite of verification tests for both the CPU and GPU versions of RMCRT on 8k, 16k, 32k, and 128K cores. For all tests a slightly modified version of the 8-corner prediction case was used. The quantities of interest (divergence of heat flux and wall heat flux) were compared against simulations that used the Discrete Ordinates radiation model. The qualitative agreement of the QOIs was very good,

giving us confidence in the methodology. Quantitative analysis of the QOIs for the different test conditions considered is currently underway.

As a way to lower the costs of RMCRT, we investigated the impact of varying the number of rays per cell (rpc) and the distance a ray would travel before moving to a coarser grid on a different level (halo distance). We ran a simulation with a high number of rpc and long halo distance and compared the QOIs against the same simulation with a low number of rpc and shorter halo distance. We used the 8-corner prediction case for these tests. The preliminary analysis showed a significant reduction in the computation time without a dramatic loss in accuracy in the QOI's. Overall, the results from the INCITE runs have demonstrated that RMCRT is a viable and competitive radiation model.

Validation, Verification and Uncertainty Quantification

Physics – VUQ integration

Based on year-2 V/UQ analysis, we found an increased importance in thermal boundary conditions. The top-down V/UQ analysis of the 8-corner echoed that analysis is showing high sensitivity to thermal resistance and emissivity on the walls. Given what was learned, much of the work this year consisted in improving the thermal boundary conditions. This included the incorporation of an ash deposition model, an ash thermal conductivity model and the creation of an ash emissivity model. Initial simulations of the 8-corner unit demonstrated numerical dispersion issues related to the momentum discretization. We implemented a hybrid discretization scheme which dissipates un-physical Reynold's numbers. In addition, the char-oxidation model was updated, devolatilization swelling and particle shrinkage were added, and a new soot formation model was added.

In response to applicable year-3 TST review comments, we submit that the new soot models (i.e. adding more physics) appear to have increased predictivity significantly by (1) adding porosity to the char model tightened bounds and (2) consistency across the board that was not present in prior models. Although the soot model can have a significant impact on smaller systems, it is possible that the soot model didn't play a significant impact in predicting QOI's for the BSF or the 8-corner simulations. We hope to better characterize in upcoming years.

Char Oxidaton V/UQ

Over this past year there has been an emphasis on the V/UQ effort for char oxidation. The model form and model parameters were being evaluated against the Sandia solid-fuel char-oxidation database for a coal of interest to our center. The model form used in past years was updated to decrease model form uncertainty as determined by the consistency analysis with the Sandia data. This model form exploration and development was performed in conjunction with several members of the Physics team based on the appropriate theory and assumptions, but still recieved mixed results from the consistency analysis. Upon detailed investigation, there were two causes. The first issue was a fitting error in the quadratic surrogate model, which has been mitigated but progress on a complete solution is ongoing. The second issue was conflicting values for the experimental error. The conflict results from the reported experimental error from the experimentalists themselves versus the spread in the data as observed in the

significant number of replicates performed. The experimentalist, depending on the source, were reporting temperature errors in the range of ±5-50 K. However, the experimental design sampled ~100 replicates for each measurement, and the variability in these samples give a range of ±100-300 K (depending on the conditions measured). The distinction between an error in the neighborhood of ±20 K versus one in the range of ±200 K is considerable. The model from is not consistent with the dataset within the former error, but is consistent within the latter. On the other hand, the latter errors are large enough that, if correct, we must re-evaluate whether the dataset is informative enough for our application. In light of this concern, in the last quarter the year, there was a shift to evaluation of the instrument error — for the Sandia data set as well as local experiments in the L-1500.

Bound-to-Bound Data Collaboration: B2BDC package (version 0.86)

During the past year, many new features were added to the B2BDC package. Recent additions include:

- vector consistency measure with various weighting schemes.
- implemented linear constraints to better characterize prior information. Previously, only box constraints were allowed.
- multiple parameter optimization methods were included:
    1. "1N-F" - minimization of 1-norm distance between a point in the feasible set and its nominal value.
    2. "LS-F" - minimization of least-square difference between simulation and experimental measurement over the feasible set.
    3. "LS-H" - minimization of least-square difference between simulation and experimental measurement over prior parameter uncertainty.
- a free optimization package, OPTI toolbox, was included to solve nonlinear optimization problems and served as an additional option to MATLAB's default optimization toolbox. Calculation times were also improved by restructuring the code.
- polynomial models were included in B2BDC as another option for surrogate models (in addition to quadratic and rational quadratic models). This added flexibility allows us to fit more diverse collections of data and enlarges the application space for B2BDC. Functions to generate sparse polynomial surrogate model following the High-Dimensional Model Representation methodology were also included in the B2BDC package. The corresponding B2BDC features such as consistency measure and model prediction were expanded to be compatible with polynomial surrogate models. The underlying optimization problem is relaxed via Sum-of-Square technique and solved with the MATLAB toolbox SparsePOP.
- functions for uniform sampling of the feasible set were included in the B2BDC package.

Vector Consistency

The vector consistency measure (VCM), a recent addition to the Bound-to-Bound Data Collaboration (B2BDC) framework, aimed at resolving disagreement between models and experimental observations, has been applied to an updated DLR-SynG dataset. This dataset consisted of 159 quantities of interest (QOIs) in 55 uncertain parameters and the tool identified approximately 40 QOIs as contributing to inconsistency (Slavinskaya et. al., 2016). A paper introducing and detailing the application of the VCM to two combustion datasets, GRI-Mech 3.0 and DLR-SynG, was submitted to SIAM/ASA Journal on Uncertainty Quantification in January of this year (Hegde et. al., 2017). The VCM approach is shown to be advantageous over B2BDC's sensitivity-based usage of the scalar consistency measure when addressing inconsistency. This is particularly true when dealing with massively inconsistent datasets, i.e. datasets with numerous contributors to the inconsistency. In addition, we have established a new B2BDC workflow for model validation utilizing vector consistency. This workflow currently forms the foundation for how we are validating char oxidation models and model form uncertainty.

Sampling

Generating uniformly distributed samples from a feasible set has been explored to provide additional statistical information, e.g. 1-dimensional marginal histograms and 2-dimensional correlations among parameters, to the B2BDC framework. The proposed strategies are fundamentally based on rejection sampling within a polytope. In fact, provably uniform samples are achieved when the polytope contains the feasible set.

To improve the practicality in higher dimensional cases, an approximation strategy and dimension reduction via principal component analysis (PCA) were investigated. This approximation strategy tries to construct a polytope that contains the majority of the feasible set's volume. Quantitative analysis was conducted analyzing the effect of approximation on sampling efficiency and sample quality. MATLAB codes were developed and tested against both toy examples, with relatively low dimensionality, and more realistic, higher dimensional examples (e.g. GRI-Mech). The performance from the numerical results showed several promising features in higher-dimensional applications. A paper about this topic is currently in preparation.

CCMSC Coal Database

The CCMSC Coal Database Application, which acts as a stand-alone, MATLAB-based front-end (https://github.com/oreluk/coalDB), was released June of 2016. This application enables ease of access to the crowd-sourced coal data which is also available through the primekinetics.org website. This front-end application is used to filter, analyze, and select experimental data to be included as validation data for the char oxidation model.

Instrument Modeling

The UC Berkeley group visited the team at the University of Utah to kick-start collaboration on instrument models, physical models, and model form uncertainty. The B2BDC tools are being employed to validate multiple combinations of instrument and char oxidation physics models (reflecting the model form uncertainty). Specifically, we are using B2BDC's consistency analysis

techniques to validate datasets constructed from specified physics and instrument models, constrained by experimental measurements from the CCMSC Coal Database.

The end-goal is to develop a physical model and an instrument model which is in complete agreement with validation data. Currently, the validation data is from Utah Skyline coal experiments conducted at Sandia's optical entrained flow reactor facility (data accessed through CCMSC Coal Database) [1]. This data includes 72 initial conditions at various mixture compositions, particle sizes and an initial mass fraction of fixed carbon. An initial application of B2BDC to the first instrument-physical model combination revealed massive inconsistency, suggesting that the models under-predicted the experimental data. New/updated instrument models and physical models were developed by the University of Utah to improve the model-data agreement.

Simulations of the char oxidation model calculate the time-evolution of a single coal particle, where the particle temperature served as the quantity of interest (QOI). Each simulation was initialized with a coal particle diameter extracted from experimental data (in this case, we took a mean value). Incomplete validation data led to initial issues with the gas temperature profile, which drives the char oxidation simulation code. These issues were resolved by explicitly defining the inlet condition of the gas temperature. The resulting dataset consisted of 399 QOIs. A scalar consistency analysis of the dataset revealed disagreements between the simulation and experiments. Further analysis showed that this inconsistency was not unique to a particular measurement height, particle size, or gas condition.

A deeper investigation revealed that the validation data had significant variation in particle diameter; not all measurements were localized to the mean particle diameter. An updated approach was taken by simulating a distribution of particle sizes to adequately capture the particle size variation observed in the validation data. Considering an initial particle size distribution required changing the instrument model to properly emulate the measurement process seen in the experiment. To be clear, particles in the experiment which were too small or too dim to be seen by the optical measurement device should be neglected in our analysis. These particles unseen by experiment must also be unseen by simulation. Therefore, an instrument model for the particle light intensity was developed and implemented.

Further iterations of the char oxidation model were developed and examined in collaboration with the University of Utah. The addition of a porosity model to the char oxidation physics model has helped reduced the model-data disagreement significantly as quantified by the consistency measure of B2BDC. A vector consistency analysis had shown that the model with porosity could be brought to consistency by relaxing 29 of the 399 QOIs.

Simulations were moved from local desktop computers to the Ash cluster at University of Utah's Center for High Performance Computing in order to quickly investigate various char oxidation model forms. Sensitivity results informed us that the consistency measure was most significantly affected by the prior bounds on the kinetic parameters. Using this information, the next model form will use wider prior bounds on the parameters.

L-1500 Experimental Campaign and V/UQ

We modified the L1500 burner geometry to include star-shaped bluff body which aids in more uniform distribution of particles in the primary fuel feed. Using the modified geometry, we performed simulations of the L1500 burner using conditions from the 2015 experimental

campaign to provide updated hand-off boundary conditions for the full furnace simulation using Arches.

Our V/UQ effort has continued to focus on the analysis of and comparison with data that were collected in the 1.5 MW pulverized coal test facility (L1500) in year-two. As part of our overall V/UQ analysis, we performed both a sensitivity analysis and a consistency analysis during this year. To perform a V/UQ analysis, the QOIs and the system parameters (scenario, model, numerical) that have a first order impact on the QOIs are identified. In the L1500 experimental dataset, the QOIs were heat flux measurements from three narrow angle radiometers, five wall temperature measurements, and heat removal by eight sets of cooling tubes. From the I/U map, we identified five variables as active: two parameters related to char oxidation (activation energy and pre-exponential factor for the oxidation reaction), and three related to ash deposition (soot blowing time, deposit thermal conductivity and emissivity). We used a sensitivity analysis to reduce the number of dimensions from five to two for the consistency analysis. The most sensitive parameters across the three different types of QOIs were ash deposit thermal conductivity and emissivity.

For the consistency analysis, we combined these two parameters into a single "effective" thermal conductivity parameter and added two scenario parameters, a burner swirl parameter that was applied to the tangential components of the inlet velocity and the coal feed rate. Without the addition of these parameters, we were unable to find consistency with the experimental data. We ran 34 large eddy simulation (LES) cases exploring the effect of three parameters on the QOIs. The VUQ methodology we employed required uncertainty bounds on the experimental data that included both the sampling and systematic errors. We used an instrument model to estimate the systematic error in a device measuring the top wall temperatures of the L1500. With this procedure, we estimated a ±115 K error, which is much greater than the estimated random error of ±2.3 K. We then performed the consistency analysis and were able to reduce the experimental data error to ±26 K for the wall thermocouples. The parameter with the largest impact on the QOIs was the "effective" thermal conductivity, so the ash deposition model needs to be refined. The swirl parameter also had a large effect, indicating a need to better understand the burner design and operation.

Boiler Simulator Facility

The BSF plays a unique, central role as it is believed to be the closest system, in terms of physical regimes, to the 8-corner unit. Thus, excellent validation of the system yields higher confidence in the predictions being made. A V/UQ study was performed on the BSF during year-3 that was aimed to investigate the consistent region between the experimental and the simulations results. The variables selected for the V/UQ study, showed the most sensitive behavior to changes in the simulation system. The variables were slag temperature, thermal conductivity in the refractory, and the activation energy of the reactions considered. Twenty-one simulation cases were run for the study and a narrower consistent region was found as compared to the V/UQ studies prepared in previous years. Table 2 summarizes the reduction in our uncertain space and the uncertainty in the predicted quantities.

Uncertain Parameter	prior range	posterior range	consistent qoi range
T slag	1350 - 1600 K	1490 - 1533 K	gas temp: ±2%
k (RAM 90)	2.5 - 4.5 W/m/K	3.3 - 3.7 W/m/K	heat flux: ±6%
A char-factor	1.0 - 2.35	1.69 - 1.87 (3-4s)	O_2: ±5%

Table 2: Summary of BSF priors and posteriors

GE Power 8-corner Unit Simulations

An increased emphasis was placed in the 8-corner simulations as requested from the year-2 review. Several issues running ARCHES at large core counts (up to 256,000 cores) were encountered while running the new case; namely, poor data I/O performance, slow compute time using our DO Radiation model, running out of memory, and slow standard time-step compute times. These issues were addressed by the physics and computer science teams resulting in a 60% decrease in overall computation time. We were able to meet with GE/Alstom to present the results and receive feedback. During the meeting, we were able to confirm that the simulations met expectations as far as temperatures and heat-flux at the various planes in the boiler, even in the presence of the uncertainty that exists in both scenario and model parameters within the computation.

Using CAD provided by GE Power for the 8-corner unit, we corrected the CAD into a format read-able by commercial packages. Once imported, we have completed STAR-CCM+ simulations of the secondary over-fired air (SOFA) nozzles for the GE Power 1,000 MW boiler, including detailed representation of all nozzles and nearby upstream ducting. We have further implemented a two-way coupling procedure between Arches and STAR-CCM+: we use results from Arches full-boiler simulations as the inlet boundary conditions for the STAR-CCM+ simulations of the entire SOFA section. This, along with the detailed geometric representation of the SOFA nozzles and nearby ducting, allows STAR-CCM+ to account for the hot gasses passing from the boiler radiant section into the SOFA section, modifying the flow patterns at the exit of the SOFA nozzles. This modified procedure provides an improved characterization of the non-uniform velocity and temperature fields near the nozzles, which, in turn, can be mapped back as boundary conditions for the full-boiler simulations in Arches. To preserve the detailed geometric representation of the nozzles as well as resolve the mixing in the boiler SOFA section, the latest unsteady LES STAR-CCM+ simulations employ on the order 800 million computational cells. The STAR-CCM+ simulations fully leveraged the university computational resources.

Overarching Prediction Design / Full System Integration

As part of the full-system integration hierarchy, the team refocused the 500 MW_e (1000 MW_t) ultra-supercritical oxy-fired boiler to a 1000 MW_t oxy-fired high pressure gasifier to broaden the applications for the resulting output species. In addition to widening the applicability of the system, this new gasifier greatly reduces capital costs and system footprint. The system is designed as a dry feed injection with pure oxygen as the oxidizer. The coal remains the same (Wyoming subbituminous – Black Thunder) and the system is designed with countercurrent oxidizer-fuel injectors with a majority of the ash material exiting the bottom of the gasifier. Additionally, pilot-scale pressurized gasification tests were completed at BYU in the late 1980s and the operating conditions and data are being collected and analyzed for model verification. We also are working on instrument models and data validation from the L1500 combustion test series completed in 2016.

A sensitivity analysis was conducted for the prediction case varying several numerical and physical parameters. This study showed that the emissivity played the largest role in accurately determining the heat flux to the wall. This study required 15 million CPU-hours.

A scaling study requiring 10 million CPU-hours was conducted to show how resolution affects the quantities of interest. Our findings showed that resolution did not significantly impact the heat flux to the walls. Computational performance was also monitored in this scaling study between 32k cores and 128k cores. The increase in computational cost under weak scaling constraints for various model components is as follows:

*Overall scalability between 32K and 128K cores - **35% slower***

- Pressure Solve - 110% slower - contribution to overall **10%**
- Arches (w/o hypre) - 4% slower - contribution to overall **3%**
- Radiation Solve - 15% faster - contribution to overall **-3%**
- Disk I/O - 210% slower - contribution to overall **25%**

The above findings showed that the code was 35% slower in a weak scaling study ranging from 32k cores to 128k cores when using the old Uintah data archive format. The primary reason for this slowdown was because of disk i/o times. Our I/O strategy is being addressed by the PIDX team.

The Incite award was used to conduct a ~50 million CPU-hour 256,000 core prediction simulation of the oxy-coal 500 MW boiler was completed using the allocation granted by INCITE on ANL's Mira. This simulation pushed the limits of the Uintah-arches algorithm using all physics models at large scale. To achieve portability, the Arches look-up table was adapted to use Kokkos views, a crucial step in achieving hardware portability for Arches.

Using high-performance computing platforms, we explored a range of scenarios and designs to predict and quantify heat flux distribution inside the design boiler, aiding in material science research of steam tubes that are able to withstand the high temperatures and pressures of oxy-fired AUSC systems. We employed the use of lower-fidelity (RANS) models to explore a wide range of design scenarios, such as overall boiler dimensions, relative lengths of radiant and convective sections, burner placement, hopper and nose dimensions, as well as geometric features for reduction of temperatures on the walls. Using steady RANS models to perform

hundreds of simulations to map out a large design space we were able to quantify the effects of design parameters on boiler design and performance. By performing a multi-objective trade-off study and cross-correlation analysis, we are able to narrow down the number of design parameters with the highest impact on the local temperature and local heat flux distributions inside the design 500 MW oxy-coal AUSC boiler.

The prediction case underwent various redesigns this year, before adopting a totally new approach to maximize efficiency and minimize costs. The oxy-fired boiler has been redesigned as a high pressure oxy-fired dry coal-fed gasifier. This makes the design more attractive to industrial collaborators. The highlighted features of this new design are that it is a less expensive design, a versatile technology applicable to a wide variety of applications, high efficiency system, and a carbon capture solution.

One of the problems with adopting this new design was accounting for the pressure increase in all the physical models used in Arches. Arches has typically ran under atmospheric conditions, a wholesale review of the models, and sub-models and the extent of their pressure dependence was investigated. In most cases adding pressure dependence was trivial. Some work remains to have pressure dependence fully integrated into arches.

The high-pressure design has been simulated using new pressure dependent code. The input file for this new design has been constructed and run on 13,000 cores on Quartz, totaling in 260 million cells. This simulation took 1.5 million CPU-hours to generate 5 seconds of simulated time.

Students and Post-doctoral Fellows

- Michael Brown, MS candidate, University of Utah
- Siddartha Ravichandran, MS candidate, University of Utah
- Teri Draper, Ph.D. candidate, University of Utah
- Joshua McConnell, Ph.D. candidate, University of Utah
- John Holmen, Ph.D. candidate, University of Utah
- Damodar Sahasrabudhe, Ph.D. candidate, University of Utah
- Pavol Klacansky, Ph.D. candidate, University of Utah
- Duong Thai Hoang, Ph.D. candidate, University of Utah
- William Usher, Ph.D. candidate, University of Utah
- MinMin Zhou, Ph.D. candidate, University of Utah
- Qi Wu, Ph.D. candidate, University of Utah
- Kaitlyn Scheib, Ph.D. candidate, University of Utah
- Alex Josephson, Ph.D. candidate, Brigham Young University
- Andrew Richards, Ph.D. candidate, Brigham Young University
- Arun Hegde, Ph.D. candidate, University of California-Berkeley
- Wenyu Li, Ph.D. candidate, University of California-Berkeley
- Jim Oreluk, Ph.D. candidate, University of California-Berkeley

- Sidharth Kumar, research associate, University of Utah
- Babak Goshayeshi, research associate, University of Utah
- Oscar Diaz-Ibarra, post-doctoral research fellow, University of Utah
- John Camilo Parra Alvarez, post-doctoral research fellow, University of Utah
- Steve Petruzza, post-doctoral research fellow, University of Utah

Former CCMSC students and employees

- Christopher Earl, staff member, LLNL
- Ben Schroeder, staff member, SNL-Berkeley
- Alexander Abboud, staff member, INL
- Aaditya Landge, software engineer, Twitter
- Pascal Grosset, staff member, LANL
- Mark Kim, staff member, ORNL
- Troy Holland, post-doctoral associates, LANL

Collaborators (unpaid)

- John Marion - Director Technology and R&D, GE Power
- Bob Kunkel - Chief Engineer, GE Power
- Armand Levaser - Principle Investigator BSF, GE Power
- Tracy Midgley - Principal Consulting Engineer, GE Power
- Paul Chapman - CFD Technology Manager, GE Power
- David Sloan - CFD specialist, GE Power
- Yen-Ming Chen - CFD specialist, GE Power
- Zhenyuan Liu (UC Berkeley),
- Jerome Sacks (NISS)
- Rui Paulo (University of Lisbon, Portugal)
- Gonzalo Garcia-Donato (University of Castilla-La Mancha, Spain)
- Craig Needham (NC State University)
- Sathya Baskaran and Mani Sarathy (KAUST, Saudi Arabia)
- Michael Burke (Columbia University)
- Richard West (Northeastern University)
- Phillip Westmoreland (NC State University)
- Nadja Slavinskaya, Jan Starcke and Uwe Riedel (DLR, Stuttgart, Germany)
- Stanislav Shvartsman and Henry Mattingly (Princeton University)

Publications and Presentations

Publications

- Bhatele, A., N. Jain, Y. Livnat, V. Pascucci, and P. Bremer. *"Analyzing network health and congestion in dragonfly-based supercomputers,"* 2016 IEEE International Parallel and Distributed Processing Symposium, IPDPS 2016, Chicago, IL, USA, May 23-27, 2016, pp. 93-102, 2016.
- Bhatia, H., A. Gyulassy, V. Pascucci, M. Bremer, M. T. Ong, V. Lordi, E. W. Draeger, J. E. Pask, and P. Bremer. *"Interactive exploration of atomic trajectories through relative-angle distribution and associated uncertainties,"* 2016 IEEE Pacific Visualization Symposium, (PacificVis), Taipei, Taiwan, April 19-22, 2016, pp. 120 – 127, 2016.
- Diaz-Ibarra, O., J. Spinti, A. Fry, B. Schroeder, J. Thornock, B. Isaac, D. Harris, M. Hradisky, S. Smith, E.G. Eddings and P.J. Smith. "*A validation/uncertainty quantification analysis of a 1.5 MW oxy-coal fired furnace*." Western States Section of the Combustion Institute 2016 Spring Meeting, Seattle, WA, March 21-22, 2016.
- Diaz-Ibarra, O., J. Spinti, A. Fry, B. Isaac, J.N. Thornock, M. Hradisky, S. Smith, P. J. Smith. "*A validation/uncertainty quantification analysis for a 1.5 MW oxy-coal fired L1500 furnace using a swirling boundary condition*." American Flame Research Committee 2016 Industrial Combustion Symposium, Kauai, HI, September 11-14, 2016.
- Earl, C., M. Might, A. Bagusetty, and J. Sutherland. *"Nebo: An efficient, parallel, and portable domain-specific language for numerically solving partial differential equations."* J. Syst. Softw., vol. 125, pp. 389–400, March 2017.
- Frenklach, M., A. Packard, J. Sacks, R. Paulo and G. Garcia-Donato, *"Comparison of statistic and deterministic frameworks of uncertainty quantification."* SIAM/ASA J. Uncertainty Quantification **4**: 875-901 (2016). ISSN: 21662525 DOI: 10.1137/15M1019131
- Frenklach, M., A. Packard, J. Sacks, R. Paulo, and G. Garcia-Donato. *"Comparison of statistic and deterministic frameworks of uncertainty quantification."* *SIAM/ASA J. Uncertainty Quantification* **4**: 875-901 (2016).
- Fry, A., J. Spinti, E. Eddings. "*Pilot-scale investigation of heat flux and radiation from an oxy-coal flame*." 41st International Technical Conference on Clean Coal & Fuel Systems, Clearwater, Florida, June 5-9, 2016.
- Fry, F., J. Spinti, O. Diaz-Ibarra, I. Preciado, K. Scheib, S. Harding, E. Eddings, P. J. Smith. "*Pilot-scale investigation of heat flux and radiation from an oxy-coal flame, Part 1: Development of instrument models*." American Flame Research Committee 2016 Industrial Combustion Symposium, Kauai, HI, September 11-14, 2016.
- Hegde, A., W. Li, J. Oreluk, A. Packard, M. Frenklach. *"Consistency Analysis for Massively Inconsistent Datasets in Bound-to-Bound Data Collaboration."* Submitted to the SIAM/ASA Journal on Uncertainty Quantification. Preprint available at *arXiv:1701.04695*, 2017.
- Holland, T. and T. H. Fletcher, *"Global Sensitivity Analysis for a Comprehensive Char Conversion Model in Oxy-fuel Conditions."* Energy and Fuels, **30**, 9339−9350 (2016). DOI: 10.1021/acs.energyfuels.6b02190.

- Holland, T., T. H. Fletcher. *"A Comprehensive Model of Single Particle Pulverized Coal Combustion Extended to Oxy-coal Conditions."* Energy Fuels, 2017, 31 (3), pp. 2722-2739. DOI: 10.1021/acs.energyfuels.6b03387.
- Josephson, A.J., D.O. Lignell, A.L. Brown, and T.H. Fletcher, *"Revision to Modeling Soot Derived from Pulverized Coal."* Energy & Fuels, **30**, 5198-5199 (2016). ISSN: 08870624 DOI: 10.1021/acs.energyfuels.6b01007.
- Josephson, A. J., D. O. Lignell, A. L. Brown, and T. H. Fletcher, "Revision to Modeling Soot Derived from Pulverized Coal," *Energy & Fuels*, **30**, 5198-5199 (2016). DOI: 10.1021/ acs.energyfuels.6b01007.
- Maljovec, D., B. Wang, P. Rosen, A. Alfonsi, G. Pastore, C. Rabiti, and V. Pascucci. *"Rethinking sensitivity analysis of nuclear simulations with topology,"* 2016 IEEE Pacific Visualization Symposium (Paci_cVis), pp. 64-71, April 2016.
- McConnell, J., and J. Sutherland. *"The Effect of Model Fidelity on Prediction of Char Burnout for Single-Particle Coal Combustion."* Proc. Combust. Inst., vol.36, no. 2, pp. 2165-2172, 2017.
- McConnell, J., B. Goshayeshi and J. Sutherland. *"An Evaluation of the Predictive Capability of Various Coal Combustion Models."* Fuel, vol. 201, pp. 53-64, 1 August 2017.
- McConnell, J., B. Goshayeshi, J.C. Sutherland. *"An Evaluation of the Efficacy of Various Coal Combustion Models for Predicting Char Burnout."* Fuel, vol. 201, pp. 53-64, 2017.
- Richards, A. P. and T. H. Fletcher, "*A Comparison of Simple Global Kinetic Models for Coal Devolatilization with the CPD Model," Fuel*, **185**, pp. 171-180 (2016).
- Saad, T., and J. Sutherland. *"Wasatch: An Architecture-Proof Multiphysics Development Environment using a Domain Specific Language and Graph Theory."* Journal of Computational Science (2016). ISSN: 18777503 DOI:10.1016/j.jocs. 2016.04.010.
- Slavinskaya, N.A., M. Abbasi, J. H. Starcke, R. Whitside, A. Mirzayeva, U. Riedel, W. Li, J. Oreluk, A. Hegde, A. Packard, M. Frenklach, G. Gerasimov, and O. Shatalov. *"Development of an UQ-Predictive Chemical Reaction Model for Syngas Combustion."* Energy and Fuels, 2017, 31 (3), pp. 2274-2297. DOI: 10.1021/acs.energyfuels.6b02319.
- Peterson, B., A. Humphrey, D. Sunderland, T. Harman and M. Berzins. "*Optimizing the Uintah Framework for Large Scale Radiative Heat Transfer Calculation on GPUs."* Submitted to Journal of Computational Science, August 2016)
- Jatale, A., P.J. Smith, S.T. Smith, J.N. Thornock, J.P. Spinti and M. Hradisky, *"Multiscale Validation and Uncertainty Quantification: Toward Predictivity."* Journal of Verification, Validation and Uncertainty Quantification. (2016)
- Hegde, A., W. Li, J. Oreluk, A. Packard, and M. Frenklach. "*Consistency analysis for massively inconsistent datasets."* Submitted to the SIAM/ASA Journal on Uncertainty Quantification. Preprint available at https://arxiv.org/abs/1701.04695

Presentations

- Berzins, M. *"Solving Petascale Turbulent Combustion Problems with the Uintah Software."* 5th European Seminar on Computing (ESCO 2016), Pilzen, Czech Republic (June 9, 2016).

- Berzins, M. *"Developing Software for Multiscale Multiphysics Modelling from Petascale to Exascale."* Keynote talk at DOD PETTT Program Annual Meeting, Vicksburg, MS (October 18, 2016).
- Berzins, M. Invited talk at School of Computing, University of Leeds, UK (November 4, 2016).
- Diaz-Ibarra, O.H., M. Hradisky, J.P. Spinti and P.J. Smith. *"Simulations of 1.5 MW Oxy-coal Fired Furnace Using STAR-CCM+ and ARCHES Coupling Approach.*" STAR Global Conference, Prague, Czech Republic (March 7-9, 2016)
- Diaz-Ibarra, O.H., J.P. Spinti, A. Fry, B. Schroeder, J.N Thornock, B. Isaac, D. Harris, M. Hradisky, S.T. Smith, E.G. Eddings and P.J. Smith. *"A Vaidation/Uncertainty Quantification Analysis of a 1.5 MW Oxy-coal Fired Furnace.*" 2016 Spring Technical Meetings, Western States Section of the Combustion Institute, University of Washington, Seattle, WA (March 21-22, 2016).
- Diaz-Ibarra, O.H., J. Spinti, A. Fry, B. Isaac, J.N. Thornock, M. Hradisky, S.T. Smith and P.J. Smith. *"A validation/uncertainty quantification analysis for a 1.5 MW oxy-coal fired L1500 furnace using a swirling boundary condition."* American Flame Research Committee 2016 Industrial Combustion Symposium, Kauai, HI (September 11-14, 2016).
- Draper, T., A. Fry, L. Kolczynski, T. Ring and E.G. Eddings. *"Thermal Characterization of a 1.5 MW Pulverized-coal Reactor using Infrared Heat Flux and Measured Heat Loss."* 41st Interntional Technical Confernce on Clean Coal and Fuel Systems, Clearwater, FL (June 5-9, 2016.
- Fry, A., J.P, Spinti and E.G. Eddings. *"Pilot-scale Investigation of Heat Flux and Radiation from an Oxy-coal Flame."* 41st Interntional Technical Confernce on Clean Coal and Fuel Systems, Clearwater, FL (June 5-9, 2016.
- Fry, A., J.P. Spinti, O.H. Diaz-Ibarra, I. Preciado, K. Scheib, S. Harding, E.G. Eddings and P.J. Smith. *"Pilot-Scale Investigation of Heat Flux and Radiation from an Oxy-coal Flame, Part 1: Development of Instrument Models."* American Flame Research Committee Industrial Combustion Symposium, Kauai, HI (September 11-14, 2016).
- Fry, A., J.P. Spinti, I. Preciado, O.H. Diaz-Ibarra, M. Hradisky and J.P. Spinti. *"Towards Next Generation Simulations of Full-scale Coal-fired Boilers."* 41st International Technical Conference on Clean Coal and Fuel Systems, Clearwater, FL (June 5-9, 2016).
- Goshayeshi, B. and J. C. Sutherland, *"Hybrid Computing in Large-Scale Multiphysics Simulation: Tabulated Properties and Particle-Cell Interpolations."* SIAM Computational Science and Engineering Conference, Atlanta, GA (2017).
- Grossett, A.V.P, A. Knoll and C. Hansen. *"Dynamically Scheduled Region-Based Image compositing."* Eurographics Symposium on Parallel Graphics and Visualization 2016, Groningen, Netherlands (June 6-7, 2016).
- Hegde, A., J. Oreluk, W. Li, A. Packard, M. Frenklach. *"Sparse Resolutions to Inconsistent Datasets Via L1-Minimization."* Presentation at *SIAM Numerical Combustion 17*, Orlando, FL. (April 4, 2017).
- Holland, T., S. Bjat, P. Marcy, J. Gattiker, J. D. Kress, and T. H. Fletcher. *"Bayesian Uncertainty Quantification and Calibration of a Clean Coal Design Code,"* presented at CO2 Summit II: Technologies and Opportunities, Santa Ana Pueblo, NM (April 10-14, 2016).

- Holland, T., S. Bhat, J. D. Kress, and T. H. Fletcher. *"Extension and Calibration of Coal Combustion Models,"* poster presented at the Los Alamos National Laboratory Student Symposium, Los Alamos, NM (Aug. 3, 2016).
- Holland, T. and T. H. Fletcher. *"A Comprehensive Model of Single Particle Pulverized Coal Combustion Extended to Oxy-coal Conditions,"* to be presented at the 10th National Combustion Institute Meeting, University of Maryland (April 23-26, 2017).
- Humphrey, A., T. Harman, D. Sunderland, and M. Berzins. *"Adaptive Mesh Revindment in Radiation Problems - Challenges in Parallel Adaptive Mesh Refinement – Part III."* SIAM Parallel Processing for Scientific Computing, Paris, France (April 12-15, 2016).
- Humphrey, A., T. Harman, D. Sunderland, and M. Berzins. *"Radiative Heat Transfer Calculation on 16384 GPUs Using a Reverse Monte Carlo Ray Tracing Approach with Adaptive Mesh Refinement,"* The 17th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC), Chicago, IL (May 27, 2016).
- Iavarone, S., B. Isaac, S. Smith, P. J. Smith, J. Oreluk, A. Packard, M. Frenklach, A. Parente. *"Collaborative Simulations and Experiments for Development and Uncertainty Quantification of a Reduced Char Oxidation and Gasification Model in Oxy-Coal Combustion Conditions."* Presentation at *SIAM Numerical Combustion 17*, Orlando, FL. (April 5, 2017).
- Josephson, A. J., N.D. Gaffin and D.O. Lignell. *"Soot Nucleation and Consumption in Oxy-Coal Systems,"* 2016 AIChE Annual Meeting, San Francisco, CA (Nov. 13-18, 2016).
- Kolczynski, L., A. Fry, T. Draper, T. Ring and E.G. Eddings. *"Density, Porosity and Heat Capacity Characteristics of Ash Deposits from a 1.5 MW Coal Furnace."* 41st International Technical Conference on Clean Coal and Fuel Systems, Clearwater, FL (June 5-9, 2016).
- Li, W., A. Hegde, J. Oreluk, A. Packard, M. Frenklach. *"Uniform Sampling of Feasible Set: A Hybrid Statistical-Deterministic Method of Uncertainty Quantification."* Presentation at *SIAM Numerical Combustion 17*, Orlando, FL. (April 4, 2017).
- McConnell, J. and J.C. Sutherland. *"The Effect of Model Fidelity on Prediction of Char Burnout for Single-particle Coal Combustion."* The 36th International Combustion Symposium, Seol, Korea (2016).
- J. Oreluk, J., A. Hegde, W. Li, A. Packard, M. Frenklach. Providing structure to experimental data: A large scale heterogeneous database for collaborative model validation. Presentation at *SIAM Numerical Combustion 17*, Orlando, FL. (April 4, 2017).
- Pascucci, V. *"Extreme Data Management Analysis and Visualization for Exascale Supercomputers."* High Performance Computing from Clouds and Big Data to Exascale and Beyond, Cetraro, Italy (April 28, 2016).
- Pascucci, V. *"Extreme Data Management Analysis and Visualization for Science Discovery."* Brookhaven National Laboratory, Upton, NY (June 19, 2016).
- Pascucci, V. *"Multi-Scale Morse Theory for Science Discovery."* Invited talk at Stonybrook University, NY (December 2016).
- Pascucci, V. *"Extreme Data Management, Analysis for Oil and Gas Applications."* Invited talk at Instituto Tecnologico de Buenos Aires, Argentina (December 2016).
- Pascucci, V. *"Extreme Data Management, Analysis for Oil and Gas Applications."* Invited talk at National University of Comahue, Neuquén, Argentina (December 2016).

- Pascucci, V. *"Extreme Data Management, Analysis and Visualization: Exploiting Large Data for Government Agencies."* Invited talk at TecNap 2016, Villa La Angostura, Argentina (November 2016).
- Pascucci, V. *"Multi-Scale Morse Theory for Science Discovery."* Invited talk at University of Rome "Roma 3", Italy, November 2016.
- Pascucci, V. *"Extreme Data Management Analysis and Visualization for Science Discovery."* Consiglio Nazionale dei Geologi (National Council of Geologists), Rome, Italy (July 11, 2106)
- Saad, T. and J.C. Sutherland. *"An Explicit Variable-Density Projection Method for Low-Mach Reacting Flows on Structured Uniform Grids."* AIChE Annual Meeting. San Francisco, CA, November 13-18, 2016.
- Saad, T. and J. C. Sutherland. *"An Explicit Variable-Density Projection Method for Low-Mach Reacting Flows on Structured Uniform Grids."* AIChE Annual Meeting, 2016.
- Sanderson, A. *"VisIt and Uintah an In-situ Marriage."* Annual DOE Computer Graphics Forum, Monterey, CA (April 2016).
- Slavinskaya, N.A., J. H. Starcke, M. Abbasi, A. Mirzayeva, U. Riedel, M. Frenklach, A. Packard, W. Li, J. Oreluk, A. Hegde. *"Consistent syngas chemical mechanism from collaborative data processing."* 55th AIAA Aerospace Sciences Meeting. doi:10.2514/6.2017-0837, 2017.
- Slavinskaya, N.A., J. H. Starcke, M. Abbasi, A. Tursynbay, U. Riedel, W. Li, J. Oreluk, A. Hegde, A. Packard, M. Frenklach. *"Consistent Chemical Mechanism from Collaborative Data Processing."* Proceedings of the World Congress on Momentum, Heat and Mass Transfer. doi:10.11159/csp16.118, 2016.
- Smith, S.T. and S. Iavarone. *"Uncertainty Quantification for Coarse-Grained Modeling of Coal Devolatilation."* Clearwater Clean Energy Conference, Clearwater, FL (June 2016).
- Sunderland, D. B. Peterson, J. Schmidt, A. Humphrey, J. Thornock and M. Berzins. *"Performance Portability in the Uintah Runtime System Through the Use of Kokkos."* 2nd International Workshop on Extreme Scale Programming Models and Middleware, International Conference for High Performance Computing, Networking, Storage and Analysis (SC '16).
- Sunderland, D. *"Performance Portability in the Uintah Runtime System Through the Use of Kokkos."* Super Computing 2016 Workshop, (November 18, 2016).
- Sutherland, J.C. and T. Saad, *"Case Studies in Using a DSL and Task Graphs for Portable Reacting Flow Simulations."* SIAM Computational Science and Engineering Conference, Atlanta, GA (2017).
- Sutherland, J. C. *"Tools and Techniques for Multiscale Simulation of Reacting Flows."* Ma'anshan, China: Anhui University of Technology. Invited talk (2016).
- Sutherland, J.C. J. McConnell, and B. Goshayeshi. *"An Assessment and Comparison of Various Coal Combustion Models."* First International Workshop on OxyFuel Combustion. (2016).
- Thornock, J.N., P.J. Smith, B. Isaac, S.T. Smith, O.H. Diaz-Ibarra, M. Hradisky and J.P. Spinti. *"Towards Next Generation Simulations of Full-scale Coal-fired Boilers."* 41st International Technical Conference on Clean Coal and Fuel Systems, Clearwater, FL (June 5-9, 2016).