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, 2015 through February 29, 2016

May 31, 2016

Annual Technical Report

DE-NA0002375

March 1, 2015 through February 29, 2016

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 designed a course that was taught in fall 2015. The course was designed for first-year graduate students at the three CCMSC universities: Brigham Young University, University of California-Berkeley and University of Utah. Enrollment was 4 from UCB, 3 from BYU and 5 from UU, plus five non-registered observers. Instructors were Philip Smith and Sean Smith (UU) and Andrew Packard and Michael Frenklach (UCB). The course was presented at the University of Utah or University of California-Berkeley and telecast to all three universities. 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 will be addressed in designs for the fall 2016 course. Members of the NNSA community will be invited to participate in the 2016 course.

Several task investigators (Tom Fletcher, Martin Berzins and Eric Eddings) attended and presented progress at the Advanced Scientific Computing Investigators' Meeting that was held in Las Vegas in February 2016. Other key meetings at which presentations were made include the Stewardship Science Academic Programs (SSAP) Symposium in Bethesda, Supercomputing 2015 (SC15) in Austin, American Physical Society (APS) Meeting in San Antonio and Workshop on Exascale Software Technologies (WEST) meeting in Albuquerque.

The UCB team collaborated with Professor Stas Shvartsman at Princeton University on application of the B2BDC methodology to system biology and hosted a visiting graduate student from the Shvartsman Lab, Henry Mattingly. Mr. Mattingly expressed interest in applying the B2BDC framework to his research during summer 2015.

Cameron Christensen, Peer-Timo Bremer, Duong Hoang, Valerio Pascucci, and Steve Petruzza Presented tutorial, "Big Data and I/O for Scientific Simulations." Lawrence Livermore National Laboratory, California, March 1-2, 2016.

Cameron Christensen, Peer-Timo Bremer, Rooh Khurram, Sidharth Kumar, Bilel Hadri, Doang Hoang, Georgios Markomanolis, Valerio Pascucci, Steve Petruzza, and Madhu Srinivasan. Presented tutorial, "Big Scientific Data Made Simple A Hands-on Tutorial in Data Generation, Processing, and Delivery for High Performance Computing and High Resolution Imaging." KAUST, Saudi Arabia, June 28-29, 2015.

Valerio Pascucci, Presented "Visualization" to the Topology and Feature Analysis Summer School, STFC Hartree Centre SciTech Daresbury Warrington WA4 4AD United Kingdom, June 14-19, 2015.

Four graduate students completed internships during year-two; namely, **Teri Draper** (Scott Skeen at Sandia National Laboratory – Livermore), **Oscar Diaz-Ibarra** (Chris Shaddix at Sandia National Laboratory – Livermore), **Mark Kim** (Peter Lindstrom at Lawrence Livermore National Laboratory) and **Troy Holland** (Joel Kress and T-1 group at Los Alamos National Laboratory).

John Holmen has scheduled his internship for May through August 2016 at Sandia National Laboratory with Jonathan Hu and Ray Tuminaro**.** **Allen Sanderson** (professional staff member) scheduled his research experience with Eric Brugger at LLNL for May**.** **Joshua McConnell** has scheduled his student internship for fall 2016 with Stefan Domino at SNL in Albuquerque.

During year-two the several students graduated and moved into post-graduate positions. **Mark Kim** and **Sidharth Kumar** (Ph.D. 2015) are now working as postdoctoral fellows at the University of Utah and **Aaditya Landge** (M.S. 2015) is working at Twitter. Alex Abboud (Ph.D. 2015) is working at Idaho National Laboratory. **Troy Holland** has accepted a position at Los Alamos National Lab to begin in June 2016.

Computer Science: I/O, EDSL, Visualization

Runtime System and Infrastructure

As we started year two, large-scale production runs on more than 128K cores on Vulcan were not able to run because of excessive memory use. Typically, runs consume more than 1.0GB per core, exceeding the machine's capacity. The excessive memory usage was due to the Hypre library that is used for the Pressure solve and the Radiation calculation using the Discrete Ordinates method. This year, we optimized the use of Hypre and OpenMP threads working in conjunction with Uintah's MPIOnlyScheduler to significantly reduce the memory use for both the Pressure solve and the Radiation calculation. The order of magnitude decrease in memory usage permits us to run jobs up to 512K cores on Mira while using just half the memory per node.

Both the CS infrastructure team and the Physics application/VUQ teams worked on the integration of Kokkos with Arches, the production code for this center. Comparisons have been promising and work is ongoing to assess the viability and performance characteristics of this integrated approach for performance portability for current and future architectures.

This year, Titan resources have been used to do testing on an infrastructure improvement development branch, aimed at reducing overhead within the Uintah infrastructure. Specifically, it was focused on the decrease in weak scalability caused by the increase in local MPI communication time (the sum of MPI receives and sends on a multi-core node). We have determined the growth in this local communication to be due to contention for shared resources, specifically mutex-protected vectors of MPI communication records and other related critical sections that unnecessarily serialized sections of code. These mutex-protected communication lists have now been replaced with thread-scalable lock-free lists, significantly reducing this local communication time. We are still testing these changes.

The principal result to note is the near order of magnitude reduction in local communication time. Additionally, this work has fixed a critical memory leak which was caused by a race condition in the previous MPI communication records that managed outstanding MPI requests, a mutex protected vector that caused multiple threads to simultaneously try and process a received message. Many threads would allocate memory for a message but only one thread would de-allocate its memory. This memory leak was most pronounced with RMCRT, as it produces significantly more MPI messages than typically expected. This memory leak was fixed by the addition of the aforementioned, thread-scalable data structures used to reduce the local communication costs.

I/O

During the year we organized a demonstration of the technology developed and presented at the Super Computing conference. This allowed us to complete a number of tasks including a first complete prototype of the data-streaming infrastructure including direct I/O from the simulation, access via ViSUS server and remote visualization ViSUS client. The demonstration has been successful and received a lot of attention during the Super Computing conference and we now have a Visus server installed on the Remus machine in University of Utah's Center for High Performance Computing that holds a number of Uintah datasets accessible remotely.

During this period we focused on scaling studies for the I/O infrastructure. This has led to scaling studies that at the moment have had an emphasis on the BlueGeneQ architecture. In particular, we have accomplished a run on the full Mira machine at the Argonne Leadership Computing Facility. The plot below shows the scaling of the PIDX library compared to the IOR standard benchmark both with single shared file and with single file per process. This is a very encouraging result since indicates that the utilization of the advanced IDX file format does not come with the downside of a slower performance. In fact, in many cases, especially at large scale, the PIDx dumps are faster than IOR.

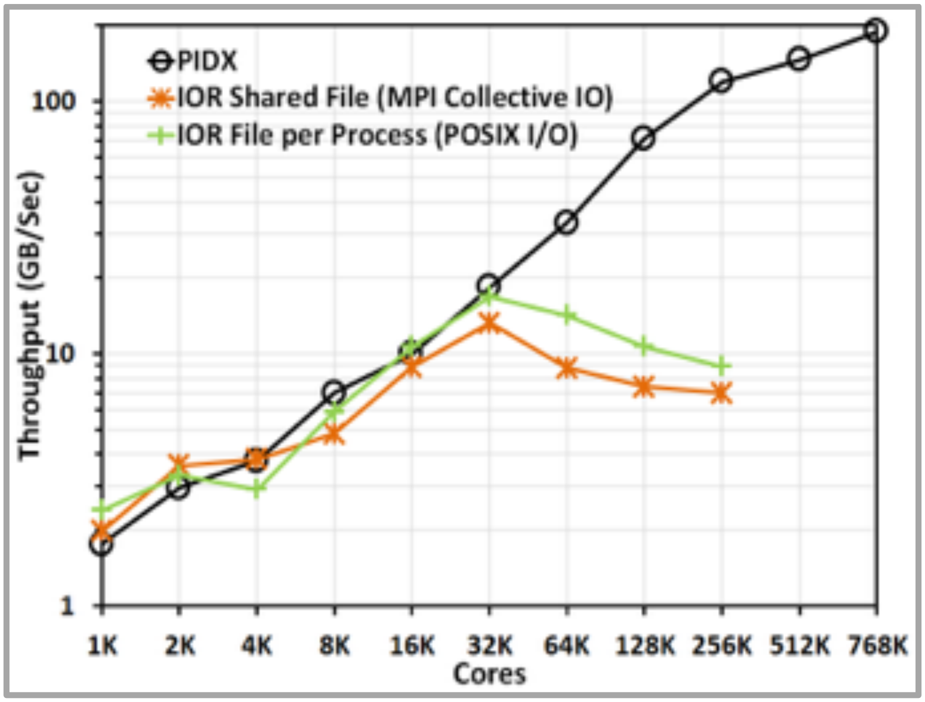

This activity of pushing the limits of the I/O performance has been complemented by the deployment of the PIDX I/O for the Unitah code used for the science runs in this project. TO this end we are working at improving the usability of the library for science users, with specific focus on the CPHC center at the University of Utah and later on at DOE facilities. In particular, the ALCF we have also tested a more sophisticated data access mechanism that should be better suited for usage within a restricted environment.

Finally, we have started to generate documentation, examples, and other material for distribution with the code. To facilitate adoption of the technology. This has resulted in a first tutorial for the users of the KAUST high performance computing facility that we used as a testing ground for a tutorial that we plan to provide to DOE scientists within the scope of the PSAAP project. Some material from the tutorial is being distributed at the site: https://sites.google.com/site/bigdatahpc/

Visualization

We continue to integrate Uintah and VisIt for in-situ analysis and visualization. The in-situ parallel implementation has been completed and we have completed some very preliminary scaling tests (~400 processes) on our local cluster (ash.chpc.edu). In addition to exploring simulation data users are now able to explore process data i.e. per process memory usage, task times, mpi communication times, etc.(top image). Such exploration is critical to debugging large scale problems. The main simulation interface has been expanded to allow users to add tasks such as re-griding or saving unscheduled checkpints (right side of the second image). Users can stop, resume, and abort Uintah simulations. Uintah-VisIt Simulation Dashboard provides a snapshot and an interface to global state variables - some of which can be modified by the application scientist allowing for computational steering and visual debugging. Additional state variables can be added by the component developer – i.e. the computer scientist's help is not needed. The In-Situ Dashboard also provides real time access to run time (on-the-fly) analysis such as Min/Max values. Normally, this information is written to disk as series of files which the users must navigate.

During the past year we have continued support of VisIt on various platforms. Most recently VisIt 2.14 has been released with new functionality leveraged from the DOE SDAV program for more comprehensive tools for flow field analysis. We continue to work to integrate Uintah and VisIt for in-situ analysis and visualization. The current bottleneck is with an MPI communication issue that we continue to investigate.

We continue to work on the closest point sparse octree. The closest point method is a PDE solver, where normal three dimensional stencils can be used to solve partial differential equations on two-dimensional embedded surfaces. By implementing a sparse octree for the closest point embedding, we can scale up to a grid size of 16,384^3. Previous implementations used a two-level grid to save space, but none have scaled to this high of a resolution. Further, to test our implementation, we perform unsteady flow line integral convolution or UFLIC using the sparse octree and the closest point method.

During the past year, we have published a paper on volume rendering and image compositing on GPU accelerated supercomputers to the IEEE Transactions on Visualization and Computer Graphics journal. In this work, we used CUDA and GPU Direct RDMA for image compositing. GPU Direct RDMA allows us to transfer data between GPUs in only one copy operation compared to at least five that would be required if we did not use GPU Direct RDMA . This results in a much faster inter-node communication that is essential for fast image compositing. Also, since GPU Direct RDMA only works from device memory and not texture memory, we have introduced a new rendering pipeline that allows us to render offscreen to an OpenGL buffer object instead of rendering to a texture. This allows us to transfer data from the OpenGL context to the CUDA context without any copy operation. We did strong scaling studies on the Piz Daint supercomputer for up to 4096 GPUs (Piz Daint has a total of 5272 GPUs) for image sizes: 2048x2048, 4096x4096 and 8192x8192 pixels and we got at least comparable performance to CPUs for 2048x2048 images and much better performance on the GPU for 8192x8192 images. We also presented part of this work at the GPU Technology Conference in April 2016.

Algorithms for sort-last parallel volume rendering on large distributed memory machines usually divide a dataset equally across all nodes for rendering. Depending on the features that a user wants to see in a dataset, all the nodes will rarely finish rendering at the same time. Existing compositing algorithms do not often take this into consideration, which can lead to significant delays when nodes that are compositing wait for other nodes that are still rendering. We have developed an image compositing algorithm that uses spatial and temporal awareness to dynamically schedule the exchange of regions in an image and progressively composite images as they become available. Running on the Edison supercomputer at NERSC, we showed that a scheduler-based algorithm with awareness of the spatial contribution from each rendering node can outperform traditional image compositing algorithms. This work will be published at the Eurographics Symposium on Parallel Graphics and Visualization 2016.

<u>Embedded Domain Specific Language (EDSL) and Framework</u>

This year, the EDSL team worked closely with the framework team to achieve large-scale CFD calculations extended to full Titan (over 18,000 GPUs). This involved significant development in the framework to improve efficiency and provide appropriate support for the large number of tasks involved in full CFD calculations. We have demonstrated scalability of basic CFD algorithms to over 18,000 GPUs for:

- Low-mach algorithms with a global poisson solve for pressure. This involves CPU linear solves with GPU kernels for the balance of the calculation.
- Fully compressible algorithm deployed entirely on GPU.

This is directly facilitated by the DSL Nebo, which also has been extended to support boundary conditions, convective flux-limiter treatment, as well as particle-cell interpolation operators on GPU. This proof of concept is the first step toward scalable reacting flow calculations using GPUs within Uintah.

Year-one's review identified linear solvers as an area where we have risk since we rely heavily on them and there are not presently linear solvers available for new architectures such as Titan. We have begun exploring alternative algorithmic approaches for LES of turbulent reacting flow that rely on compressible algorithms which do not require global linear solves. Coupling this with a dual-time integration technique that provides point-implicit treatment together with preconditioning to accelerate convergence in dual-time, we anticipate that we can deploy an algorithm that is competitive with low-mach solvers while being platform-portable.

We have added a fully compressible algorithm to Wasatch in anticipation of the dual-time approach that avoids the need for a pressure solver. This included a complete refactor of the momentum transport abstraction in Wasatch, development of new expressions for compressible flow, addition of new boundary conditions (nonreflecting conditions are not yet implemented). We have demonstrated scalability out to >18,000 GPUs on Titan for the core compressible CFD algorithm. This also involved extension of all Wasatch flow-related expressions to function on a collocated grid rather than the staggered grid previously used for momentum.

The compressible and low-Mach algorithms also have been evaluated for scalability on Mira up to 524,188 cores. For the low-Mach algorithm, acceptable scaling behavior was observed for patch sizes of 32x32x32 and larger, and approximately 80% of the time for each time step was spent in the linear solve. For the compressible algorithm, acceptable scaling behavior was observed for patch sizes of 32x64x64 or larger. Poor scaling results at small patch sizes, were due to an MPI reduction of the time step performed by Uintah

Nebo support for boundary conditions (applying operators over "mask" subsets of points) was hardened and deployed in Wasatch to enable support for boundary conditions on GPUs. This was an important step to demonstrating that the full CFD algorithm could be deployed on GPU.

In support of particle moment methods as well as point-implicit algorithms, the EDSL now supports dense linear algebra applied pointwise over a grid. This includes regression testing and performance testing to ensure efficient serial, multi threaded, GPU performance. These features included the Eigen decomposition and linear solve of lists of matrices and vectors on the GPU.

The EDSL supports backends for serial CPU, multicore CPU and GPU. We have developed and maintained these backends, but have been observing Kokkos development to consider it as a portable backend layer for the DSL. We began an exploratory effort at Kokkos integration for a subset of the EDSL features, and will continue this effort in the coming years. Results at this point are mixed, with Kokkos performing better in some instances and the native EDSL backend performing better in others.

<u>Physics: LES, Multiphase Flow, Particle Combustion, Radiation</u>

<u>Large Eddy Simulation Integration</u>

In general, the efforts of the LES Integration team over the course of the year included:

- coordination with the physics team, enabling the integration of coal and related physics into the LES algorithm;
- coordination with the V/UQ team, aiding trouble-shooting and profiling large production runs;
- development of algorithmic stability for the intended application;
- production code transition in adopting an exascale programming paradigm;
- providing proof of reliability for the main LES algorithm.

This year effort was spent on providing some additional robustness to the verification of the LES algorithm. As part of this effort, the baseline CFD transport of scalars (e.g., mass, momentum, energy) was evaluated for order of accuracy and conservation. In particular, the conservation of energy and mass, when coupled with the particle and wall models, was of particular concern. Test problems were developed and run. It was shown that the code conserved mass and energy in the full coal problem with the conjugate wall heat transfer. In addition, the constant and variable density algorithm was tested using the method of manufactured solutions [1-3]. It was shown that the algorithm formally verified up to third order for constant density problems but remains first order for the higher-order Runge-Kutta method. It has been shown that the benefit of the higher order RK methods was a reduction of overall numerical error of the scheme, which can add stability to the larger production runs. The reduction in overall accuracy for the higher-order RK-schemes has been linked to the pressure projection. There is, however, some work left to be done to accommodate the coal physics in the higher-order RK schemes.

Nightly regression test coverage was evaluated this year. To this end, a tool was developed to analyze the nightly regression suite coverage of the Arches code. The intent of the tool is to aid developers in identifying deficiencies in the nightly testing for existing options and new options as development is occurring. As a result, nightly regression coverage has increased over the course of the year.

Work continued on what has been called the Arches Task Interface, which provides a light weight abstraction layer to minimize Uintah boilerplate for the Arches code and enable easy addition of new physics while offering some flexibility in how algorithmic components are assembled. The interface was improved from a performance standpoint and from a developer's usage standpoint. This work is ongoing.

The Arches team in collaboration with the Computer Science Team has made significant progress on adopting Kokkos into the Arches code. This efforts supplants the previous effort to adopt the center's DSL scheme into Arches. The reasoning is two fold: 1) The adoption of Kokkos into Arches presents less work and 2) Some of the algorithmic/modeling schemes and developer preferences were not well suited for a highly abstracted DSL. Given the long history of the Arches code, the transition to Kokkos constitutes in the easiest terms a simple replacement of the Traditional Uintah iterator loops in the form of a C++ functor or Lambda constructs. As a result, large portions of the coal physics have been reformulated into C++ functors and can be run through the Kokkos framework support. While this work requires some verification, we anticipate an impact on the production runs soon.

A wall-deposition model for the coal ash transformation coupled with the wall heat transfer approach in Arches was also developed. The particle deposition model utilizes the particle viscosity criteria to determine the particle deposition probability and also considers the particle temperature-history and the wall temperature effects. The deposition model coupled with wall heat transfer model on the ARCHES code can predict the formation of powdery deposits on the boiler heating surfaces. The model predicts three regimes for the ash deposits stage on the boiler surface, which strongly impact the prediction of temperature field and heat flux of the wall.

Other points of interest/accomplishments include:

- Reworking of the discrete convective schemes to increase code performance
- Exploring time-splitting methods for DQMOM, including Strang and Lie splitting to enable higher order methods for the coal cases specifically
- Instrumented Arches with runtime statistic modules to enable the gathering of QOI statistics
- In conjunction with the Particle Physics team, CQMOM with the set of required coal physics was demonstrated in a proof-of-concept.

Variable Density Algorithm Development

We have developed a new, consistent variable density, low-Mach pressure projection method. The method is suitable for explicit time integrators and is shown to be robust over a wide range of density ratios up to 1000. This has been thoroughly tested against existing and new MMS and analytical solutions. This work was significantly improved through a week-long workshop hosted at Sandia that involved researchers from the Illinois and Stanford PSAAP centers as well.

Multiphase Flow

During the reporting period there have been several major shifts in responsibility for this subsection. The core responsibility, development and implementation of cutting edge Eulerian quadrature-moment methods, has remained the same. However, due to the learning that has occurred based on the year on V/UQ effort there, and due to personnel shifts the following changes have resulted. First, based on the lessons learned in year one, the primary validation case for which we were responsible — the oxy-fired combustor (OFC) — was dropped from the hierarchy. At the same time, a new responsibility was added — modeling of wall deposition and ash transformations. Also, throughout the year there have been several personnel changes that have dictated the progress in certain aspects of the project. The individual primarily responsible for the development of the new multiphase-flow algorithm, Alexander Abboud, graduated with a Ph.D. in chemical engineering and moved-on to Idaho National Laboratory. As a result technical progress has halted, and emphasis has shifted to recruiting a replacement. After one strong candidate fell through, a second has been identified and we are hopeful that the hire will occur in the next few months.

Our center has been fortunate to have a retained the services of Professor Yuxin “Martin” Wu as a part-time employee year-round and a full-time employee during his leave-of-absence from his home university. The Department of Thermal Engineering, Tsinghua University in China is a most respected institution. During his on-site visit, Martin has assumed responsibility for all simulations of Alstom's boiler simulation facility (BSF) as well as the associated V/UQ. Due to the timing of Martin's visit and overlap with V/UQ, these results will be reported in other sections of this annual report.

For the core multiphase-flow effort, the stream crossing proof of concept was extended to three dimensions with four crossing streams. A pseudo-second order kinetic convection scheme, that takes advantage of both the multi-moment and multi-environmental aspects

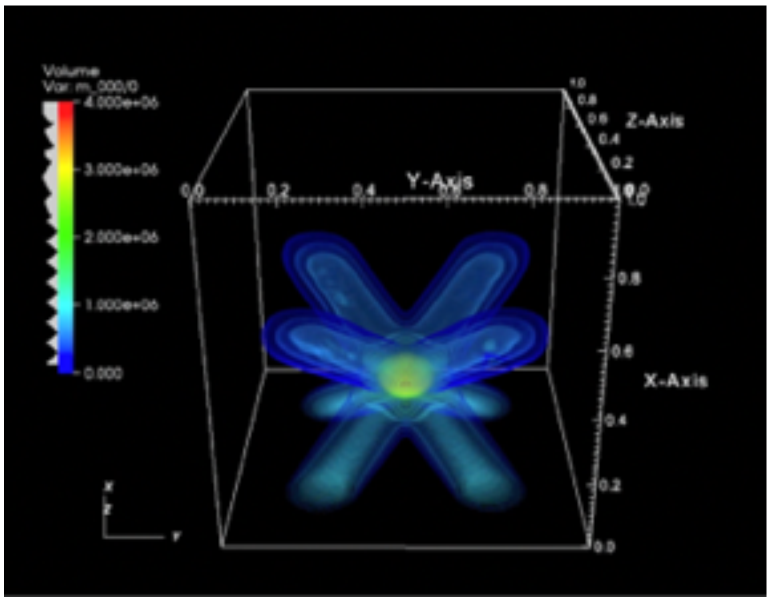

of the approach, was implemented. This greatly reduces the numerical dissipation incurred with the method. A proof of concept has been created demonstrating the methods ability to handle wall reflections in a physically realistic manner. This implementation is general enough to

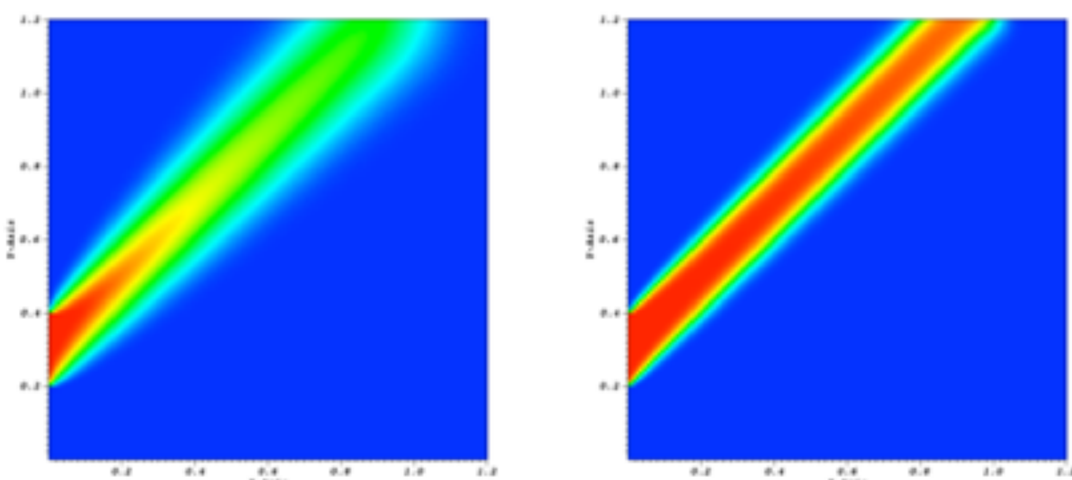

handle a variable restitution coefficient for the impact with the wall. The Arches has found that stair-stepped representation of complex geometry sufficiently accurate for the applications of interest when is applies to the fluid mechanics and scalar transport. However, when it comes to wall reflections stair steps results in a distinctly different reflection behavior compared to an original sloped surface. As such, this new method was extended to handle a reflection that more closely represents the underlying surface — no minor task. Even more, this method was

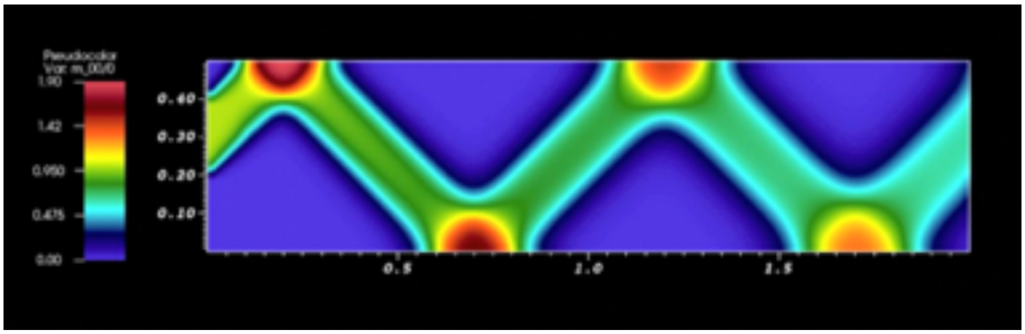

extended beyond its core kinetic capability to simultaneously include scalar properties of the dispersed phase. This extension allowed a small proof-of-concept simulation of the full multi-physics problem for devolatilizing and char oxidizing coal particles. Work needs to continue to scale up this proof of concept to the production boilers.

The ash deposition from coal combustion is of vital importance for the boilers operators. Firstly it affects heat transfer on the boiler, which would deteriorate the thermal efficiency of boiler surface and eventually cause the problem for the boiler operations. Secondly, shedding of a massive deposit would damage the boiler ash hopper. Physically, lots of physical processes are involved in the ash deposition process. Specifically, the deposit mechanism depends on the coal composition, particle size distribution of original mineral matter, thermal conditions and flow dynamics in the gas phase. Moreover, slagging in the molten surface also act the eroding effects on the deposition. The present research aims at using the deposition probability model coupled with the wall heat transfer model on the Uintah code to predict the ash deposition in the pulverized coal boiler. The particle deposition model utilizes the particle viscosity criteria to determine the particle deposition probability and also considers the particle temperature-history and the wall temperature effects. The deposition model coupled with wall heat transfer model on the ARCHES code can predict the formation of powdery deposits on the boiler heating surfaces. This part research successfully predicts three regimes for the ash deposits stage on the boiler surface, which strongly impact the prediction of temperature field and heat flux of the wall.

<u>Lagrangian Particle Transport</u>

A new and simpler algorithm for particle-in-cell interpolation method that reduces the computational cost by 60% is developed, verified and implemented. Previously, we implemented trilinear interpolation for particle-cell interpolants, which maintains second order accuracy for a range of particle sizes. However, for the grid sizes and particle sizes in our coal combustion simulations, this is overkill and we can safely treat the particle as interacting with only one cell at a time. This reduces the cost of particle-cell interpolation by as much as 60-80%.

We have added support in Uintah for injecting particles at runtime, which is required for coal simulations and has been leveraged by other efforts as well. The computing facilities at Argonne National Labs (Mira) were used to investigate the scalability of Lagrangian particle transport within Wasatch. The number of CPU cores used in the scaling study ranged from 1 to 524,288. Patch sizes considered ranged from 16x8x8 64x64x32 to grid cells. Acceptable weak scaling was achieved for patch sizes greater than or equal to 32x32x16 for cases with 1 particle per cell. Excellent weak scaling was observed for all patch sizes considered for 10 particle per cell cases. Good strong scaling was seen for domain sizes equal to or less than 256x128x128 grid cells for 1 particle per cell cases while outstanding strong scaling was achieved for all 10 particle per cell cases. Erratic scaling behavior was observed when both particle and fluid transport were considered. Execution timings indicate that this is due to a spike in the time required for the Hypre matrix solve in certain cases. The exact cause for this is not yet known. In general, both strong and weak scaling improved as the particles per cell was increased.

<u>Coal Devolatilization</u>

Six simple devolatilization model forms were implemented and optimized in Matlab by the BYU team for Sufco coal and compared devolatilization models with CPD model. Evaluative elements included:

- Heating rates: 5×10^3 K/s, 1×10^4 K/s, 1×10^5 K/s and 1×10^6 K/s.
- Developed modified two-step model with distributed activation energy; accurate and less coefficients than RF model.
- Optimization in the works to verify multiple coal types for all model forms.

Literature review of elemental analysis of pyrolysis products (especially char and tar) by BYU is in progress with the goal to formulate a set of correlations to predict product elemental composition from parent coal properties. Lab experiments in progress at BYU to gather data for verification and validation of primary pyrolysis, secondary pyrolysis, and char conversion processes in oxy-fuel conditions for coals of interest in CCMSC simulations.

<u>Coal Char Oxidation and Gasification</u>

As a result of the V/UQ efforts in year 1, we shifted emphasis from devolatilization modeling to char oxidation/gasification modeling. This involved detailed examination of char models ranging from very detailed to relatively simple. These were paired with simple and complex devolatilization and gas-phase chemistry models to assess the effects of various model pairings. Models were assessed using single-particle data from Sandia National Laboratories. Specifically, the nth-order Langmuir Hinshelwood (simple model) and the Char Conversion Kinetics (CCK) (detailed) models were considered. The Chemical Percolation and Devolatilization (CPD) and a two-step model were considered for coal devolatilization. A gas phase chemistry model based on the GRI 3.0 mechanism (detailed kinetics, or DK) was used in addition to infinitely fast gas phase chemistry.

In general, the high-fidelity model combinations were more capable of accurately predicting particle temperature and mass loss. Use of the infinitely fast gas-phase chemistry model often resulted in significant underestimation of the particle temperature, while use of detailed kinetics resulted in overestimation at elevated oxygen concentrations. Choice of gas phase chemistry had little effect on the calculated mass loss of the coal particle. The CCK model produced particle mass histories that matched experimental observations, however it overestimated particle temperatures at elevated oxygen concentrations. Conversely, the simpler Lanmuir-Hinshelwood model predicted particle temperatures fairly accurately, but fell short on particle mass predictions. Use of CPD for devolatilization gave better results for mass predictions compared to the two-step devolatilization model, which overestimated the volatile mass loss. Particle temperature calculations were not sensitive to the chosen devolatilization model.

A new model for the distribution of the liberated heat during char heterogeneous reactions is developed. This model determines the fraction (α) of the liberated heat by heterogeneous reactions (mostly char oxidation) released into the gas phase. The convective heat transfer coefficient, coal heat conductivity, capacity and coal size are incorporated in this model. The new model is evaluated in conjunction with three char oxidation models: first order, Langmuir–Hinshelwood and CCK of oxygen compositions and coal types. The new model is more successful than previous model in capturing particle temperature profiles observed experimentally over a wide range of oxygen compositions and coal types.

The BYU team rewrote the backbone of the Char Conversion Kinetics (CCK) code to facilitate time and tolerance convergence. Components were added to the CCK code (now the CCK\oxy code) to allow it to execute in the extremes of oxyfuel conditions in seconds instead of hours or days.

We completed a global sensitivity analysis of the comprehensive char conversion code by 3 methods (very consistent findings for the most important sub-models in need of attention). We determined that besides the actual chemical kinetic rate coefficients, the char annealing model coefficients were the most sensitive parameters. BYU researchers constructed and calibrated a new model for char-annealing (overwhelmingly the most sensitive sub-model in the comprehensive char conversion model). Literature data were collected, potential forms of the new model evaluated, and the new model is currently in the final stages of testing.

Intern Troy Holland made a presentation to an Industrial Advisory Board at Los Alamos to advise simulation science group leaders in the electricity generation industry on the capabilities and usefulness of both the PSAAP center and the Carbon Capture Sequestration Initiative (CCSI).

<u>Soot in Coal Flames</u>

Arches simulations of the oxy-fired combustor (OFC) with the base soot model were compared with previous RANS simulations. Similar soot and tar concentrations were found between the two cases, as well as the qualitative flame structure. Validation testing of the soot model was performed by comparison to instantaneous line-of-sight soot measurements in the OFC performed by Dale Tree at BYU. Initial comparisons proved promising. New soot oxidation model was derived with parameters tuned to data gathered from eight different published experiments in the literature. New oxidation model includes oxidation by more species than O2 only. New soot gasification model was derived with parameters tuned to data gathered from six different published experiments in the literature. New soot gasification model includes gasification by both CO_2 and H_2O. Soot and tar coded into pressure solvers used by input files in the Arches software. Parameters for both oxidation and gasification models tuned using

Bayesian methods, allowing for better quantification of uncertainty. Probability density functions derived for each model parameter. Use of soot gasification model eliminates the small amounts of soot observed in fuel-lean regions in base soot model.

Radiation

We have been exploring improvements in both the Discrete Ordinates (DO) radiation model and in the Reverse Monte Caro Ray Tracing (RMCRT) model. The improvements made during year two of this contract include both improved accuracy and improved efficiency. The improvements were often made in the form of including multiple options to optimize the options for the specific scenario.

We began by developing a high-resolution 3D integral radiation benchmark for a domain containing hot and cold clouds to help quantify error in RMCRT for multi-grid development. New radiation solver options were added for use with discrete-ordinates. This included using a more efficient implementation of the interface to the hypre library and the shedding of some legacy FORTRAN code. Geometry coarsening for multi-level RMCRT was coded. This multi-level RMCRT algorithm is intended to reduce error as well as improve the speed of the solve relative to the single-level approach.

We have explored the viability of using Uintah infrastructure to handle a sweeping discrete ordinates algorithm. Using the infrastructure/framework as is, we found sweeps to be more efficient up to a scale of approximately 16,000 processors. At this point the DO radiation solution was more computationally efficient when Hypre was used for non-liner solves as opposed to the sweeping algorithm.

During this year we have coupled the RMCRT radiation model to the entire multi-physics model with only a few minor adjustments left outstanding. We have improved RMCRT multi-level approach and have shown that it now converges on the analytical solution for multiple benchmark problems. We found that RMCRT was making unnecessary approximations in sampling the solid angle. Several alternatives were explored and the most effective approach was selected, greatly reducing the errors made by RMCRT in computing the fluxes.

Validation, Verification and Uncertainty Quantification

Vector Consistency

The University of California-Berkeley team developed a new vector consistency measure for datasets involving linear and quadratic surrogate models. This consistency measure can be used to identify a small number of likely incorrect assertions in an inconsistent dataset. They tested the recently developed vector consistency measure on the GRI-Mech dataset for methane combustion, which consists of 77 experiments in 102 variables and is inconsistent. The vector approach identifies 2 experiments as contributing to the inconsistency and suggests relaxations to resolve the dataset. These results confirm prior results found by utilizing the scalar consistency measure.

UCB developed variants and further modifications of the *vector consistency measure* in order to better analyze the properties of a dataset. For instance, introducing weighting allows the user to include scientific opinions regarding the reliability and accuracy of certain experiments when resolving an inconsistent dataset. Evaluating the vector consistency measure for multiple weighting configurations provides different relaxations, i.e. different paths to consistency, and leads to interesting tradeoff problems. This feature enables a much richer framework for consistency analysis. For example, the weighted approach reveals that the GRI-Mech dataset can be made consistent by adjusting only 1 experiment. Among new features, *vector consistency* now allows relaxations to variable bounds, accounting for inconsistency due to improper prior information.

They tested the *vector consistency measure* on the DLR dataset for syngas combustion. This effort was conducted in collaboration with combustion scientists at DLR, Stuttgart, Germany. The DLR dataset contains 83 experiments in 26 variables, with associated uncertainties. Our consistency analysis demonstrates that this dataset is massively inconsistent — 53 of the experimental bounds need to be adjusted to reach consistency. The results showed that the vector consistency measure can provide a more effective and insightful way to solve dataset inconsistency as compared to our previous tools.

They conducted random trials of a linear programming implementation of vector consistency. In these trials, vector consistency was tested on hundreds of thousands of random systems of linear inequalities with known errors. Vector consistency performs significantly better than several comparable strategies, but fails in many cases to exactly recover the error. These results indicate the importance of domain knowledge (i.e. vector consistency with properly assigned weights) in resolving inconsistency and have opened several further avenues for investigation.

Bound-to-Bound Data Collaboration

The UCB team is creating new, reworked version of the Bound-to-Bound-Data-Collaboration (B2BDC) toolbox. The capabilities being implemented include B2BDC calculations for consistency and prediction with both quadratic and rational-quadratic surrogate models. They have included an automated tool to render an inconsistent dataset into a consistent one by expanding the dataset unit bounds detected from sensitivity calculation. This methodology, however, is insufficient to address massively inconsistent datasets, such as the initial DLR example, which can be more effectively analyzed by vector consistency, to be available in a future release.

UCB has developed heuristic rules to minimize computation time in using "redundant" constraints to improve outer-bound results of SDP relaxations to nonconvex quadratic programming programs. This method has been tested on the GRI-Mech dataset for gas-phase methane combustion and achieved a 50% decrease in the optimality gap. The latest release of B2BDC implements this strategy. They have tested the new features of the B2BDC methodology in a collaborative study of H_2/O_2 combustion system with several research groups in combustion.

CCMSC Coal Database

The UCB team has developed an open-source coal database in the PrIMe cyber infrastructure; in collaboration with the UU and BYU teams, archiving 1100 data records for coal devolatilization, char oxidation and nitrogen release. Experiments span across 269 types of solid fuels including fossil, chars, biomass and blends. Using a cloud-based infrastructure, which supports crowd-sourcing, Sandia (Livermore) has contributed additional data records including coal oxidation experiments from their Combustion Research Facility.

They have created tutorials to support users in converting experimental data into the database format. These tutorials were used by students at BYU and UU when adding a large set of char oxidation experiments from Sandia (Livermore). UCB created front-end user application to the CCMSC Coal database to support users by allowing a simple way of finding relevant coal data. This application supports easy browsing, data comparison, plotting, and search queries based on text and quantitative data. The CCMSC Coal Database Application is available through the primekinetics.org website. They added search refinement capabilities to the front-end database application. This gives users the ability to filter the database by multiple search queries, supporting easier access to relevant experimental data that was not previously available.

<u>Reduced Order Model Tools</u>

UCB has created a tool to discriminate between models in QOI space using support-vector machines. A forward additive algorithm was used to find sets of QOIs where two models are overlapping. These tools were applied towards the CPD and DAEM models of coal devolatilization developed by CCMSC.

<u>Experimental Campaign</u>

The majority of our experimental effort has continued to focus on the 1.5 MW pulverized coal test facility (L1500) in order to provide the L1500 simulation team with the best possible data set for the V/UQ task. Efforts to characterize the thermal properties of the deposits coating the various surfaces inside the L1500 continue. Limited data sets of total emissivity and thermal diffusivity were taken prior to the campaign in Year 1. More extensive sets of data were taken in Year 2 after the end of the Year 1 campaign. Four hundred deposit samples were collected in a 1'x1' grid from the burner down to the heat exchanger section (a distance of 45 feet). An FTIR was used to measure spectral reflectivity at room temperature, from which the total emissivity was calculated. To further heighten our understanding of the characteristics of the deposits, we aim to measure the emissivity of the deposits at a range of temperatures comparable to those of the furnace in Year 3. The thermal diffusivity measurements taken in a 1'x1' grid for the first four feet of the reactor were repeated in Year 2 after the campaign to provide a before and after comparison in the deposit properties. Repetitions were increased in order to aid understanding of variance.

The infrared heat flux measurements taken with the high-speed, infrared camera from Year 1 were analyzed. The infrared heat flux was converted to total heat flux measurements to compare with radiometer measurements. The trends of the two measurement techniques matched as the conditions varied; however, the magnitudes of the camera measurements were about half of the radiometer measurements. This has led to an in-depth analysis of the assumptions made in calibrating both the camera and the radiometers and the assumptions made in both instrument models. Work on these instrument models is ongoing.

Extensive planning and efforts have gone into preparing for the Year 2, L1500 campaign. Major modifications to the furnace were made, including the replacement of the first four cooling coils in Sections 1 and 2 with large, 6"x36" rectangular ports. These ports are modular, with the ability to house quartz glass windows for greatly enhanced optical measurements (previous measurements were made through ports with 2.5" diameter ports) or to house soot-blown, cooling panels that are flush with the reactor interior. These panels will provide a background of known temperature for any optical measurements. Much effort was made to calculate the heat loss to these windows/panels in order to design the ports at a size where the heat loss would be similar to that of the former cooling coils. Due to various unavoidable delays resulting from the furnace modifications, the campaign for Year 2 will actually take place in Year 3, in June 2016.

Comments from the peer review team have led us to develop highly detailed instrument models in an effort to better quantify and characterize the uncertainty in our measurements. This pursuit has led to a much more fundamental understanding of what exactly we are measuring, the inherent assumptions associated with these measurements, and the validity of those assumptions. We are currently in the process of creating models for the radiometers, the cooling coils, the newly installed cooling panels, and the infrared camera. These models include methodology for studying in detail the properties and calibration of the components of these diagnostic techniques: the thermocouples that measure the inlet and outlet water temperatures on the L1500, the mass flow meters that measure the flow rate of the cooling water in the L1500, the surface thermocouples used to measure heat flux in the cooling panels and ceiling of the L1500, and the blackbody source used to calibrate the radiometers and infrared camera.

Boiler Simulator Facility

Given the position of the Alstom Boiler Simulator Facility (BSF) within the V/UQ hierarchy, the BSF plays a crucial role in validating the integrated physics of this center. During the last year significant progress has been made in adding more physics-based models, to achieve better consistency with our quantities of interest. The peer review process from year one recommended: "… that some focus be placed on how to transfer properties of the deposition to the wall heating model." There has been a significant focus within the physics and V/UQ groups during the last year in adding appropriate models to represent the deposition of ash, and add the contribution of the ash layer to the energy balance at the wall. In addition, V/UQ efforts in investigating sensitive parameters have highlighted the importance of the char oxidation reactions. Accordingly, we have implemented a char oxidation model, which has the ability to represent some of the physics, which were missing in the previous char oxidization model. Given the extent of work in exploring sensitive parameters in the BSF, and the addition of new physics, year 2 V/UQ consistency analysis proved successful. The BSF showed consistency with 22 O2 measurements, 96 temperature measurements, and 25 heat flux measurements. The consistent region was found with ash slagging temperature between 1490-1580 K (experimental value 1513 K), refractory thermal conductivity between 3.7 and 4.4 W/m/K (experimental value 3.6 W/m/K), and char oxidation activation energy between 20,000 and 26000 cal/mol (experimental value 22,000 cal/mol). Year 1 VUQ provided an effective thermal conductivity for the wall, which was not applicable to the prediction cases. Year 2 V/UQ showed consistency with data given a physics based deposition model (and char oxidation model), and provided insight into the input parameter uncertainty for ash slagging temperature, char oxidation activation energy, and for the specific refractory used in the BSF. The range of consistency in the ash slagging temperature, and activation energy can be used directly in the 8-corner prediction case.

Response to applicable year 2 review comments:

- Consider continuous recording during calculation (e.g. velocity field) to obtain characteristics of turbulent field. Obtain mean profiles across the mixing layer to get a better sense of the flow field. Do this as a matter of course to collect valuable information. This can provide basis for comparison between your use of LES and RANS models that your industrial partners are using.

- It was not clear how data requirements for QOI evaluation compare with requirements for general graphic support/visualization, e.g. vis dump every 100-200 time steps.

During year 2 the physics team implemented a variable statistics module in the production code, which has the ability to collect variable statistics, which are integrated in time. This allows us to analyze time-averages over any window in time while requiring very little I/O. We are interested in analyzing the characteristics of the turbulent field as the reviewers recommend in providing a better basis for comparison between LES and RANS. Using the variable statistics module we are able to easily evaluate the QOI of our system, in computing time-averaged values, which include averaging statistics for every time-step. The I/O frequency of 100-200 time-steps is used for generating volume rendered movies (not for time-averaging QOI).

GE Power 8-corner Unit Simulations

GE Power's 8-corner unit was introduced to the center early in year 2. The 8-corner unit is a 1000MW ultra-super critical air fired PC boiler. The unit is particularly interesting to the center because of the potential impact the center's simulations can have in working with our industrial partner. The details of the geometry for the 8-corner unit, and running conditions were sent to the center near the end of year 2. Preparation work for the simulation has included the creation of a base-case containing all of the pertinent geometry details, and running conditions. In order to run the simulation at a realistic scale a lot of work has been done in getting the production code to scale in terms of radiation solves, I/O, and increasing the speed of the calculation.

Response to applicable year 2 review comments:

- Prioritize the 8-corner simulations versus the design boiler given the availability of experimental data.
- Identify risk mitigation if needed cycles are not available.
- Consider using a higher-order CFD method or document the adequacy of your low-order method.
- Develop a plan for dealing with the loss of INCITE cycles.
- You described two large runs at different resolutions but performance of the runs - with respect to scalability - was not clear.
  - You appear to be addressing performance issues one at a time, e.g. I/O, linear solver.
  - Our experience is bottlenecks will shift as you scale up.

There may have been a miscommunication with regards to the 8-corner unit having experimental data. GE Power has provided us with various steam-side measurements, but no measurements were made on the combustion side. As such, the 8-corner unit simulations are predictions without data.

In order to mitigate risk associated with obtaining cycles, and with the impending completion of the INCITE award, we are profiling our code at scale to determine where additional savings can be made to speed up our computations.

With regards to higher-order CFD methods, we are currently reading about the nLES method to see if there are any significant benefits compared with our current approach.

Although both of the prediction cases (8-corner unit and oxy-coal boiler) are large, there are significant difference between the geometries and some of the models used between the cases. As such, the scaling performance varies between the cases. In particular, the 8-corner unit has hundreds of internal intrusions objects representing platens and pipes, whereas the oxy-design boiler design is targeting heat removal only through the walls. These internal intrusions affect the fluid mechanics and thermal boundaries of the system and as results the linear solve times for both the pressure solve and the DO radiation solves appear to take much longer. We are currently profiling the production code to determine if there is anything that can be done to decrease the computation time associated with the linear solves.

Awards to key persons

- Valerio Pascucci received the Distinguished Mentor Award from The Graduate School of the University of Utah, May 6, 2016.
- Valerio Pascucci received Best Paper Award 2016 from the IEEE Pacific Visualization Symposium.
- Valerio Pascucci received the Integrated Partnership Award from DOE Federal Laboratory Consorium for Excellence in Technology Transfer, April 29, 2015.

Students and Post-doctoral Fellows

- Michael Brown, MS candidate, University of Utah
- Oscar Diaz-Ibarra, Ph.D. candidate, University of Utah
- Teri Draper, Ph.D. candidate, University of Utah
- Pavol Klacansky, Ph.D. candidate, University of Utah
- Lauren Kolcynski, Ph.D. candidate, University of Utah
- Joshua McConnell, Ph.D. candidate, University of Utah
- John Camilo Parra Alvarez, Ph.D. candidate, University of Utah
- Steve Petruzza, 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
- Troy Holland, Ph.D. candidate, Brigham Young University
- 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
- Babak Goshayeshi, Post-doctoral Fellow, University of Utah

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

- Philip, S., B. Summa, J. Tierny, P. Bremer, and V. Pascucci. "Distributed seams for gigapixel panoramas." Visualization and Computer Graphics, IEEE Transactions on, 21(3):350-362, March 2015.
- Liu, S., B. Wang, J. J. Thiagarajan, P.-T. Bremer, and V. Pascucci. "High-dimensional visualization: Visual exploration of high-dimensional data through subspace analysis and dynamic projections." Computer Graphics Forum, 34(3):271-280, June 2015.
- Edwards, J., E. Daniel, V. Pascucci, and C. Bajaj. "Approximating the generalized Voronoi diagram of closely spaced objects." Computer Graphics Forum, 34(2):299-309, May 2015.
- Bremer,P.-T., D. Maljovec, A. Saha, B. Wang, J. Ganey, B. Spears, and V. Pascucci. "ND2AV: N-dimensional data analysis and visualization analysis for the national ignition campaign." Computing and Visualization in Science, 17(1):1{18, 2015.
- Summa, B., A. A. Gooch, G. Scorzelli, and V. Pascucci. "Paint and click: Unified interactions for image boundaries." Computer Graphics Forum, 34(2):385-393, May 2015.
- Bhatia, H., B. Wang, G. Norgard, V. Pascucci, and P.-T. Bremer. "Local, smooth, and consistent Jacobi set simplification." Computational Geometry: Theory and Applications, 48(4):311 - 332, 2015.
- Earl, C., M. Might, A. Bagusetty, J.C. Sutherland. (2016) "Nebo: An efficient, parallel, and portable domain-specific language for numerically solving partial differential equations." Journal of Systems and Software. (In press) doi:10.1016/j.jss.2016.01.023

- Saad, T., and J.C. Sutherland, "Wasatch: an Architecture-Proof Multiphysics Development Environment using a Domain Specific Language and Graph Theory," Journal of Computational Science (2016). In press. doi:10.1016/j.jocs.2016.04.010
- Gyulassy, A. Knoll, K. C. Lau, B. Wang, P.-T. Bremer, M. E. Papka, L. A. Curtiss, and V. Pascucci. “Interstitial and interlayer ion diffusion geometry extraction in graphitic nanosphere battery materials.” IEEE Trans. Vis. Comput. Graph, 22(1):916-925, 2016.
- Big Data and High Performance Computing. Lucio Grandinetti, Gerhard Joubert, marcel Kunze, Valerio Pascucci, editors. IOS press. October, 2015.
- McConnell, J., & Sutherland, J. C. (2015). “The Effect of Model Fidelity on Prediction of Char Burnout for Single-Particle Coal Combustion.” In Western States Section of the Combustion Institute (pp 1-15). Provo, UT.
- Cyr, E. C., E. Phipps, M.A. Heroux, J. Brown, E.T. Coon, R.C. Kirby … R. Trott. (2015) “Algorithms and Abstractions for Assembly in PDE Codes: Workshop Report.” SAND2015-1379, Albuquerque, NM.
- Skraba,P., P. Rosen, B. Wang, G. Chen, H. Bhatia, and V. Pascucci. “Critical point cancellation in 3d vector fields: Robustness and discussion.” IEEE Trans. Vis. Comput. Graph, May, 2016.
- Richards, A. and T.H. Fletcher. “A Comparison of Global Kinetic Models for Coal Devolatilization.” (In review)
- Holland, T., T.H. Fletcher, et al. “Bayesian Uncertainty Quantification and Calibration of Char Thermal Annealing.” (In preparation)
- Holland, T., T.H. Fletcher, et al. “Global Sensitivity Analysis for a Comprehensive Char Conversion Model in Oxy-fuel Conditions.” (In preparation)
- Josephson,A. , D.O. Lignell, A. Brown, and T.H. Fletcher. “Revisions to Modeling Soot Derived from Pulverized Coal.” (In preparation)
- 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 (In press).
- Oreluk, J., C.D. Needham, S. Baskaran, S.M. Sarathy, M.P. Burke, R.H. West, M. Frenklach and P.R. Westmoreland. “Dynamic Chemical Model for H2/O2 Combustion Developed Through a Community Workflow.” (In preparation)
- Slavinskaya, N.A., J.H. Starcke, M. Auyelkhankyzy, U. Riedel, W. Li, J. Oreluk, A. Hedge, A. Packard and M. Frenklach. “Development of an UQ-Predictive Chemical Reaction Model for Syngas Combustion.” (In preparation for submission to Energy & Fuels)

Presentations

- Holland, T. and T. H. Fletcher. “Coal Particle Combustion.” Poster presented at the Stewardship Science Academic Programs (SSAP) Symposium, March 11-12, 2015, Santa Fe, NM.
- Richards, A. and T. H. Fletcher. “A modified two-step model of devolatilization.” Presented at the 15th International Conference on Numerical Combustion, April 19-22, 2015, Avignon, France.
- Lignell, D. O., B. J. Isaac, A. Josephson, T. H. Fletcher and J. Thornock. "Large eddy simulation of soot formation in an oxy-coal combustor." Presented at 15th International Conference on Numerical Combustion, April 19-22, 2015, Avignon, France.

- Josephson, A. J., B. J. Isaac, D. O. Lignell and T. H. Fletcher. "Large eddy simulation of an oxy-coal combustor." Presented at the 9th Meeting of the U.S. Joint Sections of the Combustion Institute, May 17-20, 2015, Cincinnati, OH.
- Richards, A. and T. H. Fletcher. "A Comparison of Global Kinetic Models for Coal Devolatilization." Presented at the Western States Section of the Combustion Institute, October 5-6, 2015, Brigham Young University, Provo, UT.
- Holland, T. and T. H. Fletcher. "Comprehensive Char Conversion Global Sensitivity Analysis." Presented at the Western States Section of the Combustion Institute, October 5-6, 2015, Brigham Young University, Provo, UT.
- Fry, A., J. Spinti, I. Preciado, O. Diaz-Ibarra and E.G. Eddings. "Pilot-scale investigation of heat flux, radiation and CO distribution from an oxy-coal flame." Paper presented at the 2015 American Institute of Chemical Engineering (AIChE) Annual Meeting, November 9-13, 2015, Salt Lake City, UT.
- Fry, A., J. Spinti, I. Preciado, O. Diaz-Ibarra and E.G. Eddings. "Pilot-Scale Investigation of Heat Flux and Radiation from an Oxy-coal Flame." Presented at Clearwater Coal Conference, May 31-June 4, 2015, Clearwater, FL.
- Diaz-Ibarra, O., J. Spinti, A. Fry, J.N. Thornock, B. Issac, 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." Paper presented at the American Flame Research Committee 2015 Industrial Combustion Symposium, September 9-11, 2015, Salt Lake City, UT.
- Fry, A., J. Spinti, I. Preciado, O. Diaz-Ibarra and E.G.Eddings. "Pilot-scale investigation of heat flux and radiation from an oxycoal flame." Paper presented at the American Flame Research Committee 2015 Industrial Combustion Symposium, September 9-11, 2015, Salt Lake City, UT.
- Fry, A., J. Spinti, I. Preciado, O. Diaz and E.G. Eddings. "Heat Flux, Radiation and CO Distributions from a 1.5 MW Oxy-Coal Flame." Presentation at the 2015 American Institute of Chemical Engineering (AIChE) Annual Meeting, November 9-13, 2015, Salt Lake City, UT.
- Draper, T., E.G. Eddings and T. Ring. "Thermal Characterization of Ash Deposits in a 1.5 MW Reactor." Presentation at the 2015 American Institute of Chemical Engineering (AIChE) Annual Meeting, November 9-13, 2015, Salt Lake City, UT.
- Fry, A., J. Spinti, I. Preciado, O. Diaz and E.G. Eddings. "Predicting Heat Transfer Characteristics of a 1.5 MWth Oxy-Coal Flame." Paper presented at 5th Meeting of the IEAGHG International Oxyfuel Combustion Research Network Meeting, October 27-30, 2015, Wuhan, China.
- Diaz-Ibarra, O., J. Spinti, A. Fry, B. Schroeder, J. Thornock, B. Issac, D. Harris, M. Hradisky, S. Smith, E. Eddings, and P. Smith. "A Validation/Uncertainty Quantification Analysis of a 1.5 MW Oxy-Coal Fired Furnace." Paper presented at the American Flame Research Committee (AFRC) 2015 Industrial Combustion Symposium, September 9-11, 2015, Salt Lake City, UT.
- Fry, A., J. Spinti, I. Preciado, O. Diaz and E.G. Eddings. "Pilot-scale Investigation of Heat Flux and Radiation from an Oxy-coal Flame." Paper presented at the American Flame Research Committee (AFRC) 2015 Industrial Combustion Symposium, September 9-11, 2015, Salt Lake City, UT.
- Fry, A., J. Spinti, O. Diaz, I. Preciado and E.G. Eddings. "Predicting Heat Transfer Characteristics of a 1.5 MWth oxy-coal Flame." Paper presented at The 40th International Technical Conference on Clean Coal & Fuel Systems, May 31 - June 4, 2015, Clearwater, FL.

- Draper, T., P. Toth, T. Ring and E. Eddings. “Optical heat flux and temperature measurements on a 100 kW oxy-fuel combustor." Paper presented at the 7th European Combustion Meeting (ECM 2015), March 30 - April 2, 2015, Budapest, Hungary.
- Isaac, B., J. Thornock, S.T. Smith and P. Smith. “Simulation and Validation of 15 Mwth Oxy-Coal Power Boiler.” Paper presented at 2015 American Institute of Chemical Engineering (AIChE) Annual Meeting, November 9-13, 2015, Salt Lake City, UT.
- Sutherland, J. C., J. McConnell, and B. Goshayeshi. “An Assessment and Comparison of Various Coal Combustion Models.” Paper presented at First International Oxyflame Workshop, February 10-11, 2016, Montabaur, Germany.
- Saad, T., C. Earl, A. Bagusetty, M. Might, and J.C. Sutherland. (2015) “Uintah/Wasatch: Addressing Multiphsyics Complexity in a High-Performance Computing Environment.” Presented in SIAM Computational Science and Engineering Conference, March 14-18, 2015, Salt Lake City, UT.
- Saad T., A. Bagusetty and J.C. Sutherland. “Wasatch: A CPU/GPU-Ready Multiphysics Code Using a Domain Specific Language.” Poster presented in SIAM Computational Science and Engineering Conference, March 14-18, 2015, Salt Lake City, UT.
- Saad, T., and J.S. Sutherland. “Nebo: an Embedded Domain-Specific Language for Platform-Agnostic PDE Solvers.” Presented at the 24th International Conference on Parallel Architectures and Compilation Techniques, October 18-21,2015, San Francisco, CA.
- Saad, T. and J.C. Sutherland. “Wasatch: Addressing Multiphysics and Hardware Complexity in a High-Performance Computing Environment, in Workshop on Software Development Environments for High-Performance Computing.” Presented at the 24th International Conference on Parallel Architectures and Compilation Techniques, October 18-21, San Francisco, CA.
- Frenklach, M. and A. Packard. “Data-Model System Analysis and Prediction under Uncertainty.” Princeton University, Department of Chemical and Biological Engineering, October 7, 2015, Princeton, NJ.
- Frenklach, M. and A. Packard, “Active Subspace Identification and Utilization,” West Coast ROM Workshop, Sandia National Laboratories, November 19, 2015, Livermore, CA.
- Bhatia, H., A. G. Gyulassy, V. Pascucci, M. Bremer, M. T. Ong, V. Lordi, E. W. Draeger, and J. E. P. andq Peer-Timo Bremer. “Interactive exploration of atomic trajectories through relative-angle distribution and associated uncertainties.” In Proceedings of IEEE Pacific Visualization Symposium. IEEE Computer Society, 2016.
- Landge, A., S. Kumar, V. Pascucci, I. Rodero, M. Parashar, and P.-T. Bremer. “Evaluation of in-situ analysis strategies at scale for power efficiency and scalability.” In Proceedings of 16th IEEE/ ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID). IEEE Computer Society, 2016.
- Maljovec, D., B. Wang, P. Rosen, A. Alfonsi, G. Pastore, C. Rabiti, and V. Pascucci. “Topology-inspired partition-based sensitivity analysis and visualization of nuclear simulations.” In Proceedings of IEEE Pacific Visualization Symposium. IEEE Computer Society, 2016.
- Pascucci, V. Invited talk at the Salishan Conference, “Extreme Data Management Analysis and Visualization for Exascale Supercomputers April 28, 2016.

- Pascucci, V. Plenary talk at XSEDE 15: "Extreme Data Management, Analysis and Visualization: Exploiting Large Data for Science Discovery." July 28, 2015, St. Louis, MO.
- Valerio Pascucci, Plenary talk at XSEDE 15: "Extreme Data Management, Analysis and Visualization: Exploiting Large Data for Science Discovery." July 28, 2015, St. Louis, MO.
- Livnat, Y. Invited talk at The Salishan Conference on HIGH-SPEED COMPUTING: "Towards Interactive Analysis and Exploration of the HPC Performance Landscape." April 30, 2015.
- Landge, A. Poster at The Salishan Conference on HIGH-SPEED COMPUTING: "Scalability and Power Efficiency of In-Situ Analysis Workflows" April 30, " 2015.