Annual Technical Report

DE-NA0002375

The University of Utah
Institute for Clean and Secure Energy

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

April 1, 2018 through March 31, 2019

May 31, 2018

Annual Technical Report
DE-NA0002375
April 1, 2018 through March 31, 2019
Carbon Capture Multidisciplinary Simulation Center

The Carbon Capture Multidisciplinary Simulation Center (CCMSC) was established to demonstrate positive societal impact of extreme computing by deployment of low-cost, low-carbon energy solution for power generation. The overall strategy includes collaboration with our industrial partner, General Electric Power, with an inter-disciplinary focus on development of high-performance computing technology. Three teams contribute to the overarching predictive science: the computer science team, the physics team and the validation/UQ team.

The Center has been driven by the mission of predicting the heat flux profile for the design of a new technology for a full-scale pulverized solid-fuel (coal) thermal power generation boiler to a proven level of uncertainty using large-eddy simulations (LES) on the largest computational resources available to us. In FY19 the center demonstrated:

- Improvement in our uncertainty quantification process for extrapolating uncertainty from the pilot scale data to the full-scale prediction by identifying scenario parameter uncertainty and by propagating not only model parameter uncertainty but also model form uncertainty into the prediction.
- Reduced the model form bias in our overarching simulation from 30% to 8% by identifying the source of the largest bias to be in the ash deposition models, and thus reducing that bias by improving the quality of the deposit models. This model improvement included models for deposit particle sintering and the effect of sintering on surface emissivity.
- Demonstration that the Asynchronous Many Task (AMT) runtime system (Uintah) makes it possible to run complex, multi-phase, multi-physics applications such as the Atikokan boiler using Arches at the largest processor counts available to us.
- Additionally, we have made the most computationally challenging task, thermal radiation via ray tracing, to strong and weak scale on the same architectures and have shown how performance portability is achieved by using the Kokkos system via a machine independent loop layer in Uintah.
- The PIDX data I/O library has matured from a proof of concept prototype to a fully integrated and supported file format within Uintah.
- Incorporation of in situ visualization in to the Uintah framework using VisIt's in situ interface, libsim, by expanding Uintah's runtime infrastructure in the form of a more centralized collection of performance data, simulation parameters, debugging mechanisms, and runtime controls and presenting this collection to users via a simulation dashboard, and providing them with the ability to do interactive parameter exploration, visual debugging, and computational steering.
- A coupled simulation-machine layout making it possible to identify computational bottlenecks that would otherwise be difficult to diagnose.

OUTREACH AND EDUCATION

Current Students and Post-doctoral Associates

- Teri Draper, Ph.D. candidate, University of Utah
- Jebin Elias, 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
- Qi Wu, Ph.D. candidate, University of Utah
- William Usher, Ph.D. candidate, University of Utah
- Damodar Sahasrabudhe, Ph.D. candidate, University of Utah
- Kaitlyn Scheib, M.S. candidate, University of Utah
- Kamron Brinkerhoff, M.S. 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
- Oscar Diaz-Ibarra, post-doctoral research fellow, University of Utah
- John Camilo Parra Alvarez, post-doctoral research fellow, University of Utah

Post-degree CCMSC Students' Employment

- Christopher Earl, staff member, Lawrence Livermore National Laboratory
- Ben Schroeder, staff member, Sandia National Laboratory-Berkeley
- Alexander Abboud, staff member, Idaho National Laboratory
- Aaditya Landge, software engineer, Twitter
- Pascal Grosset, staff member, Los Alamos National Laboratory
- Mark Kim, staff member, Oak Ridge National Laboratory
- Troy Holland, post-doctoral associates, Los Alamos National Laboratory
- Oscar Diaz-Ibarra, post-doctoral research fellow, University of Utah
- John Camilo Parra Alvarez, post-doctoral research fellow, University of Utah
- Siddharth Kumar, asst. professor, Computer Science, Univ. of Alabama, Birmingham
- Siddartha Ravichandran, staff member, Expedia
- Michael D. Brown, staff, Hi-Rez Studios
- Alex Josephson, post-doctoral fellow, Los Alamos National Laboratory
- Babak Goshayeshi, senior scientific software engineer, Merck & Company
- Alex Josephson, postdoctoral researcher, Los Alamos National Laboratory
- MinMin Zhou, senior engineer, Reaction Engineering International
- Daniel Gunderson, engineer, Big West Oil Refinery

Education

A full semester training course was presented to Center personnel interested in learning hybrid validation, uncertainty quantification, and machine learning methods employed and/or developed at the center. Class projects included center deliverables including a range of instrument models used across the center for validation of the center hierarchy.

Internships and Staff Lab Visits

Kaitlyn Scheib completed a student internship at Los Alamos National Laboratory under the mentorship of Dr. Gowri Srinivasan. She worked on a project that runs simulations of flow through discrete fracture networks in subsurface rock. These simulations have many different applications including fracking, nuclear waste disposal, geothermal energy and CO_2 sequestration. Kaitlyn learned about Machine Learning (ML) methods and performed sensitivity analyses of different parameters on the breakthrough times from the simulated fracture networks. She used regression to predict the breakthrough times from given input parameters for a fracture network.

Mokbel Karam joined Dr. Fady Najjar and Dr. Ming Jiang for an internship in Machine Learning (ML) at Lawrence Livermore National Laboratory during summer 2018. Mokbel's focus was on assimilating basic knowledge in machine learning algorithms and how they can be applied to fluids. In particular, he investigated the predictive capabilities of several machine learning models applied to the Sedov-vonNeumann-Taylor blast wave problem, including Polynominal Regressors, Random Forest with Decision Trees Estimators and Multilayer Neural Networks, using Tensorflow. He has recently changed projects under his faculty advisors, James C. Sutherland and Tony Saad.

Damodar Sahasrabudhe completed his student internship at Sandia National Laboratory in summer 2018 working with Eric Phipps and Sivasankaran Rajamanickam. Damodar focused on creating a portable version of the SIMD primitive using the existing KNL back-end from the stk package in Trilinos and created a new CUDA back-end for the primitive. To test it, three kernels were converted – gemm, spmv and Uintah's charoxidation kernel – to use the portable SIMD primitive and were able to run the same code on a CPU and a GPU.

Will Usher and Duong Hoang worked at LLNL with Peter Lindstrom on particle compression strategies using ZFP during summer 2018.

Pavol Klacansky also spent summer 2018 at LLNL where he worked on a prototype VR tool for inspection of a scanned 3D printed parts. The prototype allows visualization of the data and provides an ability to measure distances and angles. The user can also mark defects for later inspection.

Danny Gunderson completed his internship with Russel Whitesides at Lawrence Livermore National Laboratory from September to December 2018. His project dealt with using the latest data science techniques to determine which physical features were most helpful in predicting octane number (RON). Some of his findings were that the ignition delay time and fuel density were the most predictive in determining RON. He reported that his efforts were helpful in the ongoing efforts to create a closed loop optimization problem to find the best fuel mixtures that result in synergistic octane boosts.

Arun Hegde enjoyed a summer internship at Sandia National Laboratory in Albuquerque working with Bart Van Bloemen Waanders assembling a forward model for direct write additive manufacturing based on the peridynamic theory. Additionally, Arun is investigating the impact of using linear surrogates in situations where the ground truth is quadratic.

Jim Oreluk completed his internship with Dr. Habib Najm at Sandia National Laboratory at Livermore. He worked on two projects. One project dealt with parameter estimation for dynamical systems and the second one investigated a probabilistic embedding of model error, parameterized by random field. He continued work on identifying potential model discrepancy functions and their generalizability for quantum chemistry predictions.

Allen Sanderson ***(research computer scientist)*** visited and presented at Sandia National Laboratory, Los Alamos National Laboratory Summer School, and Argonne National Laboratory this year. He visited LLNL in October to work with the VisIt team on ***in situ*** and visualization infrastructure. The collaboration with LLNL assisted in the formalization of the triggers and deployment of the selection inclusion lists. The lists allow users to quickly sub-select large groups of data, such as viewing patches that are on a particular MPI rank or node.

Ben Isaac ***(professional engineer)*** and Jeremy Thornock ***(research professor)*** spent a week at LLNL at the invitation of Greg Burton to discuss the work at Utah's PSAAP Center with lab personnel. Two presentations were given and efforts were made to develop new relationships and to work on the NLES closure in Arches.

COMPUTER SCIENCE: EXASCALE RUNTIME, VISUALIZATION, I/O

PIDX support for particle data

During the year we have continued to make progress on the development of new features for particle data and I/O support in PIDX. In particular, we are experimenting with different data aggregation strategies and comparing with IOR benchmark performance.

Sub-filing and two-phase I/O are commonly employed aggregation techniques to balance the trade-off between file per-process and single file (i.e., collective) I/O. However, existing I/O systems typically treat particle data as a stream of bytes, and ignore spatial correspondence in the data when applying sub-filing and two-phase I/O. We tackle these challenges by first building a correspondence between a particle's spatial location and its position in the file. This way we built a two-phase I/O for particle data that maintains spatial locality of the particles on disk.

The following figures report a series of experiments run on Mira and Theta writing 32K and 64K particles per core using different aggregation configuration (e.g., 2x2x4 stand for aggregating every 2 cores over the X axis, 2 over Y and 4 over Z). Our two-phase I/O outperform file per process and shared file I/O performance with IOR and HDF5, while writing data in a layout that is more suitable for analysis and visualization tasks (i.e., maintaining spatial locality).

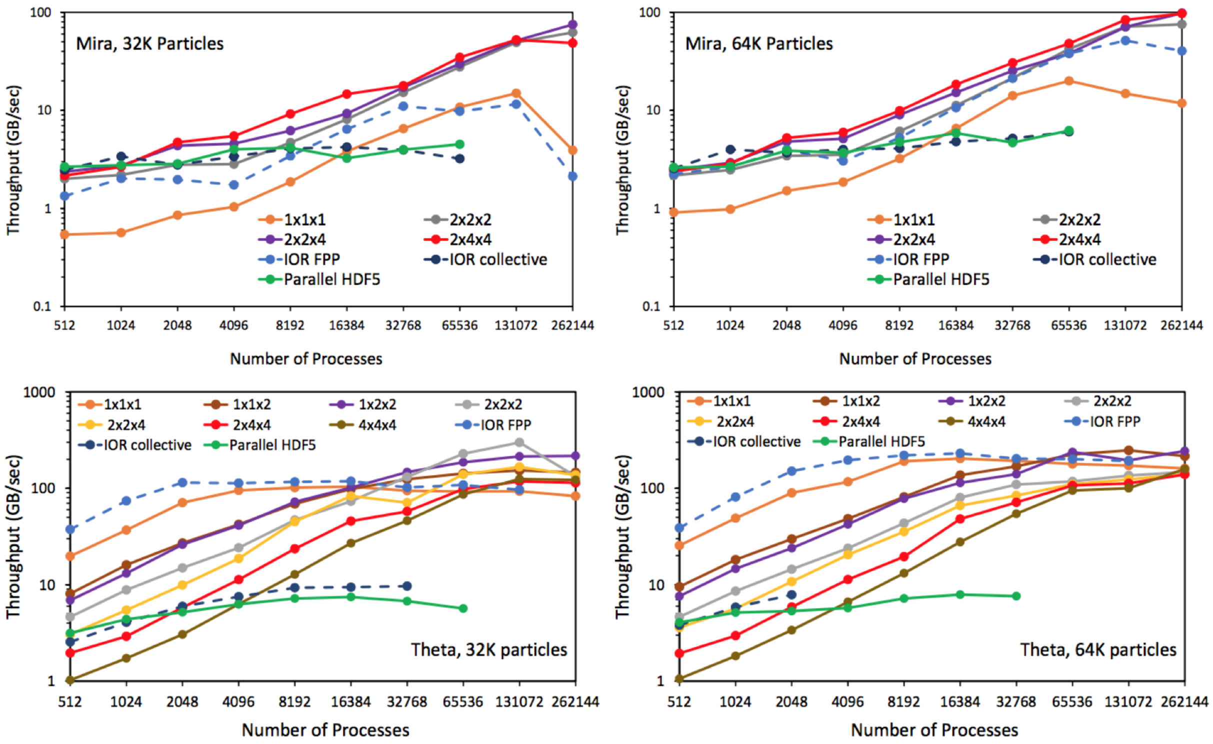

Figure 1. Computational Experiments on Mira and Theta

Furthermore, we performed read experiments on a smaller number of cores to demonstrate the importance of spatial locality and metadata in our data format. In the following figure, we report the result of those experiments performed on Theta and on an SSD equipped workstation. Our approach, using spatial correspondence, provides good strong scaling and read performance, even at far fewer cores than the data was written with. On the contrary, reading from the formats with poor or no spatial locality are significantly slower and their performance decrease at scale.

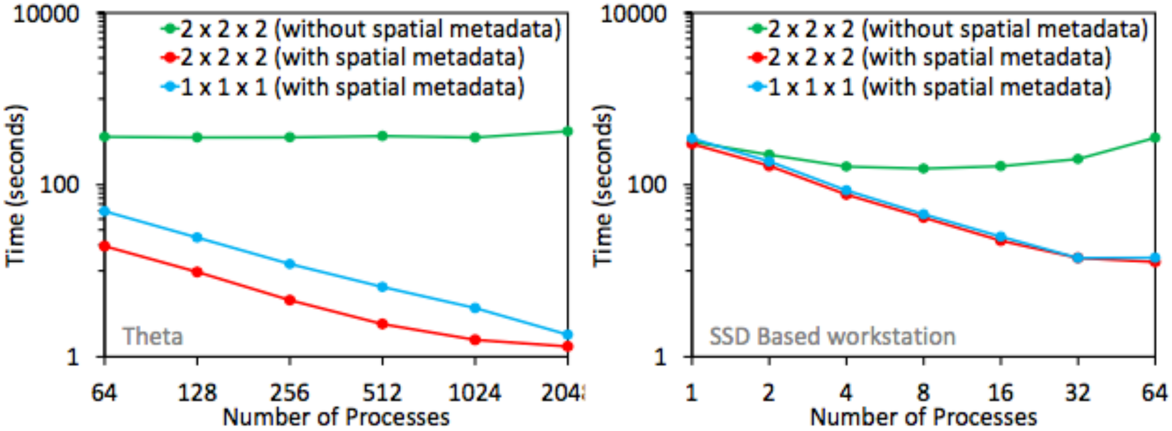

Figure 2: Computational Experiments on Theta and on an SSD equipped workstation.

Our I/O system for particle data can also reorganize the data into Level of Details (e.g., groups of particles identifying different level of resolution of the data). For example, we perform a random reshuffling of the particles before writing them to disk and we are able to access different levels progressively. We experiment reads at different Level of Details (LOD) on the same two platforms (Theta and SSD equipped workstation) demonstrating that with our I/O system LOD reads can be performed efficiently and provide interactive data access. In the following figures, we report LOD read performance on an SSD Based Workstation (a) and Theta (b) using 64 cores to read progressively higher levels from a 2 billion particle dataset. Our ordering allows fast reads for low levels of detail and does not increase significantly until reading large portions of the dataset.

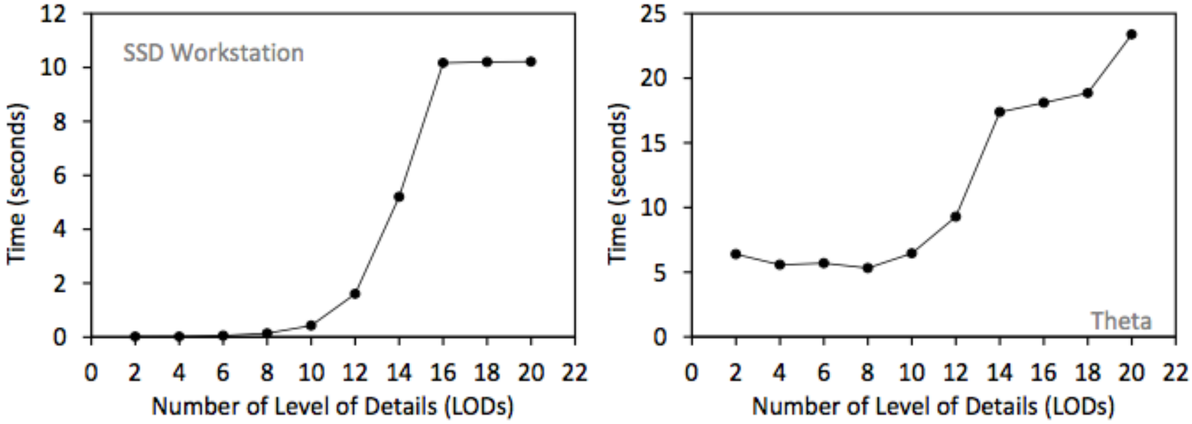

Figure 3: LOD read performance on an SSD Based Workstation (a) and Theta (b) using 64 cores.

Finally, we provide in the following figure a closer look at the visualization results to show the practical effectiveness of the approach for a specific zoomed-in region for a coal particle injection simulation dataset with 55 million particles, written using a random reshuffling for level of detail. In particular, the figure shows four visualizations obtained while progressively more particles are read, from (a) 25% to (d) 100% of the data. Lower resolutions using this LOD ordering can still provide a good representation of the data and be read quickly using our I/O strategy.

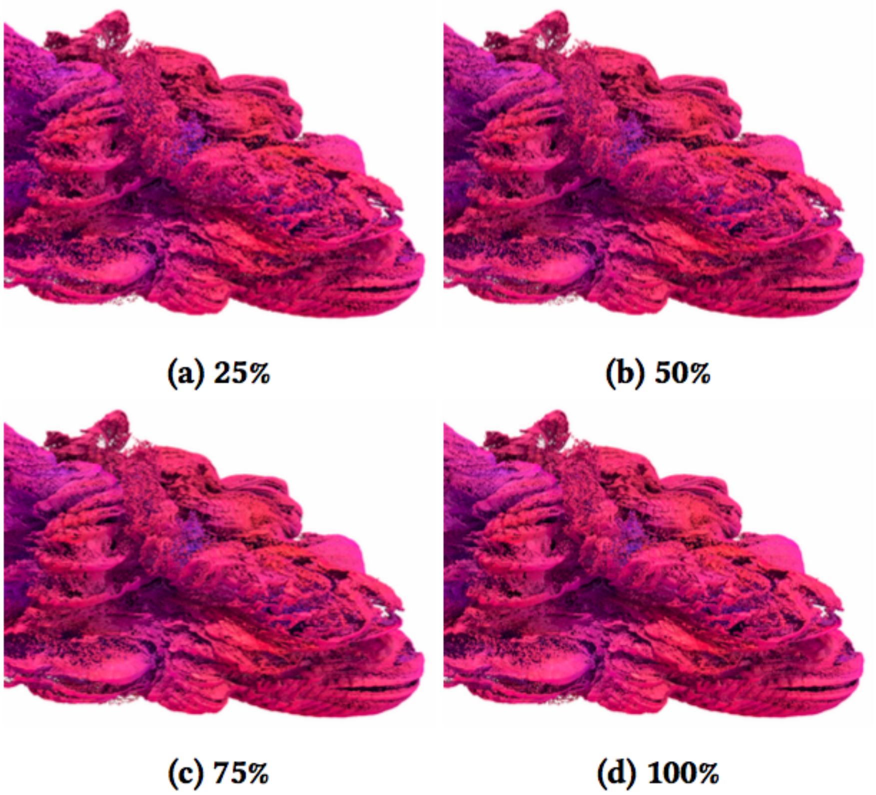

Figure 4: Visualization Results for Particle Injection

<u>PIDX deployment and integration with Uintah and VisIt</u>

PIDX deployment has been improved to support Windows platforms. Continuous integration tools are used to build and test the library at every commit on the GitHub repository on every major OS platform (i.e., OSX, Linux, Windows). The compression library “zfp”, used by PIDX for lossy compression, has been updated to the latest version and integrated as a git submodule in order to improve traceability and maintainability of the code. Two new versions of the PIDX library had been released adding support for integer compression, enhanced data partitioning, templated time series and particle data two phase I/O.

The Uintah build system can automatically build, link and set PIDX to be used as I/O system. New scripts have been developed to allow generation of global metadata files which represent a collection of several restarts of the simulation. The PIDX library together with Uintah and VisIt with the latest IDX reader have been deployed on Pascal and Surface and used, by all ARCHES’s users in the Uintah team, for several simulation runs. We supported the simulation team in their runs, discussing and fixing bugs related to metadata (e.g., how to produce comprehensive time series to make a movie), visualization (e.g., ghost cells and extra cells management), data aggregation (e.g., how to limit the number of files produced) and data streaming.

The following figure depicts one timestep of an Uintah simulation saved in PIDX and rendered using VisIt on Pascal.

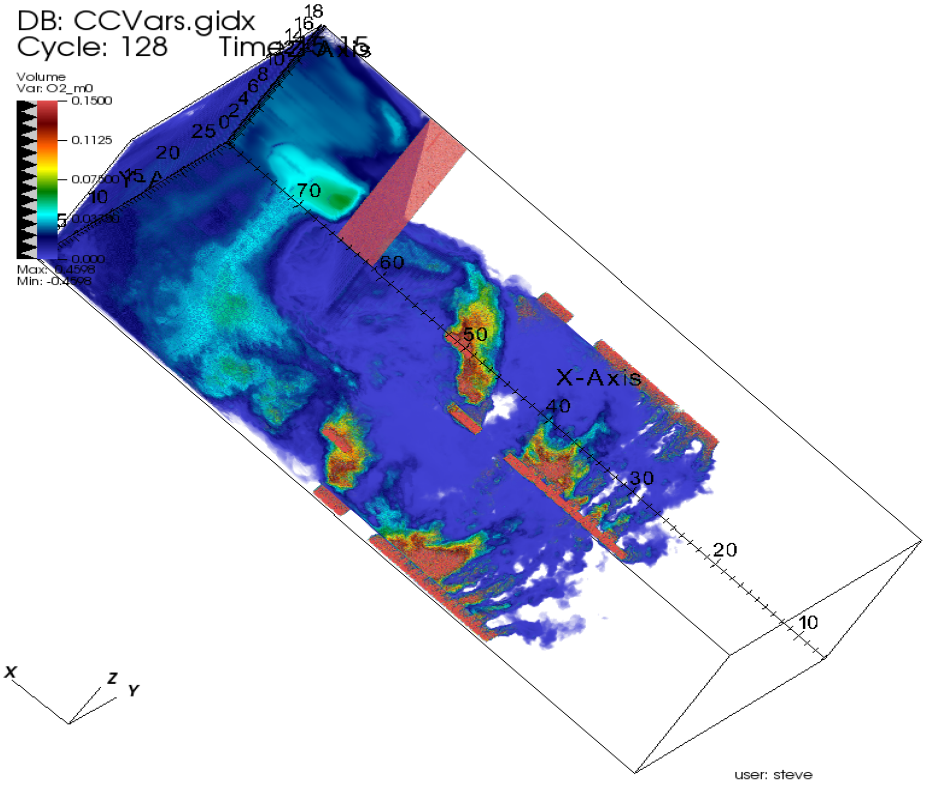

Figure 5: One timestep of an Uintah simulation saved in PIDX and rendered using VisIt on Pascal.

The latest VisIt IDX reader has been merged into the new GitHub repository of the VisIt visualization framework and will be distributed in the next 3.0 release.

Topology-based Analysis

Topology-based approaches have proven particularly useful in studying properties of combustion. For instance, Reeb graphs, contour trees, and merge trees are used to identify ignition kernels, track the merging of flame components, and extract vortical structures. Morse-Smale complexes are used to identify ridge-surfaces corresponding to flame fronts, dissipation elements that measure turbulence scales, and partition the flame into zones. We have focused on scaling the underlying computation of the topological data structures to scales relevant to analyze combustion data from the PSAAP use case. First, we developed a shared-memory parallel approach to extract Morse-Smale complexes with the accuracy needed to analyze flame fronts, improving the speed by a factor of 30x and size of data that could be addressed by 10x over previous approaches. This approach was published in IEEE Visualization 2018 and presented at the conference. We also developed a new data structure for rapid computation of merge trees, that improves by an order of magnitude the speed of computation on shared-memory parallel systems and has been demonstrated on two orders of magnitude larger data than previously possible. This work is currently under review for IEEE Visualization 2019. It also marks a launching-off point for faster in-situ topological queries, as the localized data structure avoids the V-cycle communication pattern usually used to resolve large-scale topological features. We are working on extending the localized data structure to support adaptive precision and resolution grids with dynamic refinement to enable analysis of large data sets on shared-memory systems. This direction requires progress on an in-memory data structure to support the forest construction, and on deriving heuristics and theoretical bounds for lower resolution and precision topological structures.

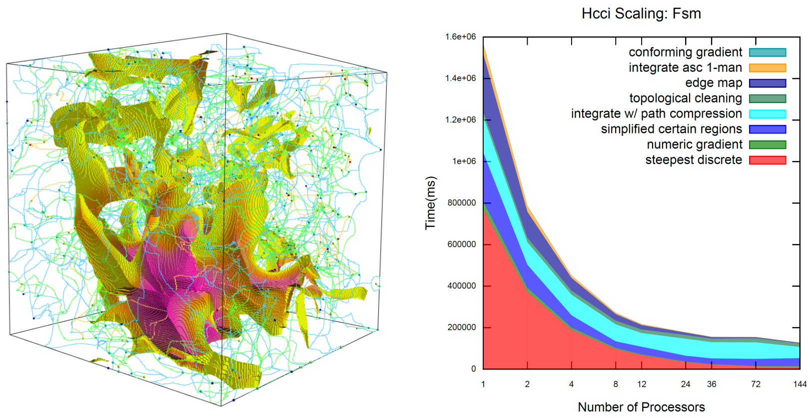

Figure 6: Leveraging a combination of numerical and discrete methods, the parallel computation of MS complex is possible avoiding the discretization artifacts common to topological approaches. This allows accurate flame surface computation to use later as a scaffolding for analysis.

We are also developing tools to use topological structures like merge trees and Morse-Smale complexes to extract high dimensional feature spaces for Machine Learning tasks (classification/ regression). The global topological structure (at full-scale) provides various attributes that summarize the data. For instance, number of leaves in the merge tree, number of branches, size of the branches, Betti number, number of regions in Morse-Smale complex segmentation, boundary surface area of ridges etc. are some attributes that represent the underlying structure of the data. We compute these attributes as a function of persistence i.e. we simplify the global topological structure by merging features whose lifetime (death-birth time) is below the threshold persistence and perform statistical analysis (like computing mean, variance, histograms, area under the curve, curvature etc.) to form feature spaces for Machine Learning. Our initial results show that training Machine Learning models using these feature spaces to be useful in both classification and regression tasks.

Multivariate Visualization System

We have designed a browser-based visualization system to support analysis of data such as that coming from the boiler simulation facility (BSF). The system is fully configurable allowing rearrangement of user interface components, sorting dimensions within plots, and multiple coordinated selections across different views of the data as well as a toolbar for supporting more complex operations such as aligning filters to the set ranges of a particular set of data. An example of the interface is shown in the Figure below. The system employs multiple annotated parallel coordinate plots in order to compare results stemming from physical experimentation, computer simulation, and regression modeling in order to identify and understand discrepancies occurring between various data sets and to validate when a reduced order modeling is capable of representing an experiment faithfully. This project is open source and available at https://github.com/maljovec/quetzal.

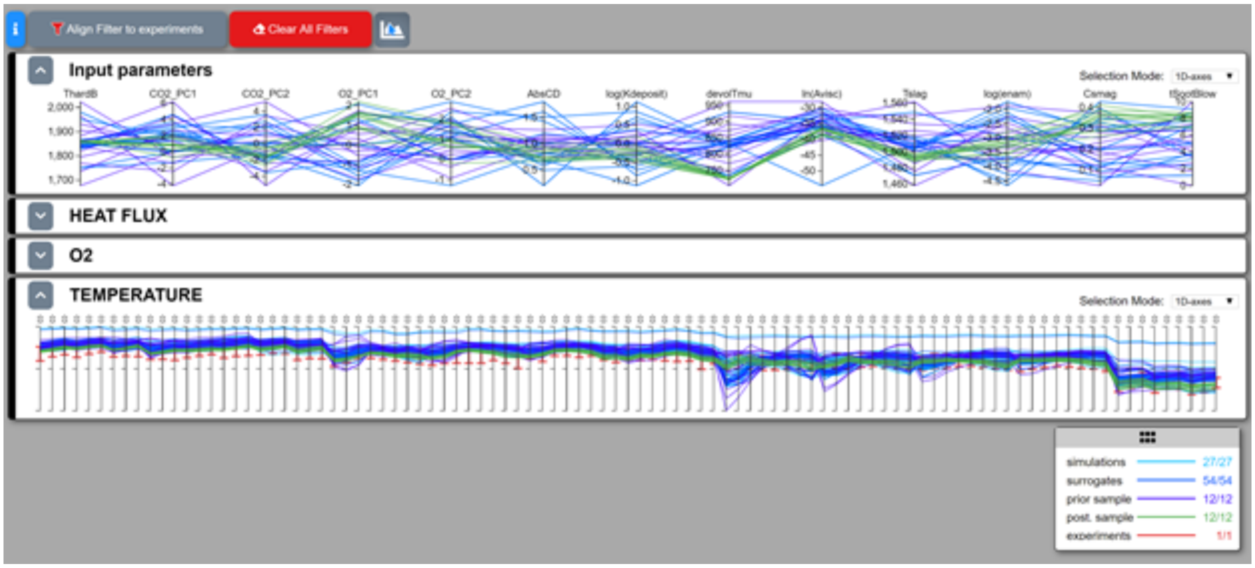

Figure 7: An example of the designed multivariate visualization system, quetzal, operating on BSF data.

An example analysis is provided in the Figure below. Here, we observe two separate parallel coordinate plots that are linked by the selection in the bottom plot. In the bottom plot, the O2 concentrations are observed at 41 different ports in a boiler tower. Each line represents a single simulation or model run and there is a corresponding line for every run in both the upper and lower plots. The different colors denote whether the run was a result of a simulation, a prior sample, a posteriori sample, or a result of a surrogate model. We have selected the experimental ranges of eight different observed O2 concentrations as shown by the collection of black boxes near the center of the lower plot. The result is that we have filtered all data that does not intersect this selection. With this information, we can then look at the top plot which shows the input parameter space that led to this set of filtered simulations to understand what ranges of input values correspond to this subset of data. Furthermore, summary information is provided in the legend showing how many runs of each dataset exist within the filter we have created. This visualization system has been useful for the engineering team to quickly and concisely convey information and to draw conclusions from their data.

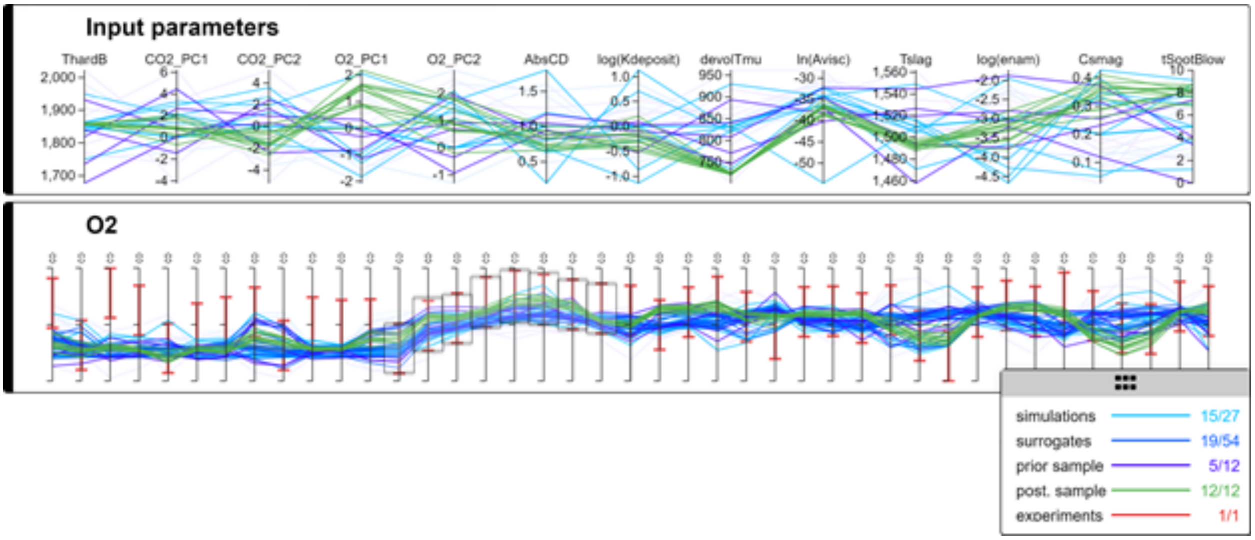

Figure 8: An Example Analysis

An analysis of O2 concentrations observed at various physical ports in the tower. Here we have filtered the data such that we only consider data points that lie within our experimental range for 8 of the 41 physical ports (each denoted by a vertical axis).

Advanced Visualization

We’ve been working on new scalable and flexible distributed rendering methods in OSPRay and using them to implement post-process and lightweight ***in situ*** visualization systems. Our Distributed FrameBuffer within OSPRay provides performance better than the state-of-the-art in common use cases and supports more flexible data distributions which are not otherwise possible. Moreover, we have implemented a data-distributed API for OSPRay, which enables its use within ***in situ*** visualization applications. With regard to ***in situ*** we've also been exploring more lightweight infrastructures, which can be easier to integrate and incur lower overhead, while still enabling the desired ***in situ*** visualization (e.g., by integrating SENSEI). We have under submission early work on developing faster spatially aware I/O strategies for particle data and are working to build acceleration structures within this format which are suitable for multi-resolution visualization and streaming. Furthermore, we have developed new geometric representations for complex or challenging data, such as isosurfaces in BS-AMR data and general tube primitives (e.g., streamlines). Along this path we are continuing work to support volume and isosurface rendering of other AMR formats, such as those used in Uintah and p4est. Finally, we have been investigating the application of the new GPU RTX hardware units to accelerate visualization and volume rendering tasks.

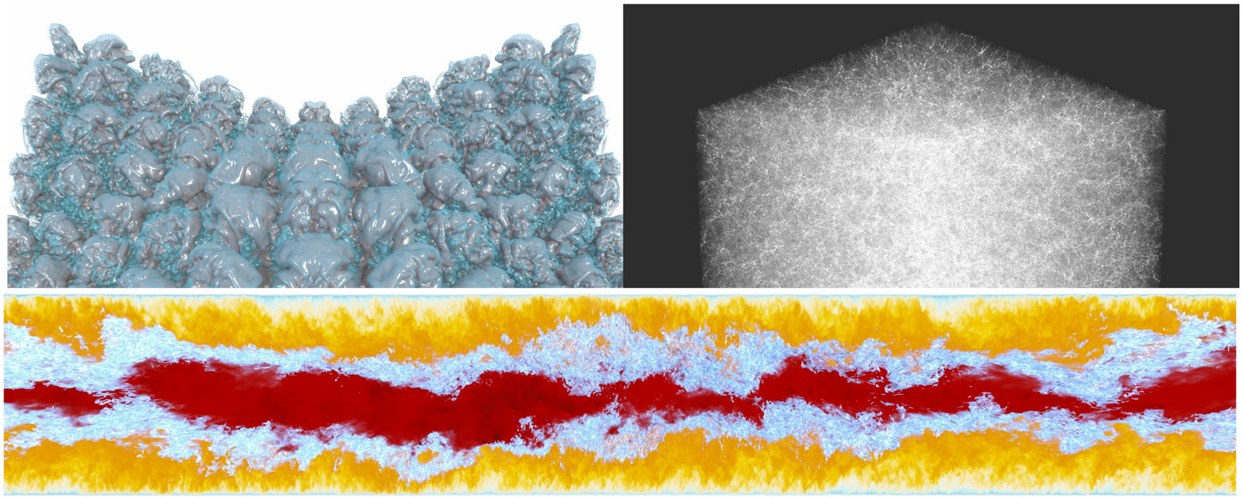

Figure 9: Large-scale interactive visualization using the Distributed FrameBuffer. Top left: Image-parallel rendering of two transparent isosurfaces from the Richtmyer-Meshkov (516M triangles), 8FPS with a 20482 framebuffer using 16 Stampede2 Intel Xeon Platinum 8160 SKX nodes. Top right: Data-parallel rendering of the Cosmic Web (29B transparent spheres), 2FPS at 20482 using 128 Theta Intel Xeon Phi Knight's Landing (KNL) nodes. Bottom: Data-parallel rendering of the 951GB DNS volume combined with a transparent isosurface (4.35B triangles), 5FPS at 4096x1024 using 64 Stampede2 Intel Xeon Phi KNL nodes.

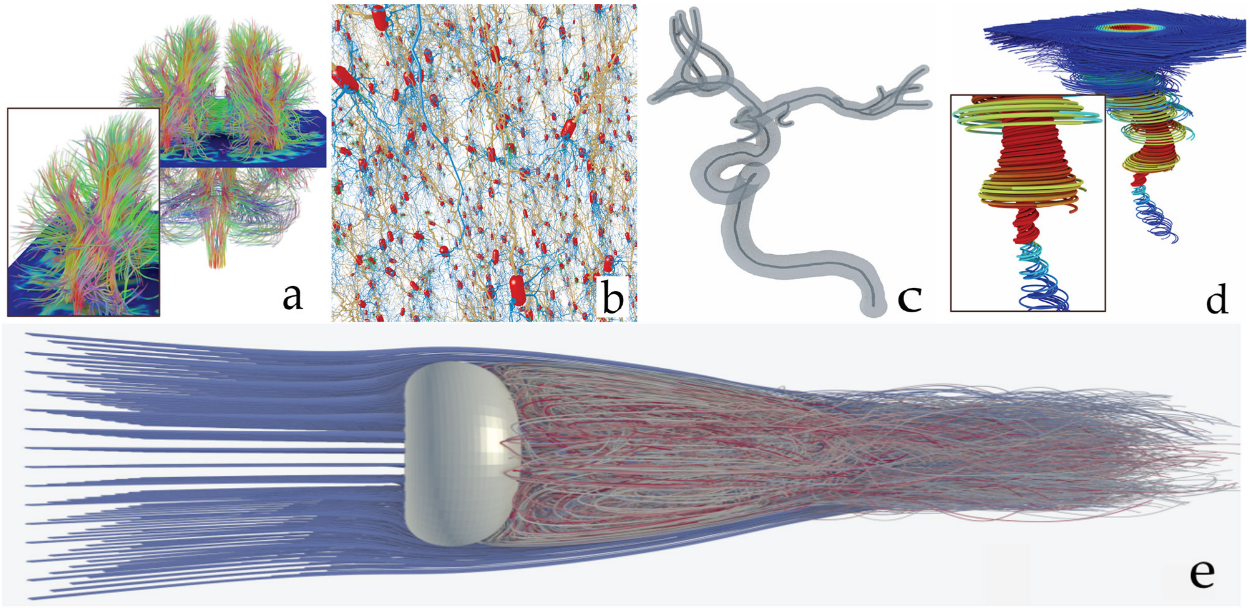

Figure 10: Visualizations using our "generalized tube" primitives. (a): DTI tractography data, semi-transparent fixed-radius streamlines (218K line segments). (b): A generated neuron assembly test case, streamlines with varying radii and bifurcations (3.2M l. s.). (c): Aneurysm morphology, semi-transparent streamlines with varying radii and bifurcations (3.9K l. s.) and an opaque center line with fixed radius and bifurcations (3.9K l. s.). (d): A tornado simulation, with radius used to encode the velocity magnitude (3.56M l. s.). (e): Flow past a torus, fixed-radius pathlines (6.5M l. s.). Rendered at: (a) 0.38FPS, (b)

7.2FPS, (c) 0.25FPS, (d) 18.8FPS, with a 2048x2048 framebuffer; (e) 23FPS with a 2048x786 framebuffer. Performance measured on a dual Intel Xeon E5-2640 v4 workstation, with shadows and ambient occlusion.

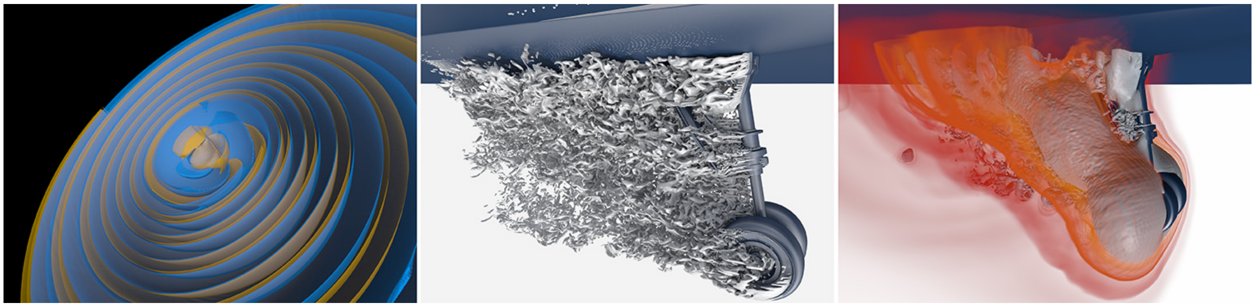

Figure 11: High-fidelity isosurface visualizations of gigascale block-structured adaptive mesh refinement (BS-AMR) data using our method. Left: a 28GB GR-Chombo simulation of gravitational waves resulting from the collision of two black holes. Middle and Right: a 57GB AMR dataset computed with LAVA at NASA, simulating multiple fields over the landing gear of an aircraft. Middle: isosurface representation of the vorticity, rendered with path tracing. Right: a combined visualization of volume rending and an isosurface of the pressure over the landing gear, rendered with OSPRay's SciVis renderer. Using our approach for ray tracing such AMR data, we can interactively render crack-free implicit isosurfaces in combination with direct volume rendering and advanced shading effects like transparency, ambient occlusion and path tracing.

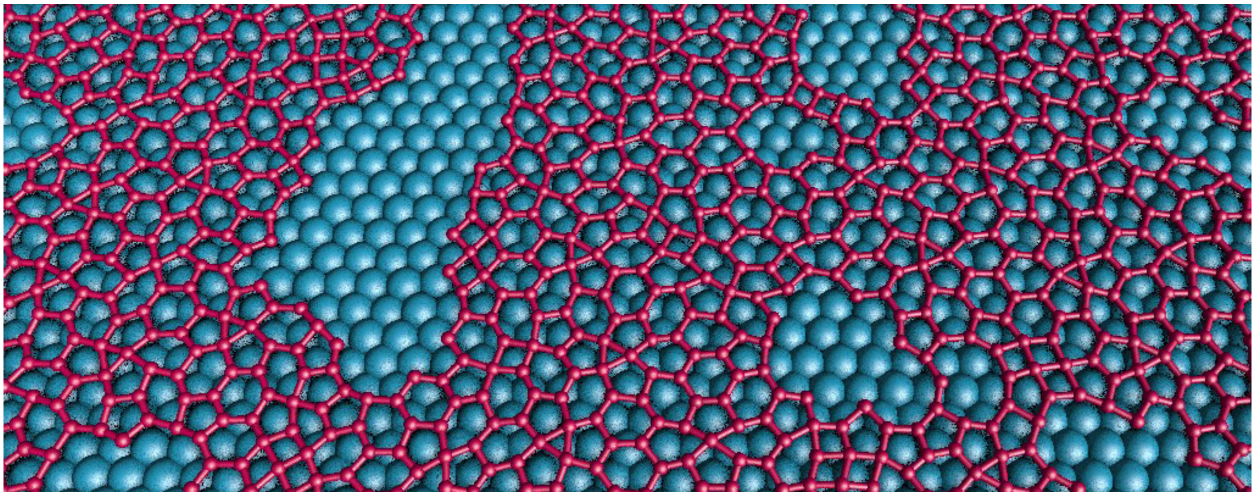

Figure 12: Interactive in situ visualization of a 172k atom simulation of silicene formation with 128 LAMMPS ranks sending to 16 OSPRay renderer ranks, all executed on Theta in the mpi-multi configuration. When taking four ambient occlusion samples per-pixel, our viewer averages 7FPS at 1024x1024. Simulation dataset is courtesy of Cherukara et al.

Data Models with Resolution vs Precision Trade-off

We explore the additional gains that could be achieved by combining reduction in precision and reduction in resolution. In particular, we developed a common framework that allows us to study the trade-off in both dimensions of data reduction in a principled manner. We represent data reduction schemes as progressive streams of bits and study how various bit orderings such as by resolution, by precision, etc., impact the resulting approximation error across a variety of data sets as well as analysis tasks. Furthermore, we compute streams that are optimized for different tasks to serve as lower bounds on the achievable error. Scientific data management systems can use the results presented in this paper as guidance on how to store and stream data to make efficient use of the limited storage and bandwidth in practice.

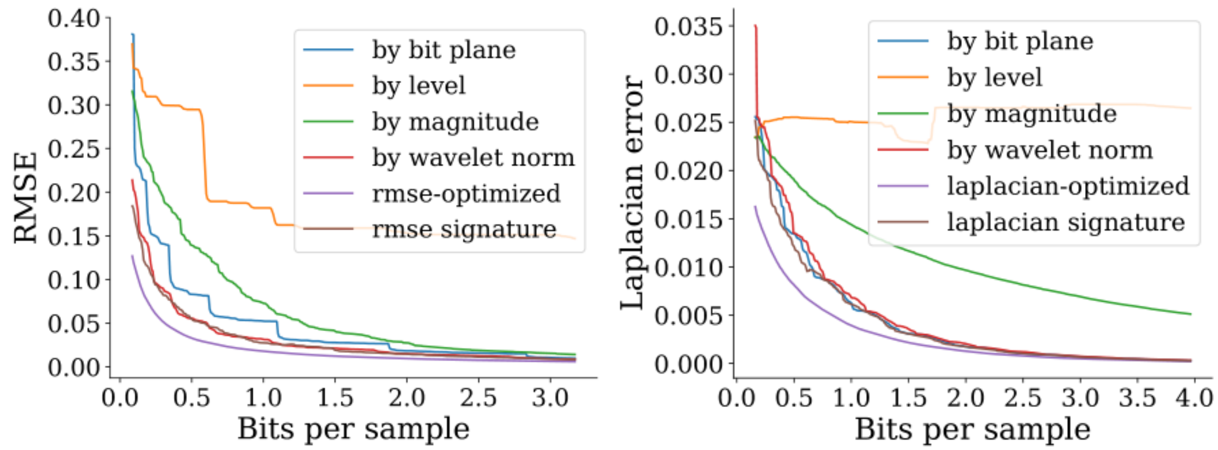

Figure 13: Left Figure—Root-mean-square error (RMSE) of reconstructed function (plasma) for different data streams. The "by wavelet norm" stream (combining both resolution and precision) performs better than both "by level" (resolution only) and "by bit plane" (precision only). Right Figure—Laplacian error induced by the same set of data streams. Here, "by wavelet norm" and "by bit plane" perform similarly, suggesting that derivative-type computations prefer resolution over precision.

File Format Extension Supporting Data Queries in Both Resolution and Precision

	Ours / State-of-the-art
alfredo	0.87
nanosphere	0.87
priya	0.89
viscontest	0.85
cosmo	1.17

Based on our IEEE VIS 2018 paper, we are developing a novel file format for scientific data that supports (region-of-interest) queries for data at low resolution and low precision, without reading redundant bits from disk. The format is based on discrete wavelet transform, compression of tiles of wavelet coefficients using zfp, and separation of every compressed tiles into equal-size chunks which are units of data to be written to and read from disk. The diagram below shows the current design for the format that we plan to deploy during the next year.

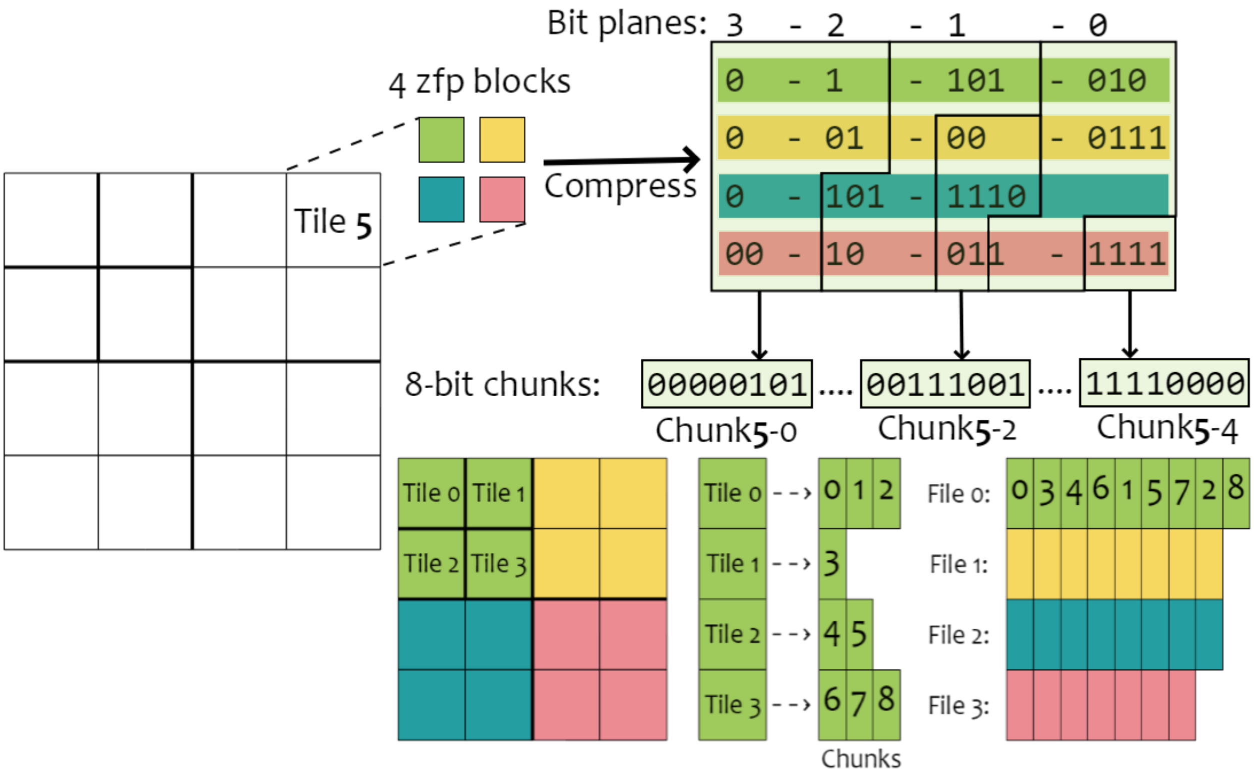

Figure 14: The current design for file format

Particle compression

We implemented binomial coding to compress particle data sets. The results are shown in the following table. In these experiments, we build a kd-tree on the particles where each splitting plane is at the exact middle in the spatial domain and each tree node stores the number of particles falling in that node. We are interested in compressing these numbers stored at the nodes only (that is, for now, we stop encoding information after reaching one particle per node).

Table 1: The ratio of data compressed with our method (based on binomial distribution) over data compressed using state-of-the-art method (based on uniform distribution). In most cases we can achieve a further reduction of about 14% in data size, except for cosmology data for which the distribution of particles is highly non-uniform.

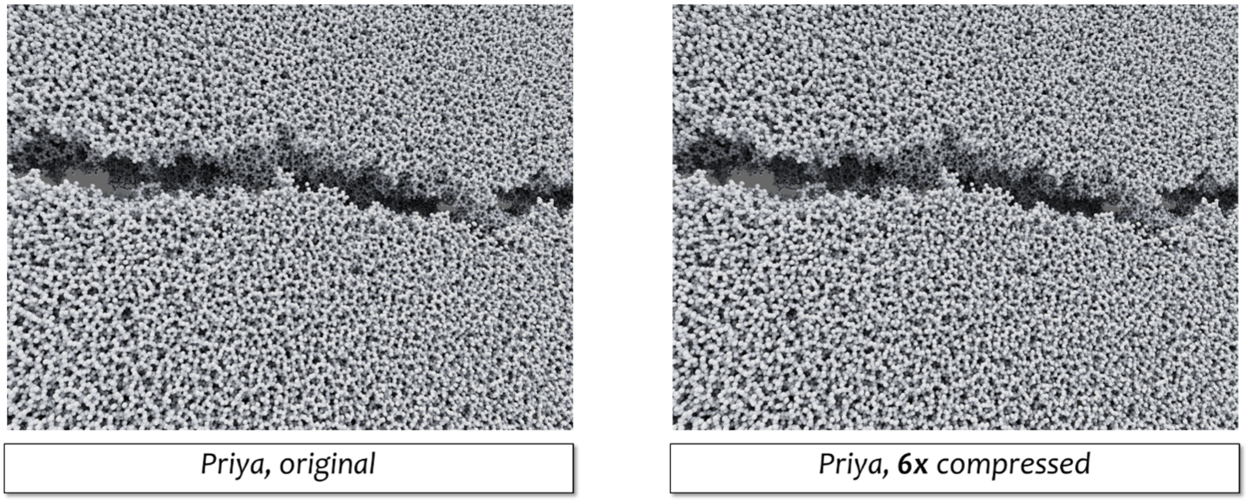

Priya, original

Priya, **6x** compressed

Figure 15: Visualization of the original particle data (Priya, 4.7 million particles) and its 6x compressed version using our method. The two renderings are visually identical.

Physics Decoupling from Runtime

During the year there was a major effort to complete the restructuring and decoupling of the Uintah runtime infrastructure from the physics applications. Thus, making the Uintah runtime infrastructure fully independent of the applications which improved long term maintainability and portability. One highlight was to formally add runtime triggers, which are events that change or augment the normal code execution. Among other abilities, the triggers allowed for dynamic solve frequencies rather than the previous static frequency. The dynamic solve frequency was used to improve the radiation solve calculations in Arches.

VisIt and In Situ Visualization

Another major effort was to continue to harden and expand the ***in situ*** performance data analysis and visualization. Performance stats can be collected, visualized, and analyzed across both the runtime infrastructure and the physics applications. This work was presented and published in the Workshop on *In Situ* Visualization at ISC (June 2018). New additions were made to collect thread performance data as well as patch-based performance data on both the simulation and machine layout. The ability to use multiple layouts has been a valuable tool in understanding different performance phenomena such as load balancing issue. In one case, it was discovered that ranks went unused. Now work has commenced to collect communication performance data.

We completed the work on scalable volume rendering research and development in May 2018 and presented the peer-reviewed, published research results at the Eurographics Parallel Graphics and Visualization (EGPGV) Symposium and at the annual review in Utah. As noted in last year's annual report, large-scale simulations can produce data easily in excess of what can be efficiently visualized using production visualization software, making it challenging for scientists and engineers to gain insight from the results of these simulations. This scalability trend will worsen with exascale, requiring modification to visualization packages to handle visualization at these large scales and new hardware architectures. We developed an effective integration of OSPRay into VisIt to provide a scalable rendering backend. This is coupled with CCMSC's work on PIDX for scalable I/O. Now we can render geometry defined by VisIt along with the scalable volume rendering.

Our implementation achieves up to 30 times higher framerates than VisIt's RayCasting:Compositing renderer, while producing equivalent high-quality images. Furthermore, we have achieved significant strong and weak scaling improvements up to 32,768 CPU cores on recent HPC platforms equipped with Intel® Xeon Phi™ (KNL) processors and Intel® Xeon® Platinum Skylake (SKX) processors. Specifically, our engineering contributions to VisIt's distributed rendering capabilities are:

- Moving to a hybrid-parallel execution model,
- Leveraging OSPRay for fast volume rendering, and
- Improving the use of IceT for image compositing.

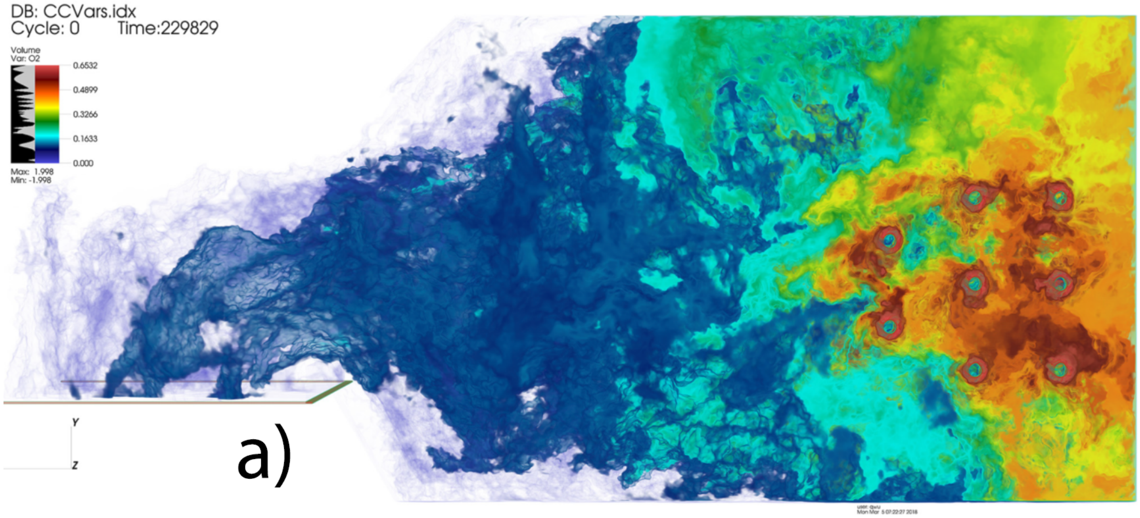

Figure 16: High-quality interactive volume visualization using VisIt-OSPRay: a) volume rendering of O2 concentration inside a combustion chamber

This research was presented and published at Eurographics Parallel Graphics and Visualization (EGPGV) Symposium, June 2018.

PHYSICS

Model Development: Emissivity Model, Production Code

Ash transformations on the walls affect heat transfer and combustion efficiency. Emissivity and thermal conductivity are primarily influenced by these phenomena, which in turn is driven by the composition of the depositing particles, and the type of physical structures they form on the hot surfaces. Past efforts have updated the way we model emissivity on the walls and last year, we improved yet again the methodology.

In the production code, we have implemented a way to compute spectral emissivities based on the scattering characteristics of coal ashes. These characteristics are determined from the optical properties of ash particles for which we have implemented models based on the works of D. Goodwin[1] and J. Ebert[2]. The algorithm uses different correlations at specific wavelengths to compute the complex index of refractions based on a large data-set of measurements. The index of refraction would be used to compute absorption, scattering and extinction properties for the particles via Mie Theory calculations[3] (Bohren & Huffman, 2008 for single scatterers. In general, Discrete Ordinates (DO) is the more correct way to compute particle emittance, but approximations were made due to the expensiveness of these computations[4]. The spectral and total emissivity are given by:

$$1 - \epsilon_\lambda = \frac{(\sqrt{1-\omega g} - \sqrt{1-\omega})}{(\sqrt{1-\omega g} + \sqrt{1-\omega})}$$

$$\omega = \frac{Q_{sca}}{Q_e xt}$$

$$\epsilon(T_s) = \frac{1}{E_B(T_s)} \int_{\lambda=0}^{\infty} \epsilon_\lambda(T_s) e_B(\lambda, T_s) \mathrm{d}\lambda$$

Where (ω) is the albedo of scatter, (g) is the asymmetry parameter, (Qi) are the scattering and extinction efficiencies. In the limit where the surfaces are smooth and reflecting (molten slags, glassy slags) a different expression is used to calculate the spectral emissivity (also known as Fresnel emissivity); the expression is given by:

$$\epsilon_\lambda = 1 - \frac{(n-1)^2 + k^2}{(n+1)^2 + k^2}$$

---

[1] *Goodwin, D.G., 1986. Infrared Optical Constants of Coal Slags. s.l.: Stanford University.*

[2] *Ebert, J.L., 1994. Infrared Optical Properties of Coal Slag at High Temperatures. s.l.:Stanford University.*

[3] *Bohren, C.F. and Huffman, D.R., 2008. Absorption and Scattering of Light by Small Particles. s.l.:Wiley Science Series; John Wiley & Sons.*

[4] *Bohren, C.F., 1987. Multiple scattering of light and some of its observable consequences. Am. J. Phys., Issue 55, p. 524 – 533.*

Where (n) and (k) are the components of the complex index of refraction and the total emissivity is calculated by integrating over all wavelengths.

It has been shown by T. Wall[5] that the total or hemispherical emissivity decreases with the temperature up to the point where particle agglomeration influences the scattering behavior of the formed structure, and increase the emissivity at higher temperatures. The model that has been selected to represent these sintering effects is the one proposed by O. Pokunda[6] (Pokluda, et al., 1997) which was formulated for viscous sintering and calculates the ratio: length of the sintering neck to the particle radius between two particles.

$$\frac{\mathrm{d}\theta}{\mathrm{d}t} = \frac{\Gamma}{r_0\mu}\Big[\frac{2^{-5/3}\cos(\theta)\sin(\theta)(2-\cos(\theta))^{1/3}}{(1-\cos(\theta))(1+\cos(\theta))^{1/3}}\Big]$$
$$\frac{x}{r} = \sin(\theta)$$
$$\frac{x}{r_0} = \sin(\theta)\Big(\frac{4}{(1+\cos(\theta))^2(2-\cos(\theta))}\Big)^{1/3}$$

The original model describes agglomeration of two particles and efforts in the literature have added the capability to describe systems of many particles by considering the effect of fractal aggregates7. In this current iteration of the sintering model, we have developed a “staging” mechanism that allows particles to agglomerate sequentially, growing in size until only a few particles or one big particle remains. The process can be described more specifically in what follows. Assume that N number of particles arrive to the wall (a given computational cell) and the agglomeration takes place for pairs of particles simultaneously, once the sintering process between pairs finishes, the number of particles has reduced to half, Np/2. The newly sintered and grown particles would sinter again reducing the number of particles to Np/4, this process would go on until there’s no particles left to sinter. Each time the current number of particles reduces in half corresponds to a “stage” and the theoretical, total number of stages can be computed as: st = log(Np)/ log(2). This process is depicted i18The concept of staging allows us to eliminate the coordination number as an uncertainty parameter and rely on the number of particles arriving to the walls, this quantity is better characterized in the simulations. The equivalent diameter of the particles after sintering at each stage is calculated based on the total surface area and the total volume of all particles at a given computational cell:

---

[5] *Wall, T.F. et al., 1993. The Properties and Thermal Effects of Ash Deposits in Coal-fired Furnaces. Prog. Energy Combust. Sci., Vol. 19, pp. 487 – 504.*

[6] *Poklunda, O., Bellehumeur, C.T. and Machopoulos, J., 1997. Modification of Frenkel’s Model for Sintering. AICHE J., Vol. 43, pp. 3253-3256.*

[7] *Eggersdorfer, M.L., Kadau, D., Herrmann, H.J., and Pratsinis, S.E., 2011, Multiparticle Sintering Dynamics: From Fractal-like Aggregates to Compact Structures. Langmuir, Vol. 27, pp. 6358-6367.*

$$
\begin{aligned}
D_0 &= D_i\Big(\frac{4}{(1+\cos\theta)^2(2-\cos\theta)}\Big)^{1/3}\\
S_0 &= \pi D_0^2(1+\cos\theta)\\
S_t &= S_0 N_p 2^{-st}\\
V_t &= \frac{\pi N_p D_0^3}{6}\\
D_{eff} &= \frac{6V_t}{S_t}
\end{aligned}
$$

where variables with the sub-index "0" denote fields (particle diameter and surface area) for 2 particles. Variables with the sub-index "t" denote fields for the total number of particles in the computational cell at a particular stage "s_t" (surface area and volume).

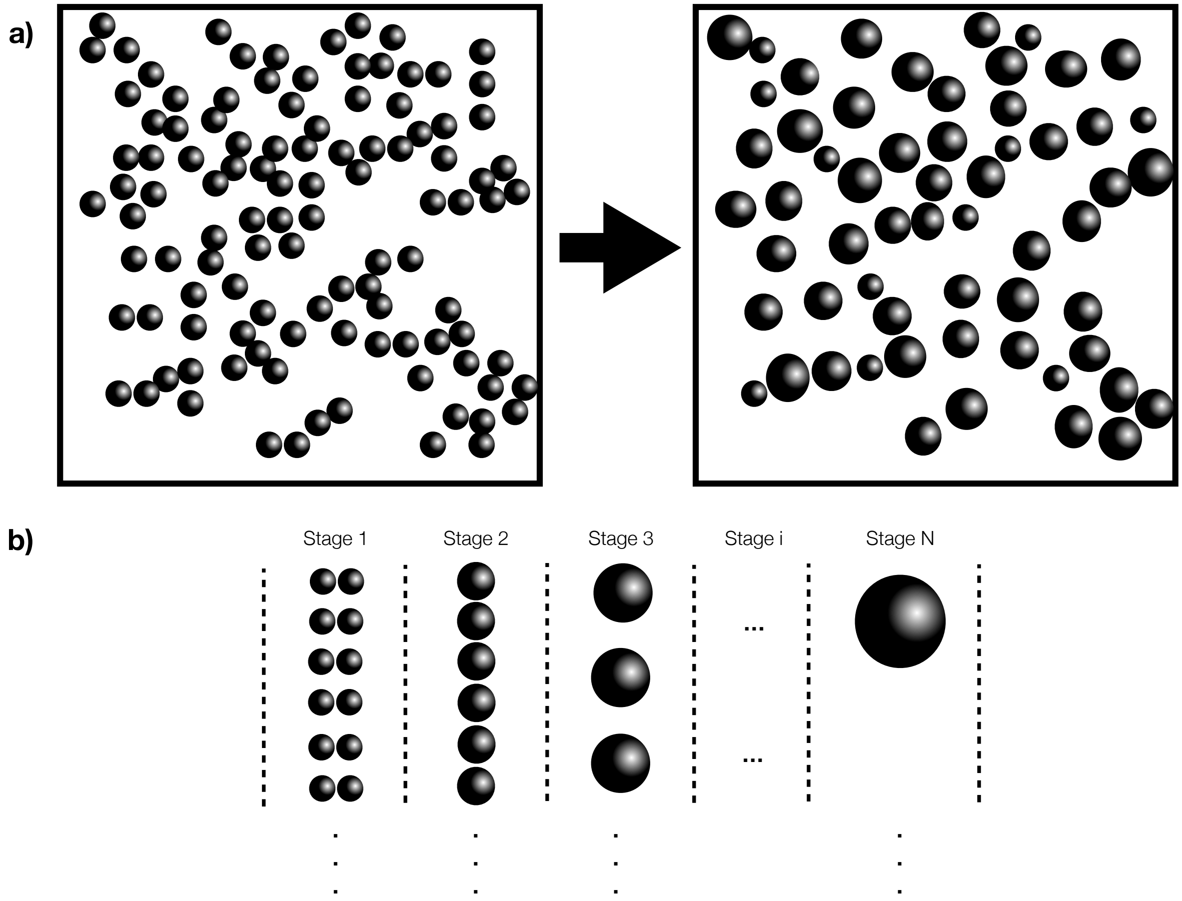

Figure 17: Sintering process represented as agglomeration stages. a) State of the particles arriving to the wall and after N stages of sintering. b) Conceptual representation of the staging process for the sintering of particles in a).

D_{eff} is the equivalent diameter of the structure at a particular stage s_t. We have applied this model to the sintering of the synthetic slags, prepared by J. Boow and P.R.C. Goard[8]. Their work outlines the most relevant behavior of slags emissivity in presence of particle sintering. A comparison of this experimental data set against our model is presented in Figure 18. For small particles, the transition of emissivity where sintering takes place (1200 K - 1400 K) is relatively smooth, this is due to the greater number of sintering events (stages) required to reach a critical size (where predictions of Mie theory are not as accurate). The transition for bigger particles is less smooth, reaching critical sizes faster than small particles. In such cases, the total emissivity transition quickly to the emissivity of glasses, polished surfaces and melts, given by the Fresnel relationship presented previously. This is certainly an improvement over the previous model where a sharp transition was set between the emissivity at the end of the sintering stage and the emissivity of glasses, bright and smooth surfaces of the deposits.

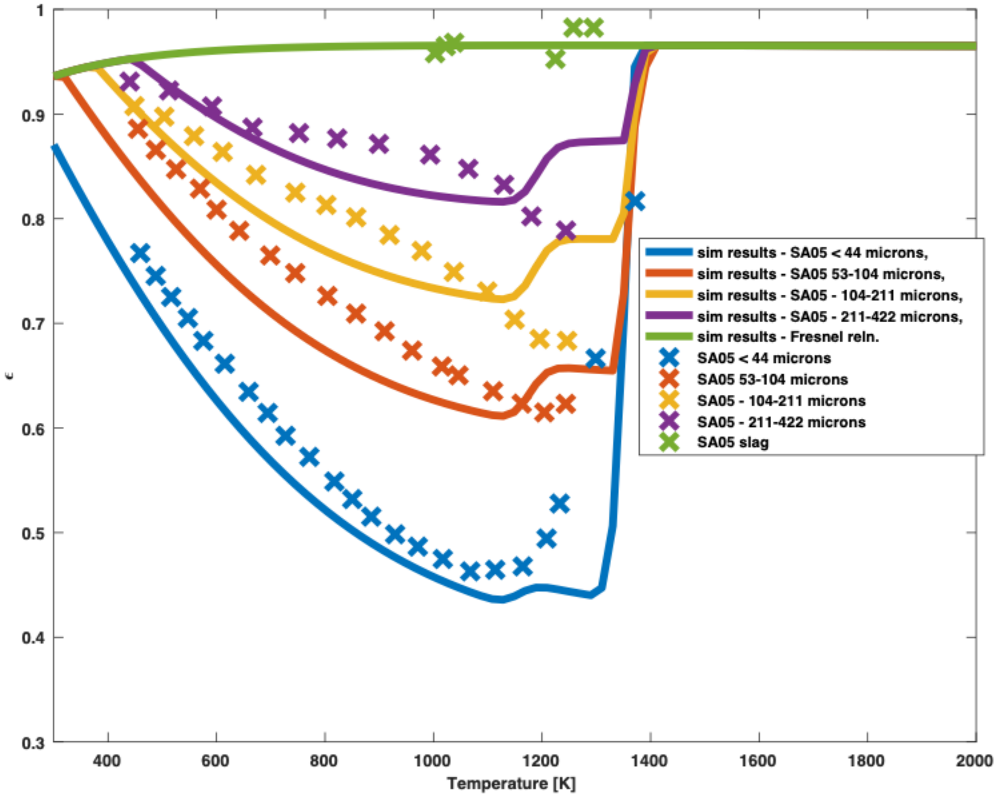

Figure 18: Comparison the total emissivity of a synthetic slag SA05 (5% iron oxide content) with the predictions of the proposed model

---

[8] *Boow,J. and Goard, P.R.C., 1969, Fireside deposits and their effect on heat transfer in a pulverized-fuel-fired boiler: Part III. The influence of the physical characteristics of the deposit on its radiant emittance and effective thermal conductance. J. Inst. Fuel, Vol. 42, pp. 412-419*

Coal Devolatilization

One of the major assumptions currently implemented in Arches is that the char and volatile off-gases have the same composition after devolatilization. This is not true, since the volatile gases become enriched with hydrogen, and the char becomes enriched with carbon. An extensive literature search was carried out to find relevant and useful data on the elemental compositions (primary organic components, CHONS) of char and tar at a variety of experimental compositions. These data were used to develop a set of correlations describing the elemental compositions of these devolatilization products as a function of several coal specific parameters and experimental conditions. The table below shows the sources used in the elemental composition correlation analysis.

Table 2: Data sources used for elemental composition correlations

Author(s)	**Institution**	**Apparatus**	**Gas Temperatures (K)**	**Char/Tar**	**Coal Types**
Freihaut, et al.	United Technologies Research Center	Entrained flow reactor	780-1069	Tar	hvA bit, sub C, LVB
Hambly	Brigham Young University	Drop tube reactor	820 and 1080	Char and Tar	ligA, subA, hvCb, hvAb, lvb
Perry	Brigham Young University	Drop tube reactor	895-1085	Char and Tar	Sub, hvb, mvb, lvb
Fletcher and Hardesty	Sandia National Laboratories	Entrained flow reactor	1050	Char	Lig, sub, hvBb, hvAb, lvb
Watt	Brigham Young University	Drop tube reactor	850-1050	Char and Tar	ligA, subC, hvCb, hvAb, lvb
Parkash	Devon Coal Research Centre	Laminar, entrained flow, atmospheric pressure reactor	820-980	Char	Sub B
Tyler	CSIRO	Heated fluidized bed reactor	873	Tar	Bit, sub

Using these data, figures were constructed to show relationships between the independent variables (such as maximum gas temperature, particle residence time, parent coal compositions, etc.) and the dependent variables (normalized elemental compositions, i.e. Cchar/Ccoal). Figures 19 and 20 show examples of these plots.

One of the coal specific parameters used in the elemental correlation analysis is parent coal aromaticity. Several correlations from the literature to describe coal aromaticity were tested using a dataset found from literature. Ko, et al.[9] [10]developed a second order polynomial for a limited set of aromaticity data based on the carbon content of the parent coal. Gerstein, et al.[11] used a simpler linear correlation of f_a', also based on the carbon content of the parent coal. Carr and Williamson[12] developed a different correlation to predict the apparent aromaticity based on vitrinite reflectance. Maroto-Valer, et al.[13] proposed an updated version of a previous correlation based on the hydrogen-carbon ratio, H/C. Singh and Kakati studied four different expressions to predict aromaticity.[14]

The aromaticity models were evaluated on a set of data that includes elemental composition (proximate and ultimate analysis results) and measured and calculated chemical structure parameters from NMR analysis. Some of these data were used by Genetti et al.[15] to fit chemical structure parameters for the CPD model. The data come from the following sources: Genetti, et al.[22], Solum, et al.,[16] Hambly, et al.,[7] [8] Perry, et al.,[9] [10] Fletcher and Hardesty,[11] Watt, et al.,[12] [13]

---

[9] *Ko, G.H., Peters, W.A., Howard, J.B., Correlation of Tar Yields with Rapid Pyrolysis with Coal Type and Pressure. Fuel, 1987 66(8) pp.1118-1122.*

[10] *Ko, G.H., Sanchez, D.M., Peters, W.A., Howard, J.B., Correlations for Effects of Coal Type and Pressure on Tar Yields from Rapid Devolatilization. Symposium (International) on Combustion, 1989, 22(1) pp.115-124.*

[11] *Gerstein, B.C., Murphy, P.D., Ryan, L.M., Aromaticity in Coal, Coal Structure, Meyers, R.A., Ed. Academic Press: New York, 1982.*

[12] *Carr, A.D. and Williamson, J.E., The Relationship between Aromaticity, Vitrinite Reflectance and Maceral Composition of Coals: Implications for the Use of Vitrinite Reflectance as a Maturation Parameter. Org. Geochem. 1990, 16(1-3) pp. 313-323.*

[13] *Maroto-Valer, M.M., Andresen, J.M., Snape, C.E., Verification of the Linear Relationship between Carbon Aromaticities and H/C Ratios for Bituminous Coals. Fuel, 1998, 77(7) pp.783-785.*

[14] *Singh, K.P. and Kakati, M.C., Comprehensive Models for Predicting Aromaticity of Coals. Chem. Eng. Commun., 2003, 190(10) pp. 1335-1347.*

[15] *Genetti, D., Fletcher, T.H., and Pugmire, R.J., Development and Application of a Correlation of C-13 NMR Chemical Structural Analyses of Coal Based on Elemental Composition and Volatile Matter Content. Energy Fuels, 1999, 13(1) pp. 60-68.*

[16] *Solum,M.S., Pugmire, R.J., and Grant, D.M., C-13 Solid-State NMR of Argonne Premium Coals. Energy Fuels, 1989, 3(2): pp.87-193.*

Gerschel and Schmidt,[17] Cui, et al.,[18] Ahmed, et al,[19] Lin, et al.,[20] Suggate and Dickinson,[21] and Zhang, et al.[22]

The seven literature correlations were tested against a comprehensive set of data in two ways: first by using the provided coefficients, and second by re-fitting the coefficients using the same dataset. Predictions made using the re-fit coefficients performed better than the originals, and are shown along with a new proposed model in Figure 21.

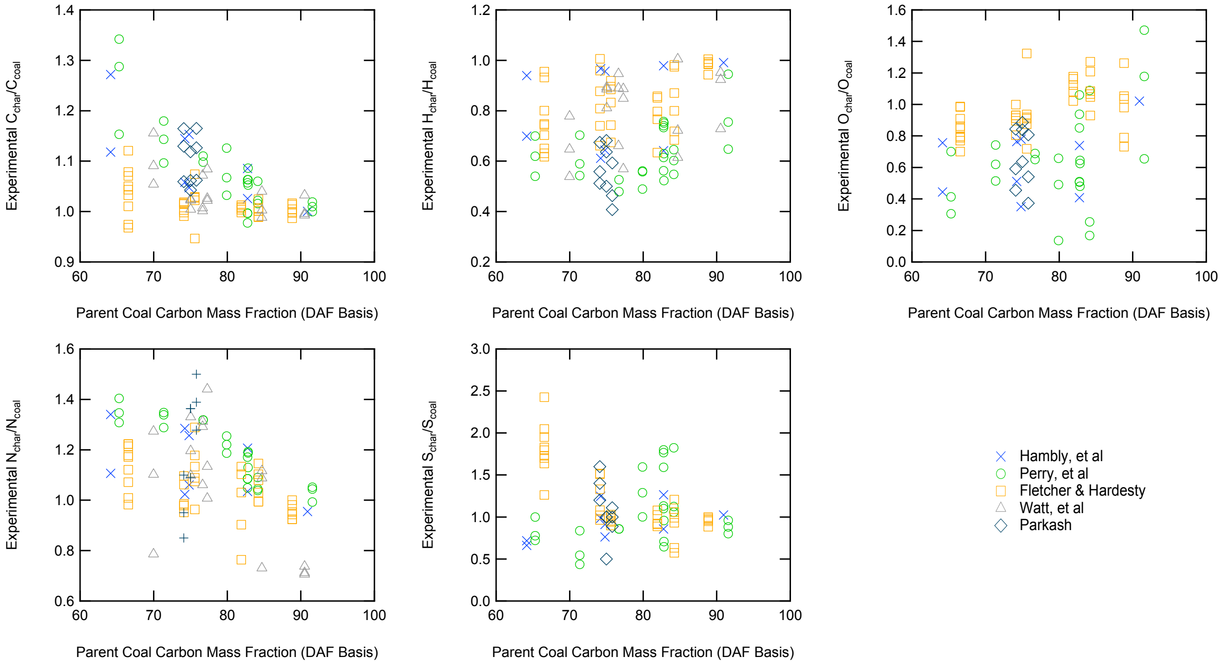

Figure 19: Experimentally determined normalized elemental mass fractions in the char compared to the parent coal mass fraction, from each reference used.

---

[17] *Gerschel, H., Schmidt, M., Modelling the Relationship between Carbon Aromaticity of Lignite Pyrolysis Chars and the Process Temperature with Petrographic Parameters. Int. J. Oil Gas Coal Technol. 2016, 11(3) pp.290-307.*

[18] *Cui, X., Yan, H., Zhao, P., Yang, Y., Xie Y., Modelling of Molecular and Properties of Anthracite Base on Structural Accuracy Identification Methods. Journal of Molecular Structure. 2019, 1183 pp. 313-323.*

[19] *Ahmed, M.A., Blesa, M.J., Juan, R., Vandenberghe, R.E., Characterization of an Egyptian Coal by Mossbauer and Ft-Ir Spectroscopy. Fuel 2003, 82(14) pp.1825-1829.*

[20] *Lin, H.L., Li, K.J., Zhang, X.W., Wang, H.X., Structure Characterization and Model Construction of Indonesian Brown Coal. Energy Fuels, 2016, 30(5) pp. 3809-3814.*

[21] *Suggate, R.P., Dickinson, W.W., Carbon NMR of Coals: The Effects of Coal Type and Rank. Int. J. Coal Geol.2004, 57(1) pp. 1-22.*

[22] *Zhang, P.Z., Li, L.Y., Ye, C.H., Solid State C-13 NMR Study of Chinese Coals. Fuel Sci. Technol. Int., 1995, 13(4) pp. 467-481.*

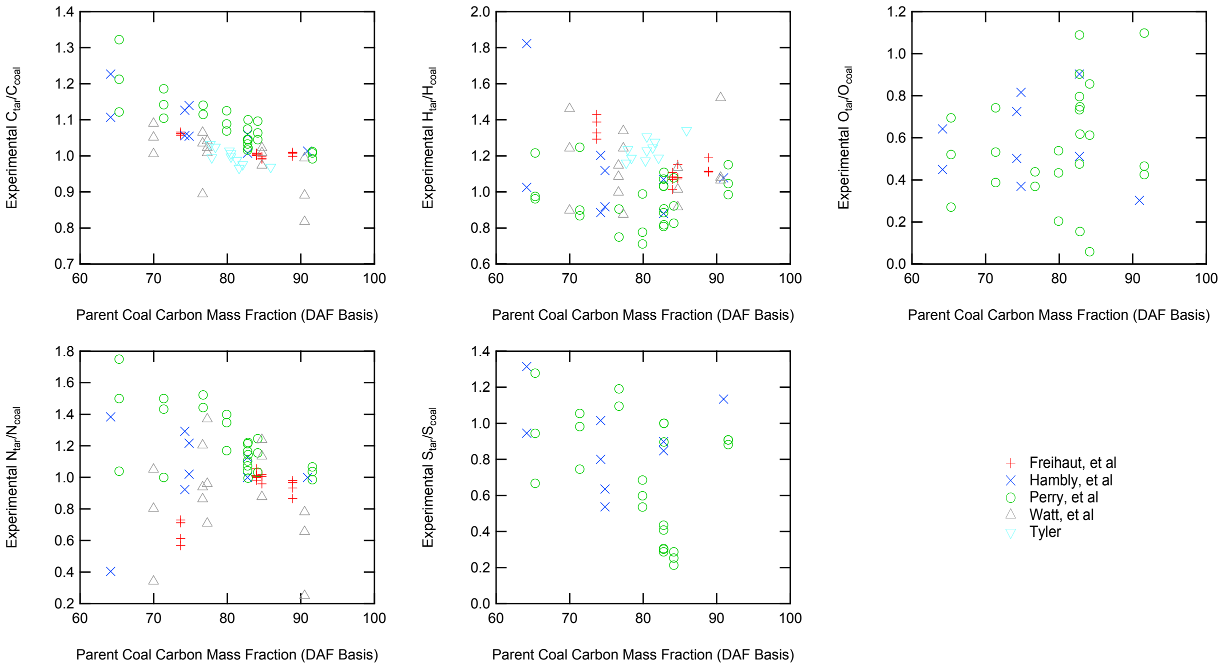

Figure 20: Experimentally determined normalized elemental mass fractions in the tar compared to the parent coal carbon mass fraction, from each reference used.

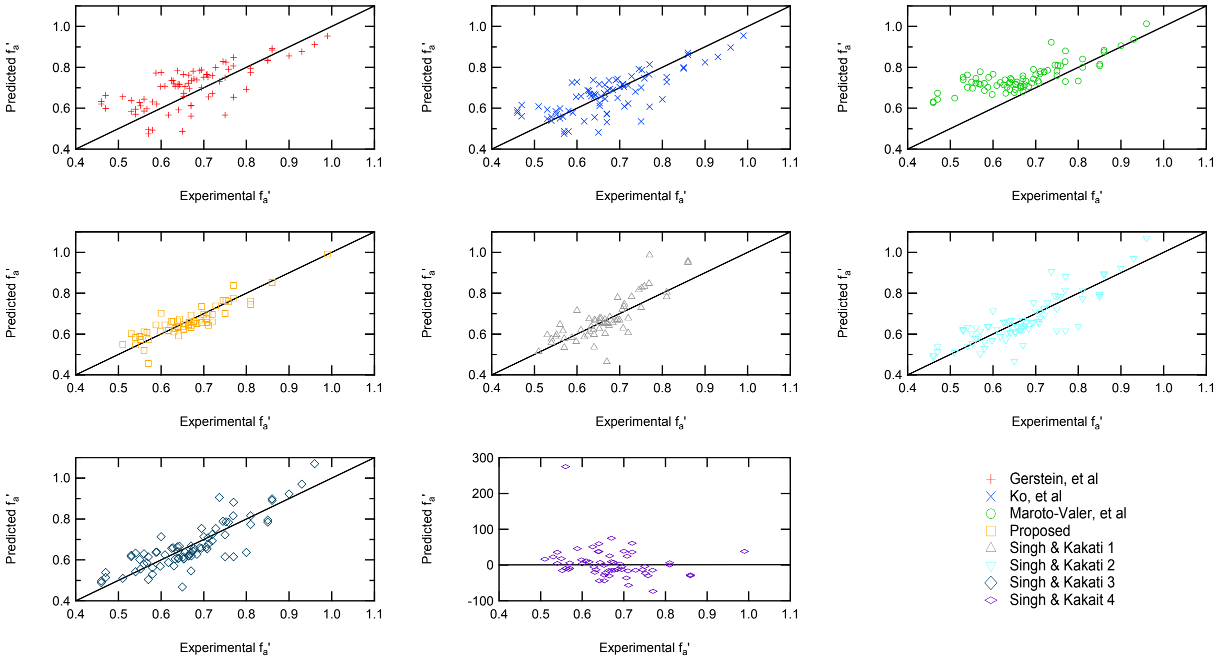

Figure 21: Parity relationship of seven aromaticity correlations from literature and a newly proposed aromaticity correlation (orange squares).

The best of the literature models was Singh and Kakati's second and third correlations (light blue triangles in the second row and third column of Figure 21 and dark blue diamonds in the bottom row and first column, respectively). The proposed model did even better.

$$f_a' = 4.384 - 8.679 \cdot 10^{-2} \cdot X_C + 5.352 \cdot 10^{-4} \cdot X_C^2 + 2.601 \cdot 10^{-2} \cdot X_H - 6.879 \cdot 10^{-3} \cdot X_H^2 + 3.525 \cdot 10^{-3} \cdot X_O - 5.710 \cdot 10^{-4} \cdot X_O^2 - 2.666 \cdot 10^{-3} \cdot X_{VM} + 5.659 \cdot 10^{-6} \cdot X_{VM}^2$$

Using the proposed aromaticity correlation and the experimental data from the sources in the table, an extensive uncertainty quantification analysis is currently being pursued to develop comprehensive correlations describing the major elemental compositions (CHONS) in the char and the tar. There are several different approaches to quantify uncertainty, including Bayesian statistics, cross validation, relative uncertainty using L_1, L_2, and infinity norms, and other similar approaches. Each approach has unique advantages and disadvantages. The goal of the statistical analysis is to find not only an overall uncertainty, but to find the uncertainty at each experimental condition. A journal article is in progress containing the details of the aromaticity correlations and the elemental composition correlations.

Soot in Coal Flames

Three approaches to modeling tar and soot formation in coal combustion systems were evaluated. The first approach (case ***A***) used the empirical Brown-Fletcher tar and soot model and assumed soot is carbon and tar is $C_{10}H_{10}$. Case ***A*** is used as the standard of comparison as soot is known to be depleted in hydrogen compared to tar. The second approach (case ***B***) also used the Brown-Fletcher tar and soot model, but assumed soot and tar are both $C_{10}H_{10}$. The motivation for assuming tar and soot are both $C_{10}H_{10}$ is that the product composition of tar and soot reactions can be parameterized by a single stream, which is advantageous when considering a mixture fraction-based chemistry mode, while adequate parameterization of tar and soot reaction products for case ***A*** requires three streams. The third modeling approach did not utilize a model for soot formation and used C_2H_2 as a "tar surrogate" instead.

Comparison of the three approaches was carried out by simulation of a turbulent coal combustion with the One-Dimensional Turbulence code. High fidelity models were used for chemical kinetics and devolatilization, while char combustion was neglected. We found that peak temperatures calculated for cases ***B*** and ***C*** were similar to case ***A***, and that the predicted behavior of gas phase quantities with respect to equivalence ratio for case ***A*** was shifted towards an equivalence of zero for lean conditions. Predicted particle temperature and coal volatiles remaining for case ***B*** were nearly identical to case ***A*** values, while case ***C*** predicted complete depletion of coal at an earlier residence time than cases ***A*** and ***B***. Overall, we found that assuming that soot has the same stoichiometry as tar was reasonable. Additionally, ***a priori*** flamelet reconstructions of various gas phase quantities were compared to finite-rate calculations for cases ***A, B,*** and ***C***. Flamelet reconstructions of O_2 mass fraction and gas temperature proved to be very accurate for cases ***B*** and ***C,*** while reconstructions for case ***A*** exhibited decreased accuracy in fuel lean conditions for O_2 and fuel rich conditions for gas temperature.

Our efforts in soot formation modeling have focused on the development, implementation, and testing of a physics-based soot model for coal combustion in oxy-coal environments. Previous models, such as the Brown and Fletcher model, while robust and easy to implement, are largely empirical, making it difficult to apply the models to oxy-coal conditions with confidence in the accuracy and performance of the model.

We have developed detailed and reduced soot models. The detailed model considers soot formation from tar species that are evolved during coal devolatilization. The tar size distribution is represented using a sectional model with nine sections. The soot particle size distribution is represented with the method of moments with interpolative closure (MOMIC).[23] Six sections were used. In addition, an additional "d" moment was included to model particle aggregation.[24] Soot is formed from tar combination via coagulation; soot growth proceeds via collision-based tar deposition and the HACA[25] mechanism. Soot coagulation and aggregation is included, along with oxidation and gasification as have been previously reported. Tar cracking is an important submodel that was specially developed here by considering tar to be composed of four component groups: phenol, toluene, naphthalene, and benzene. The soot model was validated against coal experiments previously performed at BYU. The model is also applicable to biomass combustion, and the model also was validated using a data set for biomass combustion using three different wood species. The detailed model was published in ***Combustion and Flame*** during this reporting period.

This detailed model is fairly complex and computationally expensive. A significant effort has been underway to simplify the model. A monodisperse version was developed that utilized the same physical and reaction processes as the detailed model, but only requires three species: tar, soot number density, and mass fraction. This model has been implemented in Arches and is undergoing testing. We have submitted a paper on this model to ***Combustion Theory and Modeling*** that is currently in review. The model was applied in Arches to biomass combustion successfully. We have since been testing the model under oxy-coal conditions. We have found and corrected a few errors in the monodispersed formulation, and we are working to accommodate stiffness due to fast tar coagulation/soot nucleation processes that have caused some instabilities.

In addition, simulations in the OFC are being performed to quantify the influence of the soot on radiation heat transfer, as well as sensitivity to the specific soot models used and the coupling

---

[23] *Frenklach,M., Method of Moments with Interpolative Closure. Chem. Eng. Sci. 2002, 57(12) pp. 2229-2239.*

[24] *Balthasar, M. and Frenklach, M., Detailed Kinetic Modeling of Soot Aggregate Formation in Laminar Premixed Flames. Combust. Flame, 2005, 140(1-2) pp. 130-145.*

[25] *Appel, J., Bockhorn, H., and Frenklach, M., Kinetic Modeling of Soot Formation with Detailed Chemistry and Physics: Laminar Premixed Flames of C2 Hydrocarbons. Combustion and Flame, 2000, (121(1-2) pp. 122-136.*

between the soot and gas phase. Soot is often ignored in coal furnace simulations. These tests showed some unexpected soot profiles, as shown below, which shows an instantaneous profile of soot number density. There is a distinct separation in the profiles, a so-called “lobed” structure. This structure appears in the tar, soot mass, soot number density, and in the coal gas mixture fraction. On inspection, it was found that this was due to the use of three particle environments in the coal DQMOM formulation. The particles disperse to differing extents in the reactor based on particle size, and result in the three lobes shown. We have increased the number of particle environments from three to six, with significant improvements, as shown below.

The soot number density, however, appears to be too high and this is resulting in soot particle sizes that are smaller than is expected. This high number corresponds to unphysically small particle sizes. We believe this is due to timescales for tar coagulation/soot nucleation that are much smaller than the simulation timestep, requiring explicit stiffness removal. We are investigating pseudo-steady state, semi-analytic, and sub-time-stepping methods for treating this stiffness.

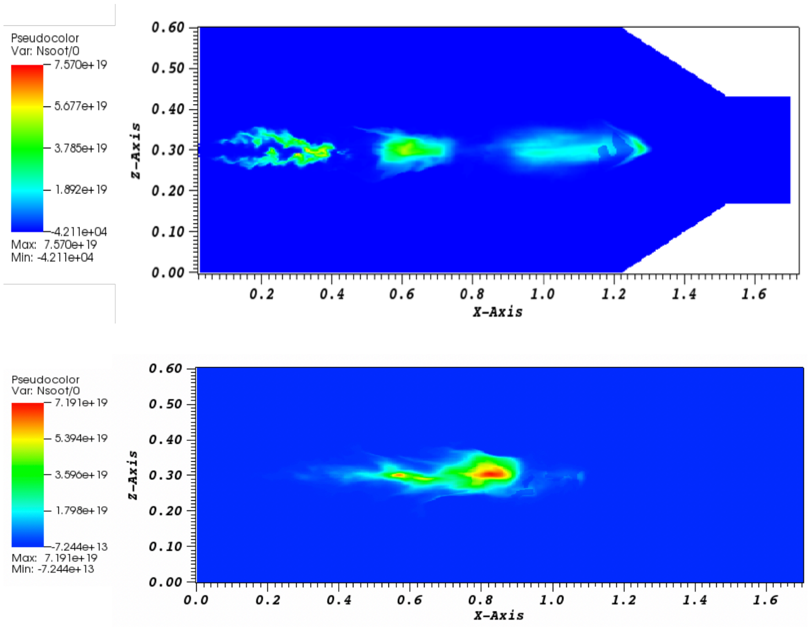

Figure 22: Soot number density on a 2-D plane in the reactor. Top: results with three particle environments; bottom: results with six particle environments.

Large Eddy Simulation Research

The efforts of the LES Integration Group for Year 5 of the PSAAP program are characterized into three categories; 1) Developments of the physical LES model with associated sub-models, 2) Verification activities along with support efforts of the VUQ simulations, and 3) Portability efforts through the adoption of Kokkos within the Uintah Framework and the Arches component. The following will briefly describe Year 5 highlights within the three categories.

The implementation of the Nonlinear Large Eddy Simulation (NLES, G. Burton) model was completed and demonstrated in Arches for periodic isotropic turbulence flows with decay. Implementation of a force turbulence scenario in Arches is proceeding for further testing. The implementation of NLES into Arches was done with the assistance of Greg Burton, LLNL, which included in-person meetings at Utah and at LLNL. The NLES model presents a nice contrast in turbulence closure to the more traditional gradient diffusion models employed within Arches. In particular, NLES offers a modeling approximation for the fluctuating subgrid velocities, which may help in obtaining a better representation of the drag for particles with low Stokes numbers. With the assistance of a visiting scholar, some ***a prior*** studies were conducted using DNS data investigating the Stokes number effects on particle dispersion of inert particles within periodic turbulence using diffusion-base models and no-model for particle drag. The model/no-model sensitivities were demonstrated. Demonstration with NLES closure for the particle drag interactions is currently underway in the ***a priori*** analysis. Additionally, studies were carried out ***a posteriori*** with Arches comparing the response of the particle energy spectra of different Stokes number particles when using the NLES subgrid velocity predictor vs. other models or no-model. The differences for certain classes of particle Stokes numbers are significant. We plan to continue to develop the NLES model and investigate the use of the subgrid model predictor for particle closure in turbulence flows.

Coal combustion in the boiler naturally has a wide range of length and time scales. LES modeling of this system at the scales of interest has allowed direct representation of several of these timescale - or at least resolution of a significant portion of these scales - which allows the reduction of modeling dependency. This concept has been applied to the particle reactions, devolatilization and char oxidation over the lifetime of this project. This year it was extended to the NOx and CO formation/destruction reactions for the gas phase. The NOx work was supported by drop tube data obtained from Tsinghua University. The NOx model is fairly complex, involving a series of NOx formation and destruction reactions through various pathways. The drop tube data also included burnout data at various temperature and stoichiometric conditions. The data set in whole represents a fairly well characterized set of data. As a result, a series of LES simulations were run performing VUQ for a set of model parameters related to the NOx model and for the VUQ of the char oxidation reactions. To this point, this work has produced an initial VUQ on the NOx model parameters. The VUQ for the char oxidation model is currently underway using the drop tube data.

The LES group also participated in the development and deployment of the wall deposition model, including the conjugate wall heat transfer model. Along with this, some work was done on exploring various turbulence wall closures. Given the extreme scenario uncertainty with the geometric wall profile. Ultimately, we settled on a fairly course wall closure model. This model was validated by comparing the LES profiles to DNS profiles of confined channel flow. Concerning the NLES model described above, work still needs to be done to handle near wall conditions considering the wide stencil required.

Regarding the second category (verification activities and VUQ simulation support), effort to support the production runs through various bug fixes, feature addition, etc. was done. This work is fairly non-specific and varied depending on the production case, but regularly involves a significant investment of the LES Integration group. Additionally, several verification studies were carried out to test code correctness (e.g., overall mass balances) and numerical verification of the various discretization schemes in Arches. More specifically, Arches code ported to the Uintah-Kokkos parallel_for expressions required complete rewrite, thus requiring code verification. Method of manufactured solutions and solutions with known analytical expressions were used and several tests were added to the nightly regression suite. Additionally, tests such as periodic turbulence decay were used to test the set of models ported to the Uintah-Kokkos paradigm. Major portions of the new code were demonstrated to show code correctness and proper numerical convergence based on the various discretization schemes offered.

Finally, significant progress was made in the third category regarding adoption of Kokkos within Arches. This has been a highly collaborative effort between the LES Integration and CS teams. The core Arches CFD algorithm was made fully ready for portability testing. This includes turbulence closure and Eulerian particle transport. Although several coal models still require porting into the Arches/Uintah/Kokkos abstraction, the most heavy-weight of them (Char Oxidation) was ported and extensively tested by the CS team and lessons learned were communicated between the CS and Physics group. Arches was demonstrated to run fully on the GPU within the Kokkos constructs. As the linear hypre solver for pressure does not offer a portable solution, an implementation of a preconditioned conjugate gradient method was implemented as part of this effort. This linear solver is useful for demonstration purposes but isn't terribly performant. Currently, there are no plans to optimize this solver. Although the portability demonstrations over the past year are encouraging, next year will require work to demonstrate performant portability of the entire CFD algorithm with particle combustion.

VALIDATION/UNCERTAINTY QUANTIFICATION

Validating Char Oxidation Models

Early in 2019, we examined an updated char model developed by Salvatore Iavarone. This new model differed from the previous iteration in a few ways. First, the initial fraction of char joined the initial particle diameter as an additional random quantity in the physical model. Second, the prior bounds on the uncertain parameters were narrowed. Third, the experimental bounds were slightly widened. A bound-to-bound data collaboration (B2BDC) analysis with quadratic surrogate models (and appropriate fitting error characterization) determined the dataset to be consistent. If the original experimental bounds were instead used, the dataset could be made consistent by removing 2 of the 399 quantities of interest.

Model discrepancy in B2BDC

We are currently investigating a new solution method to inconsistent datasets by adding a scenario-dependent discrepancy term to model outputs. The addition of this term is well motivated when we believe the inconsistency is due to an inadequate model rather than mis-specified experimental bounds. The general framework follows the Kennedy and O'Hagan structure[26], but the implementation in Bound-to-Bound Data Collaboration (B2BDC) leads to different interpretations, considerations, and consequences. The discrepancy takes the form of a linear combination of basic functions depending only on the scenario parameters. This structure leads to an extended feasible set in the space defined by the uncertain model parameters and discrepancy coefficients. In this context, prediction is performed with respect to the extended space.

Our current work is focused on developing this modified framework and examining the methodology with various discrepancy functions. For instance, polynomial discrepancy functions with different degrees were tested in both a toy mass-spring example and a realistic hydrogen combustion example. Preliminary results demonstrate the proposed method can always resolve dataset inconsistency when enough basis functions are included, but may suffer expanded prediction intervals. Comparing different choices of basic functions that resolve dataset inconsistency is still a very difficult and complex task and is the subject of future study. A manuscript summarizing this work is in preparation.

Uniform Sampling to the Feasible Set

We continued and finalized our analysis on evaluating the performance of applying the developed MCMC sampler within B2BDC to uniformly sample the feasible set. A strategy of allocating initial points employing optimization techniques showed good performance ensuring the over-dispersion condition of model parameters. In our test case, the developed sampler outperforms a Metropolis-Hastings sampler and an adaptive Metropolis sampler by displaying a larger step size and faster mixing. Uniform samples generated within the GRI-Mech dataset provide valuable complementary information to the deterministic bounds of feasible set and

---

[26] *Marc C Kenedy and Anthony O'Hagan. Bayesian calibration of computer models. Journal of the Royal Statistical Society: Series B (Statistical Methodology),2001, 63(3):425-464. doi:10.1111/1467-9868.0029*

predictions, offering a promising hybrid framework for uncertainty quantification. A manuscript summarizing this work is in preparation.

Linear Surrogate Models for Quadratic Ground Truth

In most applications of our tools, we fit quadratic or polynomial models to an underlying simulation or computer code using random samples. The accuracy of the fit can be estimated by separating the samples into training and test data, and the error can be absorbed into both consistency and prediction analyses. Different realizations of the random sample, however, can result in different surrogate models, which in turn results in different datasets and consistency measures. We are currently investigating how this sample-to-sample variability manifests when fitting linear surrogate models and datasets to approximate a collection of ground truth quadratic models. This scenario emulates the realistic case where the surrogate cannot capture all aspects of the underlying simulation. Our experiments with both the GRI-Mech and DLR datasets demonstrate that the sample-to-sample variability can produce non-negligible variation in the computed consistency measure. Fortunately, in both examples the conclusion of consistency vs. inconsistency is not significantly affected.

Inclusion of Ellipsoidal Constraints and Additional Criteria

In some cases, we may wish to include quantities of interest that are poorly represented by quadratic, rational quadratic or polynomial surrogate models in a dataset. This could also include alternate feasibility criteria, such as the notion of implausibility associated with Bayesian history matching. We are currently investigating the inclusion of sample-based feasibility assessments through the calculation of minimum volume covering ellipsoids. Importantly, this can be immediately implemented using the current B2BDC computational strategies (i.e., formulating semidefinite programs) as ellipsoids are represented through quadratic inequalities. We are also investigating the inclusion of additional constraints derived through support vector machines with degree two polynomial kernels.

Verification and Uncertainty Quantification of the BSF:

Test of surrogate models: The general assumption that allows us to use the current mathematical form of the Bayesian methodology used in the Center's work, is that of linear surrogate models. This greatly simplifies the numerical treatment and streamlines the analysis for rapid results. This also means that some of the non-linearities of the system might not be captured by the current surrogates. Introduction of non-linear surrogates might circumvent this issue. New analysis developed by the VUQ team, incorporating these changes were tested by running additional simulations for a new posterior from the previous analysis. Table 3 relates the total number of cases used to test non-linear surrogate models, and Figure 23 (Predictive posterior for the temperature, oxygen concentration and heat flux) shows the predictive posteriors for the BSF QOI's.

IU Map and new simulation cycle: The information obtained from previous VUQ cycles is extremely useful to update models and input parameters. The information gathered has been used to propose an input/uncertainty map (IU map) that relates the current values of the model/

scenario parameters with their relative importance. This information is presented in Table 4. Some of the cases rated as 5 are being currently run in the available computational resources.

Table 3: Cases and input parameters used to test non-linear surrogate models

	ThardB	CO2_PC1	CO2_PC2	O2_PC1	O2_PC2	AbsCD	log(Kdeposit)	devolTmu	ln(Avisc)	Tslag	log(enam)	Csmag	tSootBlow
min. val.	1750	-2.0	-4.0	-2.2	-1.3	0.80	-1.20	735	-54.7	1471	-4.70	0.05	1.2
max. val.	1950	5.0	4.0	2.2	3.0	1.90	1.20	930	-30.0	1549	-1.80	0.40	8.8
Case_1	1850	0.0	0.0	0.0	0.0	1.00	0.00	832	-54.7	1510	-3.00	0.20	5.0
Case_2	1850	1.3	0.0	0.0	0.0	1.00	0.00	832	-34.2	1510	-3.00	0.20	5.0
Case_3	1850	0.0	1.3	0.0	0.0	1.00	0.00	832	-34.2	1510	-3.00	0.20	5.0
Case_4	1850	0.0	0.0	1.3	0.0	1.00	0.00	832	-34.2	1510	-3.00	0.20	5.0
Case_5	1850	0.0	0.0	0.0	1.3	1.00	0.00	832	-34.2	1510	-3.00	0.20	5.0
Case_6	1850	0.0	0.0	0.0	0.0	1.00	0.00	832	-34.2	1510	-3.00	0.20	8.8
Case_7	1980	0.0	0.0	0.0	0.0	1.00	0.00	832	-34.2	1510	-3.00	0.20	5.0
Case_8	1850	0.0	0.0	0.0	0.0	1.00	0.00	832	-34.2	1549	-3.00	0.20	5.0
Case_9	1850	0.0	0.0	0.0	0.0	1.20	0.00	832	-34.2	1510	-3.00	0.20	5.0
Case_10	1850	0.0	0.0	0.0	0.0	1.00	0.00	832	-34.2	1510	-2.35	0.20	5.0
Case_11	1850	0.0	0.0	0.0	0.0	1.00	0.39	832	-34.2	1510	-3.00	0.20	5.0
Case_12	1850	0.0	0.0	0.0	0.0	1.00	0.00	930	-34.2	1510	-3.00	0.20	5.0
Case_13	1850	0.0	0.0	0.0	0.0	1.00	0.00	930	-34.2	1510	-3.00	0.33	5.0
Case_14	1850	0.0	0.0	0.0	0.0	1.00	0.00	832	-38.8	1510	-3.00	0.20	5.0
Case_15	1750	0.0	0.0	0.0	0.0	0.81	0.00	734	-34.2	1471	-3.65	0.07	1.2
Case_16	1750	-1.3	-1.3	-1.3	-1.3	0.81	-0.39	734	-38.8	1471	-3.65	0.07	1.2
Case_17	1859.42	3.73	-3.59	-0.80	1.14	1.61	-0.71	820.71	-35.46	1518.62	-3.76	0.11	7.22
Case_18	1861.05	2.91	3.59	-1.10	1.74	1.80	-0.86	843.29	-36.64	1516.05	-4.27	0.05	3.85
Case_19	1863.04	3.26	1.91	2.07	1.60	0.64	0.71	818.36	-32.44	1503.95	-1.95	0.30	7.68
Case_20	1836.96	1.60	2.26	-1.57	1.04	0.50	-0.61	841.57	-30.33	1502.86	-2.11	0.14	2.32
Case_21	1829.30	2.39	-1.91	1.10	2.87	1.43	-0.52	838.94	-31.25	1501.38	-3.89	0.12	3.64
Case_22	1865.74	2.63	-1.39	-2.07	2.37	1.50	0.86	825.06	-35.70	1504.84	-2.24	0.10	3.12
Case_23	1840.58	2.17	-1.63	-0.94	2.10	0.20	1.13	826.16	-32.13	1515.16	-4.67	0.35	6.88
Case_24	1834.26	1.97	2.73	-1.30	1.47	0.69	-1.13	822.43	-36.27	1506.31	-4.05	0.40	3.40
Case_25	1841.99	4.59	-2.26	1.57	1.90	0.57	0.61	840.16	-35.96	1505.61	-1.73	0.29	6.60
Base_Case	1850	0.0	0.0	0.0	0.0	1.00	0.00	832	-34.2	1510	-3.00	0.20	5.0

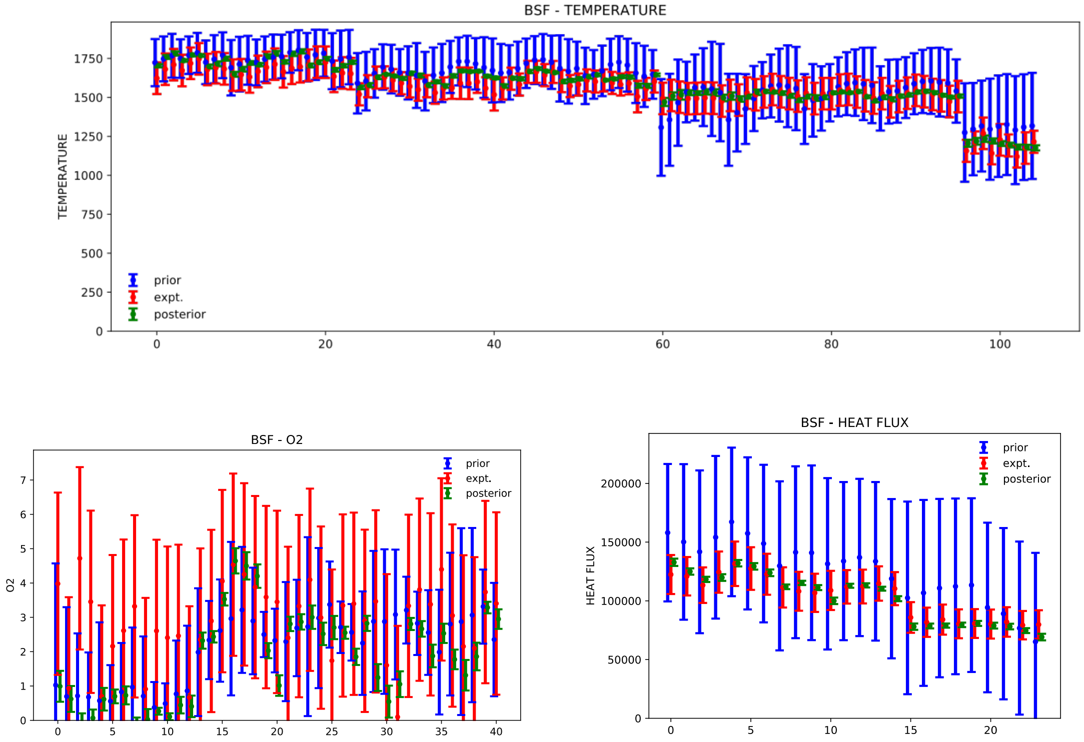

Figure 23: Predictive posterior for the temperature, oxygen concentration and heat flux

Table 4: Input/Uncertainty map (IU map) for model and scenario parameters

Parameters (H)		**Nominal Value**	**Uncertainty**	**Up Corner**	**Down Corner**	**Importance**
abs_Coef		0.636053693	0.24806094	0.884114633	0.387992753	3
PC1_O2		-1.83158207	-0.714317007	-2.545899078	-1.11726506	
PC2_O2		1.089131019	0.424761097	1.513892116	0.664369922	
Ae O2		2.021644437		0.590395073	6.92256155	5
Ea O2		11202.51147		7588.908543	14816.1144	5
PC1_CO2		2.00337885	0.781317751	2.784696601	1.222061098	
PC2_CO2		-0.183367821	-0.07151345	-0.254881271	-0.11185437	
Ae CO2		2054.992846		2666.597056	1583.664689	5
Ea CO2		63805.90996		69935.02124	57676.79868	5
ksi		0.97266	0.3793374	1.3519974	0.5933226	5
Relaxation coeff		0.005		0.05	0.001	5
k_ash_sigma			-0.277707214	-0.277707214	-0.27770721	5
T_hardened		1950.06337	380.2623572	2330.325727	1569.801013	4
v_hit		0.6202	0.241878	0.862078	0.378322	5
T_devol	T_mu	832	216.32	1048.32	615.68	5

T_devol	T_sig	134.6493799		189.4000451	79.89871477	5
SgO		1010.8	394.212	1405.012	616.588	5
porosity		0.6		0.7	0.35	5
ash fluid temperature	T_hemi					
ash fluid temperature	T_soft					
ash fluid temperature	T_fluid	1510	196.3	1706.3	1313.7	4
viscosity pre-exp. factor		1.40E-15		1.78164E-18	1.10519E-12	5
ln(viscosity pre-exp)		-3.42E+01	-6.67E+00	-4.09E+01	-2.75E+01	
viscosity activation energy		49.42441876	9.637761657	59.06218041	39.7866571	5
HHV		20952.63	4085.76285	25038.39285	16866.86715	5
particle density		1300	253.5	1553.5	1046.5	5
particle initial void fraction		0.1		0.35	0.08	5
Composition		ND lignite		German Lignite		4
Ta		19910.00	7764.90	27674.90	12145.10	4
A		2.29087E+11		3.76011E+13	1395725458	4
ln(A)		26.15736666	5.100686498	31.25805315	21.05668016	
soot density		1900	370.5	2270.5	1529.5	3
particle thermal cond.		4	1.56	5.56	2.44	4
K_main		3.45	0.897	4.347	2.553	5
enamel thickness (mm)		3.09E-02		1.5	0.5	5
t_scale dep (s)		788836.091		432000	129600	4
Wall emissivity		0.8	0.104	0.904	0.696	4
Rad solve frequency		20		40	10	4
Angular Discretization		8		16	4	4
Re_limit		10000	3900	13900	6100	4
Gas phase temporal scheme		2nd order		3rd order		3
opl		0.9488	0.37	1.32	0.58	3

Verification and Uncertainty Quantification of the L1500

Our major focus this year was the detailed budgeting and test planning for a final L1500 test campaign with improved methods for collecting and analyzing combustion data in support of the simulations. First, we designed an ash deposition probe. Based on the analysis of simulation results from prior years, we identified the ash deposit characteristics as the main quantities of interest (QOIs) for the L1500 brick. However, during the 2016 campaign, little ash was collected on the heat transfer surfaces, probably due to the low hot-side metal temperature of the cooling panels. Therefore, we designed an ash deposition probe such that the hot-side temperatures would be similar to water wall temperatures in a boiler (and sufficiently high to collect ash). We used both RANS and LES simulations to prototype detailed representations of circular and square probe designs. We completed a suite of 282 simulations with 18 input parameters, including thickness of the ash and enamel layers; density, thermal conductivity, and specific heat for ash, enamel, and steel; and the cooling water flow rate. Based on the results of these simulations, we finalized a square probe design and created detailed CAD drawings for manufacturing.

Second, we examined the narrow-angle radiometers in great detail, focusing on the design, material science of the materials of construction, and the electronics and data processing. Our goal was to develop a more physics-based radiometer instrument model and an improved calibration procedure. We are in the process of refining our new instrument model, which we will present in our Q1 report for Year 6. The millivolt signal output of the radiometers is proportional to the intensity of the radiation incident on the thermopile but is also affected by the ambient temperature of the thermopile itself. We calibrated all four narrow-angle radiometers in the temperature range from 600-2200°C using a black body radiation source. We then created a simple regression of the calibration data taken near the ambient temperature of operation during the campaign. Figure 8 shows the intensity of the flame as measured by the radiometers with this calibration for a 45-minute window on the morning of November 13, 2018. We will present this analysis at the Clearwater Clean Energy Conference in June 2019. We also calibrated the view angle for all four narrow-angle radiometers. The view angles range from 0.7-1.06°. On the simulation side, we have modified the discrete ordinates radiation model to compute narrow-angle intensities that are aligned with the coordinate axes since the three narrow-angle radiometers mounted in the L1500 were aligned with one axis and perpendicular to the other two. The orthogonal sweeps are saved in the middle of the intensity array. For example, if you use SN8, this suggests 80 sweeps, but orthogonal sweeps will run 8 extra sweeps for a total of 88 sweeps. This means that sweeps 0-9 are 10 standard x+y+z+ octant directions with #10 (the 11th element) being an orthogonal sweep. This change was recently implemented, so we will report results in the Q1 report for Year 6.

Third, we determined the volumetric flow rate of gas sampled by both a high-velocity thermocouple (suction pyrometer) and a gas sampling probe. We attached a large plastic bag over the sampling end of the probe, filled the bag with air, and measured the diameter and length of the bag when fully inflated. Then, we attached the inductor using house air and measured the time to deflate the plastic bag. The procedure was repeated multiple times. The flow rate is needed for determining the corresponding sampling volume from the simulation.

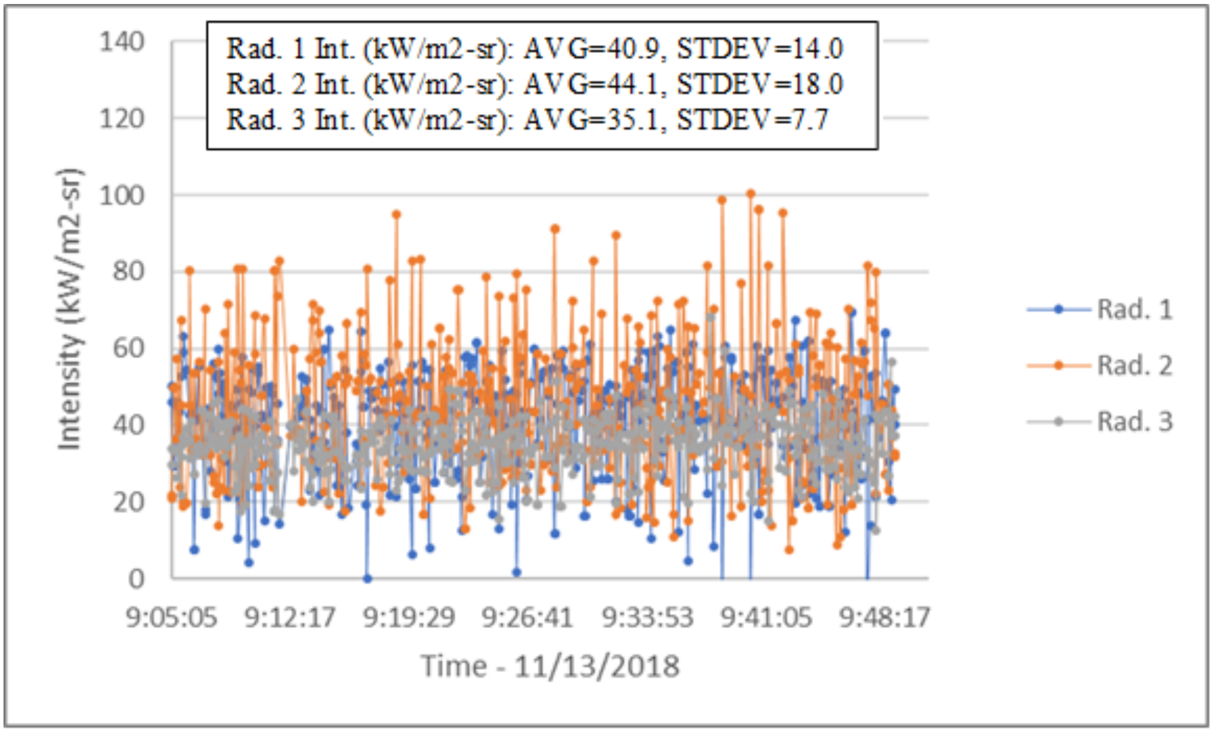

Figure 24: Intensity data taken by the narrow-angle radiometers on Nov. 13, 2018. The fuel is pulverized Sufco coal and the burner swirl setting is 0%. Radiometers 1, 2 and 3 are approximately 0.6, 1.8, and 3 m downstream from the burner, respectively.

Due to unfortunate circumstances, the proposed L1500 test campaign was scrapped. However, we completed a reduced test program by leveraging a test campaign sponsored by Pacificorp on November 12-16, 2018. We collected data only when the L1500 was firing 100% coal (initial furnace line-out) and without burner swirl.

Prior to the start of the campaign, we reviewed data from previous campaigns and identified equipment changes/modifications that would reduce experimental uncertainty with minimal capital investment. There are multiple water-cooled ports, panels and coils in the furnace. The heat removed by these surfaces needs to be well-characterized for comparison with simulation output. We identified the manufacturer-specified flow meter accuracy of ±15% as the greatest source of uncertainty in estimating heat removal. To reduce this uncertainty, we purchased and installed flow reducers for the flow meters, which increased their manufacturer-specified accuracy to ±7%. We then calibrated the flow meters on each water line (16 total) before and after the campaign using a flow meter with factory-certified accuracy of ±1.5%. Figure 25 shows calibration lines for a single flow meter taken before and after the test campaign in November 2018. The data points have a strong linear regression and are very similar between dates, which would indicate that flow meter operation remained stable throughout the campaign.

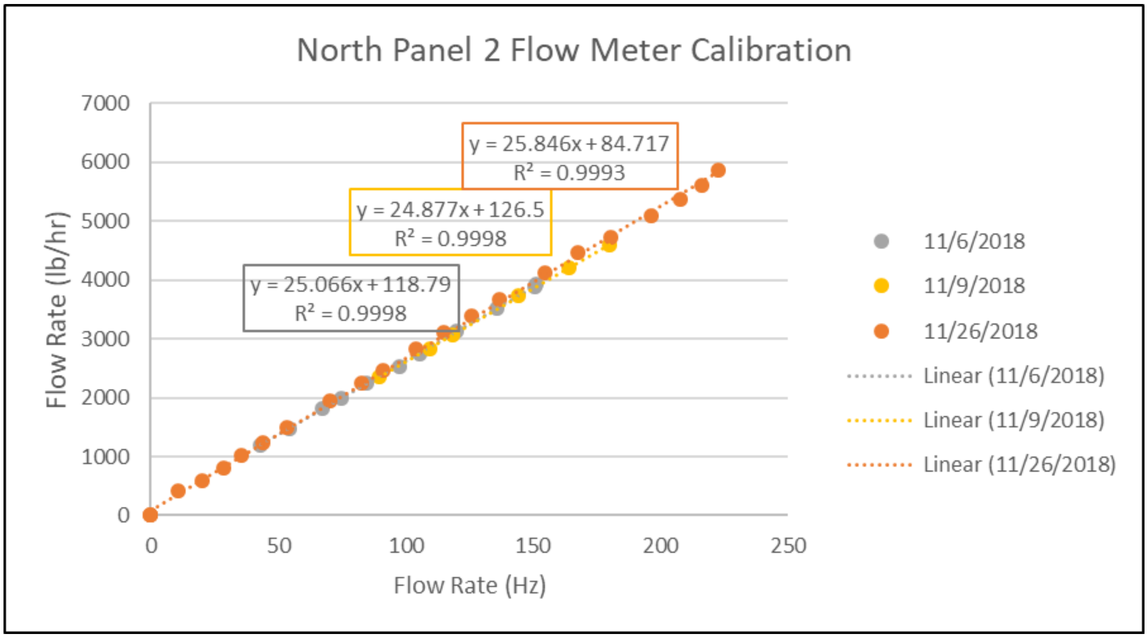

Figure 25: Calibration lines for the flow meter on the second heat flux panel from the burner on the north side of the L1500 furnace.

Because of geometric limitations of the experimental setup, we did not meet manufacturer's requirements for proper installation. Therefore, we performed additional calibration procedures to quantify these effects, including installation of two flow meters in a series.

We also collected from the walls and ceiling of the L1500 along its length. The emittance of these deposits will be characterized as a function of temperature. We presented emittance data as a function of temperature (up to 1000 °C) for L1500 coal ash. An apparatus to measure emittance data at even higher temperatures is in the final stages of fabrication.

We ran a suite of simulations of the November 2018 experiments to determine sensitivity of the QOIs to various parameters including the frequency of the discrete ordination radiation model solve, the use of a spectral model or the Hottel-Sarofim model for the gas phase absorption coefficient, the relaxation coefficient for the wall heat transfer model, and the presence (or not) of a deposition probe. Figure 26 shows a comparison of heat flux measured at four locations on the two water-cooled panels (south side of the L1500) for several of these variants. There is some sensitivity to all of these factors, but we need to optimize for efficiency and accuracy; for example, it takes longer to run a case with relaxation coefficient of 0.05 than 0.1. Before we do this optimization, we need to finish the Bayesian analysis of our experimental data as detailed in the next paragraph. The data in this figure was generated using a traditional propagation of error method and we believe that our new Bayesian analysis will more accurately capture both the mode and the uncertainty of the data.

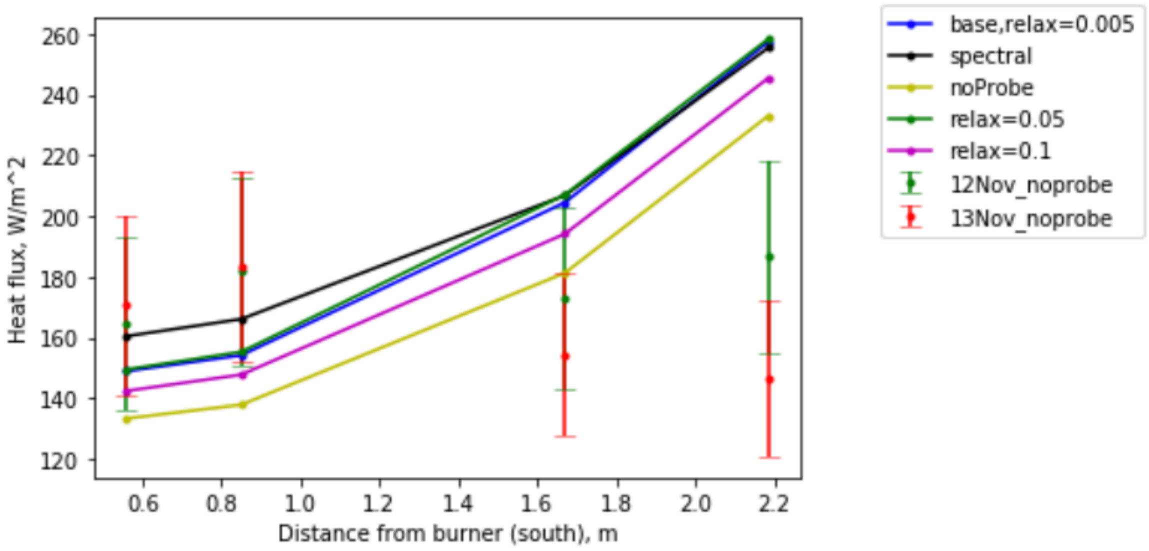

Figure 26: Comparison of heat flux data on the south side of the L1500 furnace from both simulations and the November 2018 experimental campaign.

Several members of the L1500 team took a short course on uncertainty quantification offered by CCMSC personnel (Prof. Sean Smith and Prof. Phil Smith). This rigorous course, which focuses on Bayesian methods, has transformed the lens through which our L1500 team analyzes the data. We have developed new instrument models for both radiation intensity (narrow-angle radiometers) and heat flux (water-cooled panels and multi-depth thermocouples in those panels) measurements in the L1500.

PUBLICATIONS & PRESENTATIONS

Publications

- A. J. Jospehson, R.R. Linn, D.O.Lignell. "Modeling Soot Formation from Solid Complex Fuels," ***Journal of Combustion & Flame***, 2018, 196:265-283. ***doi: 10.1016/ j.combustflame.2018.06.020***
- S. Iavarone, J. Oreluk, S.T. Smith, A. Hegde, W. Li, A. Packard, M. Frenklach, P.J. Smith, F. Contino, A. Parente. "Application of Bound-to-Bound Data Collaboration approach for development and uncertainty quantification of a reduced char combustion model." ***Fuel***, 2018. ***doi:10.1016/j.fuel.2018.05.113***.
- J. Oreluk, Z. Liu, A. Hegde, W. Li, A. Packard, M. Frenklach, D. Zubarev. "Diagnostics of Data-Driven Models: Uncertainty Quantification of PM7 Semi-Empirical Quantum Chemical Method." ***Nature: Scientific Report***s, 2018. ***doi: 10.1038/s41598-018-31677-y***
- J. Parra-Álvarez, B. Isaac, S. T. Smith, M. Zhou, T. Ring, S. Harding, and P. Smith***.*** "Radiative Properties of Coal Ash Deposits with Sintering Effects." ***Energy and Fuels. doi: 10.1021/ acs.energyfuels.8b04206.***
- D. Hoang, P. Klacansky, H. Bhatia, P.T. Bremer, P. Lindsgtrom and V. Pascucci. "A Study of the Trade-off Between Reducing Precision and Reducing Resolution for Data Analysis and Visualization." ***IEEE Transactions on Visualization and Computer Graphics*** 25(1):1-11, 2019.
- A. Gyulassy, P. Bremer and V. Pascucci,."Shared-Memory Parallel Computation of Morse-Smale Complexes with Improved Accuracy***." IEEE Transactions on Visualization and Computer Graphics***, 25,(1):1183-1192, 2019. ***doi: 10.1109/TVCG.2018.2864848***
- F. Wang, I. Wald, Q. Wu, W. Usher and C. R. Johnson. "CPU Isosurface Ray Tracing of Adaptive Mesh Refinement Data." ***IEEE Transactions on Visualization and Computer Graphics,*** 25(1):1142-1151, 2019.
- W. Usher, S. Rizzi, I. Wald, J. Amstutz, J. Insley, V. Vishwanath, N. Ferrier, M. E. Papka, and V. Pascucci. "libIS: A Lightweight Library for Flexible In Transit Visualization***". ISAV: In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization (ISAV '18***), 2018. ***doi:10.1145/3281464.3281466***

Conference Publications

- M. Frenklach. "Bound-to-Bound-Data-Collaboration: Prediction on the Feasible Set." SAMSI MUMS Program Opening Workshop, Duke University, Durham, NC, August 20-24, 2018.
- Q. Wu, W. Usher, S. Petruzza, S. Kumar, F. Wang, I. Wald, V. Pascucci, and C. D. Hansen. "VisIt-OSPRay: Toward an Exascale Volume Visualization System." Eurographics Symposium on Parallel Graphics and Visualization, Brno, Czech Republic, June 4, 2018.
- C. Hansen. "VisIt-OSPRay: Toward an Exascale Volume Visualization System." Eurographics Symposium on Parallel Graphics and Visualization, Brno, Czech Republic, June 4, 2018.

- A. Sanderson, A. Humphrey, J. Schmidt and R. Sisneros. "Coupling the Uintah Framework and the VisIt Toolkit for Parallel ***In Situ*** Data Analysis and Visualization and Computational Steering." High Performance Computing, ISC High Performance. Frankfurt, Germany, June 24-28, 2018.
- J. McConnell, and J. C. Sutherland. "An a priori Analysis of Mixture Fraction-Based Modeling of Coal Combustion." Western States Section of the Combustion Institute (pp. 1–13). Bend, OR, March 25-27, 2018.
- T. Saad, M. Karam, and J. C. Sutherland. "An Explicit Variable-Density Projection Method for Low-Mach Reacting Flows on Structured Uniform Grids." Fluid Dynamics Conference, AIAA AVIATION Forum, (AIAA 2018-4266), Atlanta, GA, June 25 - 29, 2018.
- J. Parra-Álvarez, B. Isaac, S. Smith, T. Ring, S. Harding, and P. Smith. "Radiative properties of ash deposits." Proceedings - 27th International Conference on the Impact of Fuel Quality on Power Production and the Environment, Lake Louise, Alberta, September 24-28, 2018.
- J. Parra-Álvarez, B. Isaac, S. Smith and P. Smith. "Predictivity: Definition and Application to a Tangentially-fired Combustion System." AFRC 2018 Industrial Combustion Symposium, Salt Lake City, UT, September 17-18, 2018.
- J. Parra-Álvarez B. Isaac, S. Smith and P. Smith. "Definition of Predictivity from a Machine-Learning Perspective: Application to a Tangentially Fired Oxy-coal Combustion System." 43th International Technical Conference on Clean Energy, Clearwater, FL, June 3-8, 2018.
- J. Parra-Álvarez, S. T. Smith, B. Isaac, M. Zhou, and P. Smith. "Uncertainty Quantification and Validation of a 15MW Oxy-fired Coal Combustion System." 11th Meeting of Joint Sections of The Combustion Institute, Pasadena, CA, March 24 – 27, 2019.
- K. Brinkerhoff, A.J. Josephson, B. Isaac, J. Thornock, and D. Lignell. "The effect of soot in oxy-coal combustion systems." (Poster), 11th Meeting of Joint Sections of The Combustion Institute, Pasadena, CA, March 24 – 27, 2019.
- J. McConnell and J. C. Sutherland. "Assessment of Various Tar and Soot Treatment Methods for Use in Coal Combustion Simulation." 11th Meeting of Joint Sections of The Combustion Institute, Pasadena, CA, March 24 – 27, 2019.
- A. Sanderson. "Coupling the Uintah Framework and the VisIt Toolkit for Parallel ***In Situ*** Data Analysis and Visualization and Computational Steering." ISC High Performance 2018 International Workshop in ***In Situ*** Visualization, June 24 – 28, 2018.
- J. McConnell and J. C. Sutherland. "Coupling an Explicit Variable Density Projection Method to Finite Rate Kinetics." SIAM Conference on Computational Science and Engineering, Spokane, WA, February 25 - March 1, 2019.
- V. Pascucci and A. Gyulassy. "Shared-Memory Parallel Computation of Morse-Smale Complexes with Improved Accuracy." IEEE Visualization 2018 Conference, Berlin, Germany, October 2018.
- S. Petruzza, A. Gyulassy, V. Pascucci and P. T. Bremer. "A Task-Based Abstraction Layer for User Productivity and Performance Portability in Post-Moore's Era Supercomputing." Proceedings of 3rd International Workshop on Post-Moore's Era Supercomputing 2018 (PMES2018).

Presentations

- V. Pascucci, Keynote Presentation. “Extreme Data Management, Analysis, and Visualization for Scientific Discovery and Economic Development.” IEEE Cyber Science and Technology Congress, August 14, 2018.
- V. Pascucci, Keynote Presentation. “Extreme Data Management Analysis and Visualization for Exascale Supercomputers and Experimental Facilities.” AMITY International Conference on Artificial Intelligence - AICAI’19, Dubai, India, September 25, 2018.
- V. Pascucci, Invited Talk. “Extreme Data Management Analysis and Visualization for Exascale Supercomputers and Experimental Facilities.” Big Data Meets Large-Scale Computing, Los Angeles, September 25, 2018.
- A. Gyulassy, “Introduction to topological methods for data analysis,” Overview talk for TTK Tutorial at IEEE Visualization 2018, Berlin, Germany, October 2018.
- W. Usher and A. Gyulassy. Section of the TTK tutorial at VIS18 on TTK’s build infrastructure, Berlin, Germany, October 21 – 26, 2018.
- V. Pascucci. Keynote Presentation, “Extreme Data Management, Analysis and Visualization for Scientific Discovery and Economic Development.” ISUM Congreso de Supercomputo, March 28, 2019.
- V. Pascucci. Invited Talk, “Extreme Data Management, Analysis, and Visualization for Scientific Discovery and Economic Development.” Riken Supercomputing Center, Kobe, Japan, March 18, 2019.