Annual Technical Report
DE-NA0002375
March 1, 2014 through February 28, 2015

Carbon Capture Multidisciplinary Simulation Center
University of Utah – Institute for Clean and Secure Energy

Philip J. Smith, Principal Investigator
April 30, 2015

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

Highpoints of year-one computer science infrastructure include runtime enhancements with PIDX scalability to 100K+ cores and public release. RMCRT scalability to all of Titan's CPU nodes was achieved. Dramatic improvements were made in volume rendering using compositing. We pursued concurrent rendering of VisIt animation frames with parallel visualization of each frame for improved DMAV. The DSL has been extended to provide GPU support, and the framework has been modified to improve GPU performance when using the DSL. The DSL has been incorporated into several parts of the production code (Arches) resulting in approximately 3x speedups on single-core execution timings.

Highlights of accomplishments made by our physics team include improving robustness with NDF representation as number densities approaching zero. We developed heat transfer and reaction models that relax to arbitrary specified length and time scale. LES spatial and temporal scales are defined by the local Courant No. We explored a devolatilization model for uncertainty with validation and uncertainty quantification. The team addressed discrete ordinates of particles with scattering, improved particle, gas and deposit radiate properties analyses, and made advancements in RMCRT for future implementation.

In year-one we achieved the goal of full system integration across all levels of our validation hierarchy. Experiment and simulation are coupled with sensitivity studies to determine parameters and resources. A new devolatilization model was created to address the model form uncertainty of what was estimated, prior to the year 1 analysis, to be our most uncertain and sensitive physics component. After completing one cycle through our validation hierarchy we determined that the wall thermal properties were more significant than the devolatilization model; and thus, our priorities for year two were adjusted to focus more on coal ash transformations to quantify the wall deposit properties more fully.

Our year-one experience highlighted the constraint on our problem of the limited computer resources available for our center mission. In year 2 we will work on optimizing the use of this limited resource.

ACCOMPLISHMENTS

Outreach and Education

The Computing team hosted a "Deep Dive" workshop in July 2014 with representatives of LLNL, LANL and SNL to discuss details of the programming models we employ for our framework. The workshop was held in Salt Lake City and was attended by about 40 people, overall, from Utah, Los Alamos, Sandia and Livermore. Talks and informal discussions around both DOE Lab and Utah code and future developments took place. The workshop was judged to be a success by those who attended and the format has since been exported to the other PSAAP Centers.

A "Mini-Apps" workshop was held in Salt Lake City over November 10 to 12 that focused on implementing a compressible CFD algorithm (miniAero) within the Uintah run-me system. This was follow-on to the "Deep Dive" Workshop held in Salt Lake City in July 2014. It was designed to expose all of the details related to developing a new component in Uintah. Nine staff members of Sandia National Lab participated; namely, Janine Bennett, Ken Franko, Hemanth Kolla, Steve Bova, Keita Teranishi, Paul Lin, Greg Sjaardema, Gary Templet and Matt Bettencourt. The workshop provided an introduction to Uintah's patch based domain decomposition along with the "rules" for writing a task. There was also a discussion on what the framework provides to component developers including geometry objects for domain initialization and boundary conditions. Details on requesting data from neighboring patches was discussed along with the available patch iterators. The original miniAero algorithm was broken down into tasks and the data dependencies between the tasks defined. This formed an initial software design that was iterated upon. Two coding teams were then formed that implemented software for a generalized simulation initialization specification, boundary conditions and the miniAero algorithmic tasks. The Uintah:miniAero component was running on a single processor for a single time-step. At the end of the workshop, the 1D Riemann shock tube problem was executed within Uintah's new miniAero component, performing all specified tasks, but producing the wrong answers. Since the workshop the Uintah and Sandia teams have continued development of the Uintah:miniAero component, adding higher order spatial decretization scheme and a Runge Kutta RK4 algorithm for time stepping. Given that the component had not been run a large scale previously, this is significant. We continue to work with the Sandia team on performance analysis of the MiniAero component.

The V/UQ faculty of the University of Utah and the University of California, Berkeley designed a course to be taught in fall term 2015 for students of both institutions. The course is designed for first year graduate students. It will be a project-oriented course with themes and lectures covering (1) Introduction and Motivation for V/UQ, (2) Semi-definite Programming, (3) Surrogate Modeling, (4) Experimental Uncertainty, (5) Dimensionality Reduction, (6) Kennedy O'Hagen Analysis, (7) MCMC Sampling, (8) Bounds-to-Bounds Analysis and (9) Practical Workflow. The course will be open to personnel of the DOE laboratories who are interested.

Five graduate students have completed internships at one of the national labs. **Cameron**

Christensen worked with Eric Brugger at Lawrence Livermore on heterogeneous computing for visualization and analysis pipelines. **Ben Schroeder** went to New Mexico to work with Walt Woodkowski at Sandia National Lab for studies on methods for model from uncertainty and evaluation of predictivity. Dr. Peer-Timo Bremer at LLNL hosted **Aaditya Landge** to work on building large scale in-situ topological analysis algorithms and techniques that enable extraction of feathers of interest while the simulation is running. They also began extending these techniques for Adaptive Mesh Refinement (AMR) meshes. **Pascal Grossett** worked with Al McPherson and Ben Bergen at Los Alamos National Lab and in a team of six students on a co-design project entitled "Evaluating Distributed Runtimes in the Context of Adaptive Mesh Refinement." The interdisciplinary team (3 computer scientist and 3 mathematics) evaluated three different runtime environments; Charm++, HPX and Intel CnC in the context of the Sod Shock problem.

Mark Kim's internship at LLNL comes this summer and **Allen Sanderson** (research staff) is planning to spend time at LLNL sometime in the next year. **John Holmen** is hoping to attend the Argonne Extreme Scale Summer School that runs in the summer of 2015 and is interested in completing an internship at Sandia National Lab in 2016. **Dr. Tony Saad** (research staff) is currently planning a visit to SNL Albuquerque with Dr. Stefan Domino August 3-7, 2015. This visit will focus on pressure projection algorithms for variable density flows. **Oscar Diaz-Ibarra** is currently exploring opportunities at the national labs for an internship in early 2016. He is particularly interested in combustion research related to char oxidation and in V/UQ techniques. **Troy Holland** will begin his internship at LANL in October 2015, under the mentorship of Dr. Joel Kress.

Computer Science: TASC/EDSL, DMVA, Exascale Runtime

Uintah 1.6.0 was released in January 2015. It includes the Nebo DSL, which is fully integrated into Wasatch and is being incorporated into the Arches component as features are hardened (in Wasatch) to be production-capable. This minimizes disruption in Arches. The DSL has GPU support, and this will be further expanded and improved over the coming year. We coordinated the release of VisIt 2.9.0 as the Uintah database read and is now part of the VisIt build and installation script. We have made the core Uintah System work better for Apps. We also worked on interacting with Apps solution with Visus PIDX and VisIt. Uintah now runs and scales on all major DOE and NSF platforms (Mira, Titan, Stampede, Vulcan and BlueWaters).

Runtime Systems & Infrastructure Improvements

A strong scalability study was performed on ALCF:Titan using the MPI/threaded scheduler and the adaptive mesh capabilities of the RMCRT algorithm for solving the RTE. Good strong scaling characteristics were shown up to 262K cores on three problem sizes. The grid consisted of 2 levels with a refinement ratio of two. The total number of cells on the highest resolved level was 128^3, 256^3 and 512^3. The standard Burns & Christon benchmark test problem, using 100 rays per cell, was solved.

The Reverse Monte Carlo Ray Tracing (RMCRT) based radiometer capability was extracted from the core RMCRT code and re-implemented for use in conjunction with the Discrete Ordinates

Method (DOM). Users running production cases need the ability to directly compare RMCRT:radiometer point flux calculations against experimental heat flux sensors while solving the Radiative Transport Equations with DOM. Scaling/serial studies of the radiometer code showed that the cost of the computations is a small fraction of the over all cost. By far the dominant cost is the all-to-all communications, which is also true for the RMCRT code.

Two techniques for reducing the communication costs are being investigated. The first method involves casting the communicated variables as floats instead of doubles. In theory this should reduce the cost by a factor of 2. The second method involves using Uintah's AMR infrastructure to communicate a coarse representation of the radiative properties. Depending on the refinement ratio and the number of levels used the reduction in cost could be a factor of 8 or greater.

Analysis was done on how the CPU-based version of the Reverse Monte Carlo Ray Tracing-based radiation model performs on the Intel MIC Architecture. Working towards this goal, the algorithm was ported to the Intel Xeon Phi and experiments were conducted examining native execution on a single co-processor. These experiments were used to better understand the strengths and weakness of the current algorithm while identifying optimal run configurations and gathering baselines for future efforts. These experiments have also targeted understanding how to effectively manage the Xeon Phi's reserved core through an examination of varying thread placement strategies.

Through this year's work we identified that the compute allotment is insufficient for desired V/UQ runs. Vulcan is heavily subscribed and throughput for long running production runs is challenging. Security restrictions impose challenges for accessing and moving data. We recommend that version control restrictions – svn access – ports be opened. Also it was found that every <patch> listed in a (large) XML file uses a huge amount of memory to store the XML file object when the code runs. We need to avoid creation of XML DOM data structure. By streaming dynamically parsed data directly into our data structures 8192 patches memory was reduced by 40 MB per core on Vulcan. A new scheduler solved the MPI buffer memory issue and the Load Balancer was changed to support RMCRT. Faster compile times and smaller compilation units and binaries were achieved by removing 50K lines of unnecessary code.

Uintah has been ported to all major Tri-lab clusters including: Sandia (Chama), LLNL (Vulcan, Cab, Syrah, Surface), and LANL (Mapache, Mustang, Wolf). This port included libraries such as Boost, Hypre, and PETSc. Additional infrastructure support tasks completed during year 1 include CUDA/GPU build support, reducing memory footprint of Uintah's XML datastructures, and supporting applications running at scale on various platforms.

Data Management Analysis and Visualization

The classical model of compute first, analyze later will soon be infeasible. A new model of DMAV is needed at exascale. Challenges include the widening gap between compute and I/O rates and markedly different characteristics of analysis codes compared to simulation codes. Data-dependent computations and global communications are difficult and expensive at scale. Understanding and modeling end-to-end workflows will be critical to Big Data at exascale. A flexible infrastructure is needed for empirical evaluations of the characteristics of overall data movement. Recognizing the challenges we adopted strategies that maximize scientific productivity at all stages.

Work was performed for better parallel methods for volume rendering and basic research into GPU-based flow visualization to improve algorithmic scaling. Data streaming is at the core of our Big Data management, analytics and visualization framework. The framework is evolving towards a real-time evaluation of "expressions" that can be delivered as "virtual datasets."

<u>Visualization</u>

We developed a new parallel volume rendering compositing algorithm for hybrid MPI parallelism that focuses on communication avoidance and overlapping communication with computation at the expense of evenly balancing the workload. This method is called the TOD-Tree: Task-Overlapped Direct send Tree Image Compositing for Hybrid MPI Parallelism. The algorithm has three stages: a direct send stage where nodes are arranged in groups and exchange regions of an image, followed by a tree compositing stage and a gather stage. Comparing this new algorithm with radix-k and binary-swap from the IceT library in a hybrid OpenMP/MPI setting, shows good strong scaling results. This has been published along with detailed explanation of how performance improvements were obtained over these other two well-known and widely accepted algorithms.

The goal of linking Unitah and VisIt at runtime is to allow for in-situ analysis and visualization. The in-situ analysis and visualization is for both simulation data and run time data. The simulation data includes not only the data that would saved to disk, but also intermediate data that would not normally be saved but is used as part of the simulation. The run time data could include memory usage, time spent per processor, I/O, etc. all of which are performance indicators

Isosurface extraction is a fundamental technique used for both surface reconstruction and mesh generation. One method to extract well-formed isosurfaces is a particle system; unfortunately, particle systems can be slow. In this research, we introduced an enhanced parallel particle system that uses the closest point embedding as the surface representation to speedup the particle system for isosurface extraction. The closest point embedding is used in the Closest Point Method (CPM), a technique that uses a standard three-dimensional numerical PDE solver on two-dimensional embedded surfaces. To fully take advantage of the closest point embedding, it is coupled with a Barnes-Hut tree code on the GPU. This new technique produces well-formed, conformal unstructured triangular and tetrahedral meshes from labeled multi-material volume datasets. Further, this new parallel implementation of the particle system is faster than any known methods for conformal multi-material mesh extraction. The resulting speed-ups gained in this implementation can reduce the time from labeled data to mesh from hours to minutes and benefits users who employ triangular and tetrahedral meshes in their simulations.

Based on the closest-point embedding on the GPU, we introduced a novel flow visualization technique for arbitrary surfaces. This new technique utilizes the closest point embedding to represent the surface, which allows for accurate particle advection on the surface as well as supports the unsteady flow line integral convolution (UFLIC) technique on the surface. This global approach is faster than previous parameterization techniques and prevents the visual artifacts associated with image-based approaches.

<u>Domain-Specific Language</u>
We made substantial progress in adding basic capabilities to our domain specific language (DSL) to support deployment on GPU architectures. These include (1) reduction operations (min, max, norm, etc.), (2) conditional operations - the ability to perform calculations on subsets of the domain where a condition is met, (3) masked field operations to enable boundary condition treatment through DSL, and was a major step forward to enable GPU support for a much larger segment of the application codes, and (4) boundary condition support for discrete operator inversion and using masks allowing boundary conditions to be evaluated on GPU.

These capability improvements now enable the majority of Wasatch to deploy to GPU. Wasatch is seeing modest speedups on GPU currently (~6x) while the same code outside Uintah is seeing 10-15x speedups. We will be working with the framework team during the upcoming year to improve these performance gaps. Indeed, during year-one we worked with the framework team to identify and correct several performance critical issues regarding GPU execution in Uintah. As this technology is hardened in Wasatch, it is transitioned to Arches through in-place refactoring. Arches now has demonstrated basic scalar transport using the DSL and achieved ~3x speedups as a result.

During the first year of the project, we also implemented a component-agnostic framework for Lagrangian particle transport within Uintah. Both Wasatch and Arches use this, and it took Arches only one day to implement basic particle transport using this helper.

Finally, as part of our participation in the "Deep Dive," we significantly improved the documentation of our DSL software library, including examples that are compiled as part of our nightly test system on both GPU and CPU systems and automated nightly updates of the documentation.

Among the many other improvements and novel contributions, we list the following:

1. Development of a base interface for GPU-support for boundary conditions
2. Develop algorithm for particle relocation across processor and patch boundaries that does not require the tracking of temporary particle variables
3. Develop algorithm for finding particles near boundaries
4. Allow components to access particles near boundaries
5. Develop algorithm to inject particles into the domain
6. Allow boundary conditions to compute and return their surface areas
7. Improve the Uintah boundary conditions interface by removing storage of un-necessary boundary condition data

Resource issues with ARCHES on Vulcan for prediction were unexpected and delaying. The I/O crashes loading with slow write times and unknown I/O errors. There is slow data interrogation and the File System limits the number of files. It is hard to compute with large numbers of parallel jobs (180 million hours versus 35 million hours). Vulcan is 30 times slower than Utah's

ASH on a core-to-core basis. This encourages small patch sizes, but makes file number limitations more pronounced. Moving to a threaded scheduler may help the speed.

Physics: Radiation, Particle Combustion, Multiphase Flow, LES Environment

During year one, the following tasks were accomplished within the Physics Effort:

- Established the Arches Task Interface (ATI) for better task granularity, ease of implemented new physics, and to allow access to the CCMSC EDSL. Some physics components within Arches were moved into the ATI. This work ongoing.
- End-to-end simulation efficiency was improved, enabling the VUQ work to be performed in a timely manner.
- Radiation physics were improved. For the discrete ordinates model, better gas and particle phase models were implemented. For RMCRT a complete integration into Arches was demonstrated. In addition, better sampling methods that increased the algorithmic efficiency were also introduced and demonstrated.
- A CQMOM working group was formed. The CQMOM algorithm was investigated and a first-cut implementation into Arches was completed. Demonstration of the CQMOM algorithm in Arches using three internal coordinates was shown.
- Coal particle physics models in the Arches code were reviewed, recoded, and re-implemented in the Arches algorithm to provide better reliability, stability, and tighter gas-particle coupling.
- Time-scale separation coal models appropriate for the Arches LES context were developed and deployed in the VUQ simulation.
- A coal soot model developed and demonstrated in Arches
- Two devolatilization models were developed and demonstrated in Arches.
- A comprehensive char oxidation model was developed and a parametric sensitivity was performed

Particle Combustion

The physics team's particle combustion research had three main thrusts: coal devolatilization, char conversion, and soot formation. The Chemical Percolation Devolatilization (CPD code) for coal devolatilization was translated into Matlab by the BYU team, put into a form more amenable for uncertainty quantification, and disseminated to the University of Utah and UC Berkeley. Estimates of bounds for input parameters were determined to permit the uncertainty quantification. The CPD model was used to inform simpler devolatilization models for use in the large-scale simulations. Simple coal devolatilization models studied included the one-step Biagini and Tognotti (BT) model where the ultimate yield changed as a function of temperature, the one-step Yamamoto model where the effective activation energy changed as a function of conversion, and the Modified Two-Step (RF) model where the effective activation energy of each step changed with conversion.

The char conversion kinetics (CCK) code for char oxidation was translated into Matlab and put into a form more amenable to uncertainty quantification, along with realistic bounds for code parameters. In addition, a short series of experiments was begun that will address the enormous uncertainty of coal-specific kinetic parameters. Several methods of uncertainty quantification for large models were explored. The CCK model and a simpler model were used to analyze some recent data on coal gasification by CO_2 and H_2O, resulting in a paper published in *Energy & Fuels* (see publication list below).

Soot formation physics based on the Brown and Fletcher model were incorporated into the Arches LES software, and preliminary calculations were initiated. Much of the year-one effort consisted of identifying soot models for use with coal combustion, and implementing those models into Arches, with a view towards running LES simulations in year-two. The physics and kinetics of the model were reviewed as relevant to the current research in soot formation. Areas of model improvement also were identified that are specific to oxy-coal combustion. Specifically, the effects of CO_2 and H_2O gasification, which are expected to play a more significant role at the higher temperatures and higher CO_2, H_2O concentrations present under oxy-fired conditions.

<u>Radiation</u>

The Reverse Monte Carlo Ray Tracing (RMCRT) based radiometer capability was extracted from the core RMCRT code and re-implemented to be used in conjunction with the Discrete Ordinate Method (DOM). Users running production cases need the ability to directly compare RMCRT:radiometer point flux calculations against experimental heat flux sensors while solving the Radiative Transport Equations with DOM. Scaling/series studies of the radiometer code showed that the cost of computations is a tiny fraction of the overall cost. By far the dominant cost is the all-to-all communication, which is also true for the RMCRT code. Three techniques for reducing the communication costs are being investigated. The first method involves casting the communicated variables as floats instead of doubles. In theory this should reduce the cost by a factor of 2 and preliminary scaling results have shown significant savings. The second method involves scheduling the radiometer tasks only on the time-steps when radiometer calculations are performed or “Temporal Task Scheduling.” The final method schedules tasks to be executed only on the patches that contain “radiometer” cells or “Spatial Task Scheduling.” Traditionally, component level tasks are scheduled to execute on every time-step and every patch. This is clearly not necessary for the radiometer tasks. All of the techniques have been implemented and are under testing.

A strong scalability study was performed on Titan using the MPI/ threaded scheduler and the adaptive mesh capabilities of the RMCRT algorithm for solving the RTE. Good strong scaling characteristics were shown up to 262K cores on three problem sizes. The grid consisted of 2 levels with a refinement ratio of 2. The total number of cells on the highest resolved level was 128^3, 256^3 and 512^3. The standard Burns & Christon benchmark test problem, using 100 rays per cell was solved. Infrastructure issues limiting RMCRT:CPU runs have been resolved opening the door for spatially scheduled ray tracing.

RMCRT has for some time been a stand-alone component of Uintah. Work has been done to modify RMCRT so that it can work with the CFD component, ARCHES. Now RMCRT can be used in multi-physics computations. RMCRT is not as efficient as DOM at the same level of accuracy. To make RMCRT competitive with DOM, stochastic quadrature sampling schemes have been introduced into RMCT in combination with ray tracing approximations. These methods allow RMCRT to surpass DOM in accuracy for the same computational cost. This analysis is based off

Burns Christen benchmarks and an analytical solution to the RTE for a sinusoidal flame-mimicking system developed under this project.

The Discrete Ordinance method for solving the Radiative Transport Equation (RTE) has been expanded to include reflections from walls, as well as scattering effects due to particles dispersed in the flow field. Scattering is an important phenomenon in coal systems since the scattering coefficients are often 20% larger than the absorption coefficients. The scattering coefficients are being computed based on Mie theory with a phase function to account for forward and back scattering effects. The Discrete Ordinance model handling of particle interactions has been improved by using the temperature and intensity attenuation based on each particle environment. The DOM linear solve convergence rate has been improved by storing the intensity fields from the previous intensity solves.

<u>LES Environment</u>
Within Aches, an existing implementation of the coal models was used as a starting point in modeling the coal physics of these systems. These models include acceleration of the particles due to drag, devolatilization, char-oxidation, and energy exchange due to radiation, convection, phase-change, and reaction. Various changes were made to the coal models during the first year. The changes were primarily related to initial implementation issues, the addition of sub-grid modeling for the coal models within the large-eddy simulation framework, a method for dealing with small particle weights and the addition of two new coal devolatilization models, and to couple tighter the solid and coal phase.

The initial implementation of the models contained an error in lagging the update of the coal model rates by one time-step. This was corrected, and the models were analyzed and modified in order to speed up the compute times for each of the models. Additionally, model level verification was performed by ensuring that the models behaved qualitatively correct and compared well to an independently developed, stand-alone small-scale code over a range of expected conditions.

Large-eddy simulations resolve the large energy containing scales of the systems and rely on sub-grid modeling to represent the effect of the unresolved scales. In order to march in time at the time scales dictated by the CFL number, it is appropriate to add sub-grid effects to the coal models. This was done by limiting the particle rate to be consistent with the time scales of the fluid, in the case of unresolved time scales for the coal models.

<u>Multi-Phase Flow</u>
The conditional quadrature method of moments (CQMOM) has been implemented into the Arches codebase with both first and pseudo-second order convective schemes. The method has been used to show a two-dimensional case with particle trajectory crossing that the current DQMOM is incapable of resolving. A basic case has also been demonstrated with particle trajectory crossing in three dimensions with no particle-fluid interactions. Wall interactions with a restitution coefficient have been added to the code and demonstrated in a two-dimensional example. CQMOM source terms for inert particle-fluid interactions to model the drag effects have been added to the code, and simulations of inert particles in a coaxial jet flow case have been run.

During the first year of this project, we have implemented a Lagrangian particle transport framework in Uintah. This framework is component-agnostic - any Uintah component is capable of leveraging its capabilities. As an example, ARCHES currently support basic Lagrangian transport. Wasatch, on the other hand, has leveraged all capabilities of this Lagrangian framework. Here is a summary of the supported models:

1. Particle momentum, particle size, particle mass
2. Basic drag and force models for Lagrangian particles
3. Two-way coupling between particles and fluid transport
4. Support wall boundary conditions for particles: Fully elastic, partially elastic, and Inelastic (sticky)
5. Support particle injection through arbitrary boundaries. Particles are seeded using a uniformly random distribution at the injection site
6. Initialize arbitrary shapes with particles
7. User-friendly input interface for setting up particle transport

Validation, Verification and Uncertainty Quantification

Over the first year of CCMSC, the Validation, Verification and Uncertainty Quantification group has focused effort on four projects of primary interest. The first priority was to provide tools and support as well as execution for the first-year full-hierarchy validation. Our second priority was the design of an experimental campaign based on the L-1500 device specifically for comparison with simulations (effort based on cost-sharing funds). The third priority involved a four-university collaboration on the development of a scale-bridging devolatilization model (jointly between predictive science & V/UQ teams). And finally, it has been our pleasure to include Prof. Dongbin Xiu into the center (effort based on non-mandatory cost-sharing funds).

Our consulting technical advisor, Dr. Jerome Sacks, was apprised at regular intervals of almost all of these activities. His contributions have indeed proven valuable by offering advice during the big-picture planning stages and even down to the technical level. Our verification effort has relied heavily on pre-existing checks in our regression-testing suite, while all new coal models (devolatilization, char oxidation, heat capacity, etc.) developed into the Arches framework were verified against external codes and some of them were shown to reproduce relevant quantities from the original publications. In spite of these accomplishments, our verification procedure is not yet comprehensive of our entire code and progress will continue in the remaining years of the program. Additionally, plans have been started for deep dives to be held for V/UQ (as well as particle physics with multi-phase flow).

Our general year-to-year plan has been to perform model development and validation for the fundamental physics on a one-year cycle, and to use the previous year's model in the multi-physics simulations that bridge up scale to the prediction. However, this is not possible in the first year where there is not an acceptable pre-existing model. It was projected early in the year that we would accelerate the development of the scale-bridging devolatilization model in order to incorporate it into the first-year multi-physics simulations.

Progress has been focused within the study of model-form uncertainty and its applications towards the creation of a coal devolatilization scale-bridging model. Significant effort was put into the buildup of foundational knowledge in Bayesian probabilistic based methods of approaching model-form uncertainty and understanding caveats of implementing such methods. Comprising this effort was literature review, application exploration through a pedagogical example, and consideration of novel uses of the methods. This research effort then transitioned into defining the idea of credible simulation design for scale-bridging models. Scale-bridging models are models created to capture specific characteristics of high-fidelity physics models, while functioning on different temporal and/or spacial scales. Utilizing the ideas of model-form uncertainty and following the process for credible simulation design, a coal devolatilization scale-bridging model was created. This scale-bridging model was based upon a rigorous chemical kinetics model the Chemical Percolation Devolatilization model. By quantifying the effect of propagating the uncertainty contained within parameters of the more rigorous model-form through the model to specific outputs, the resulting uncertainty was utilized as constraints for the creation of a reduced model-form. Various forms for the reduced model-form were explored and iterated upon. Discrepancy and validity of the final reduced model-form were assessed through various validation techniques. The discrepancy was then used to inform model-form improvements, effectively improving model credibility. This reduced model-form was then implemented within the large-scale CFD code and one of its parameters utilized throughout the hierarchal validation and uncertainty quantification study.

Initial effort was made to consolidate multiple experimental databases into a single database. This effort has involved collaborators at the Universite Libre de Bruxelles, Belgium who have traveled to the International Flame Research Foundation test facility (5 MWth oxy-gas furnace and 1.5 MWth oxy-coal furnace) in Italy. Simultaneously, Tom Fletcher at BYU assisted in the inclusion of experimental databases from Sandia National Laboratories and BYU. Michael Frenklach at UC-Berkeley provided the PRIME database tools to allow access and querying of the database. This effort will continue into the remaining years of the Center.

<u>Oxy-Fuel Combustor (OFC)</u>

The oxy-fuel combustor has played an important role in the center during the first year. The OFC has been used extensively as a verification test-bed, a bench-mark case for improving the efficiency of ARCHES, and a full V/UQ analysis has been performed on the simulation of the combustor. Due to the relatively small size and computational costs for this case, the OFC has been an ideal case for testing new models, such as the modified Biaghini/Tognotti devolatilization model, and new soot formation models. The OFC simulation has been used to help in understanding sensitivity to both scenario and model parameters, because of its short compute times. The V/UQ analysis conducted on the OFC included 12 simulations, which investigated the systems response to changes in wall thermal conductivity, and one of our devolatilization model parameters. The cases showed consistency with experimental data given a small range in thermal conductivity of approximately 0.4-0.5 W/(mK).

<u>IFRF Fo.Sper Furnace</u>

The 5 MWt IFRF furnace can be air- or oxy-fired with either coal or natural gas. The V/UQ effort in year one was focused exclusively on the natural gas data. This furnace is a brick in the hierarchy because of its scale and the fact that there are no particles associated with the natural

gas flame. This case allows inclusion of many of the physics present in a full-scale furnace without the complexity of particle transport and combustion. By eliminating all the parameters associated with the models required for simulating coal particles, different areas of the I/U map can explored that might have an impact on heat flux, the overall QOI.

This effort was divided into four specific tasks, each of which has been completed this first year of the program. The first task, due to the limited literature data available, was to obtain additional experimental details needed to reduce uncertainties associated with the simulations. The second task was to perform simulations of fluid flow through the complex geometry of the TEA-C industrial burner used in the experimental campaign. This required a separate STAR-CCM + calculation and a handoff file to interface results to the ARCHES LES. The third task involved a python script to coarsen the handoff boundary condition from the STAR-CCM+ results, so they could be used directly as boundary conditions in the LES. The fourth task was to perform scoping simulations of the 2 m x 2 m x 6.75 m IFRF furnace. A mesh size of 10 mm was chosen for the initial simulations, resulting in a computational mesh of 27 million cells.

L1500 Furnace

The L1500 Furnace is a single-burner 1.5 MWt furnace housed at the University of Utah. The scale of the furnace is such that it allows for detailed measurement without the burden of burner-burner interaction. It was determined in the course of the first year that the data on which the VUQ activities were planned were not appropriate for the VUQ study. However, simulations of the L1500 using the previously collected data served useful to test parametric sensitivity. In particular, sensitivity to the burner swirl was investigated.

Simulations of the L1500 included a detailed representation of the burner geometry using a coupled STAR-CCM+/Arches simulation. Near-burner flow was resolved using the unstructured STAR-CCM+ code that then served as a boundary condition to the full-scale Arches LES simulation. The burner swirl was adjusted from 100% to 0%. Both simulations were run on the Vulcan(LLNL) cluster on 8192 cores. The high-swirl simulation required 4.1M CPU hours and the no-swirl required 0.9M CPU hours to complete.

In addition to designing and performing the simulations, the experimental and simulation teams worked closely to plan and execute the new experimental campaign carried out on the L1500 furnace. Several new data points were collected including mass flows, temperatures, tube heat balances, heat fluxes, and gas compositions. Two operating conditions were explored, include running with 100% and 0% swirl. Simulations were performed exploring the sensitivity to the wall thickness and the devolatilization parameter, V_hiT. These simulations improved upon geometric representation of the L1500 within the LES simulations.

Boiler Simulator Facility (BSF)

Alstom's boiler simulation facility has been useful in demonstrating the capabilities of CCMSC in accurately modeling pulverized coal boilers. A V/UQ analysis of the BSF was performed, involving 8 simulations where the wall thermal conductivity, and one of the devolatilization model parameters were varied. Consistency between 22 heat flux measurements, 104 temperature measurements, and 38 oxygen measurements was found given a range in wall thermal conductivity of 4-5 W/(mK). The devolatilization model parameter showed very little effect in the quantities of interest. A large demonstration run of the BSF was performed in an effort to show sensitivity to mesh resolution, and demonstrate the scalability of the Uintah framework and Arches as we move our computations toward exa-scale computing. The simulation represented the physical system with 135 million cells each with a volume of 1 cm^3, and used 22.7 million cpu hours to compute an initial 5 seconds of simulated time. The case helped to identify several memory issues involved in these large computations.

Experimental Campaign

Much of our experimental effort has been focused on the 1.5 MW pulverized coal test facility (L1500) in order to provide the L1500 simulation team with the best possible data set for the V&V/UQ effort. First, we analyzed existing L1500 data to identify potential improvements that could be made in the quality of data collected. Many of these improvements were implemented for the L1500 test campaign that was scheduled for this year. The facility was visited multiple times by the V&V/UQ team (both experimental and simulation team members) to discuss many of the logistics with the supervising test engineer. Many meetings were held regarding the planning of the L1500 test campaign and simulations, some involving the experimental team, some with the simulation team, and some with both teams. The experimental team also met multiple times to discuss how get the optimal amount of data while still apportioning limited resources, such as space, time and money, amongst everyone.

A single operating condition of interest was identified with the intent to take data using two distinct burner swirl settings at this condition. The campaign was initially planned for the last week in January 2015 but due to delays involving the furnace, it was performed during February 23-27, 2015. Prior to the test campaign, many modifications and upgrades to the furnace were made in an attempt to improve data quality. A refurbished feed system for the pulverized coal stream was completed. This new system provides coal at a much steadier rate than the previous one. This was a high priority for both teams, since the coal feed rate is of primary importance in both experimental and simulation work. New equipment was added to the reactor to provide more insight into the primary quantity of interest: heat flux. New water wall cooling tubes were added in the first four sections of the reactor. Digital thermocouples and flow meters were used to monitor heat flux in three narrow-angle radiometer probes. Quartz window ports were built for the optical access of a high-speed infrared (IR) camera. The sampling ports were modified to ensure sealed sampling.

Before the campaign began, reflectivity from deposits collected from the L1500 interior walls were measured. These results were then converted into a total emissivity (a value the simulation team needed as an input). Thermal diffusivity measurements were also taken on the interior of the furnace walls before the test campaign. The ash/slag deposits at various locations in the furnace are very different in both color and structure, and we desired to examine the differences in these deposits by creating a “map” of diffusivity measurements.

During the campaign, infrared heat flux measurements were taken by focusing on various surfaces inside the furnace with the high-speed IR camera. This data is currently being scrutinized in detail for the insights it can give into the radiation from the various surfaces within the reactor (i.e., cooling tubes coated with deposit, hot refractory, the cool surface of a heat flux gauge face). Heat flux data was also taken by radiometers, but since they were point measurements, it is more difficult to estimate which type of surface they were averaging over.

We also worked on the identification and organization of existing OFC experimental data to be included on the CCMSC wiki, and we will continue to do so to ensure that data set is complete.

Two of the comments from the annual peer review indicated an interest in seeing “smaller scale/reduced physics experiments,” particularly ones that would yield data regarding wall deposition. We are planning to implement these suggestions. We plan to meet with the simulation team in the coming weeks in order to design experiments that would best suit the needs of the V&V/UQ effort. An additional comment asked for a more in-depth look into the properties of the deposit and how those properties affect the wall heating model. We have already begun characterizing some properties of the deposits inside the L1500 furnace with our reflectivity and thermal diffusivity measurements. We will continue with this effort, as we plan to take more measurements of the reactor walls to examine conditions post-campaign when it is available to be opened up in May. Thus, we will have a before and after comparison of the wall deposits in the reactor.

PROJECT PLAN / GOALS

<u>Outreach and Education</u>

As noted in the introduction, several internships are planned for the coming year and a few others are considering their opportunities. One staff research experience is planned. The jointly designed, first-year, graduate course discussed earlier will be taught. Participant evaluations will be requested and considered for potential improvements to delivery and effectiveness of content design. Additionally, a workshop is being planned that will address particle transport and reactions and will involve all of the six PSAAP centers. We will seek to augment the workshop with a few leading scientists in the field. As the workshop details are defined, announcements will be placed on the Center’s website and be distributed to our partner labs through our TST and AST representatives. We will host the semi-annual TST and annual reviews with broad participation of researchers and students of the Center.

<u>Exascale Scalable Infrastructure: TASC/EDSL, DMVA, Exascale Runtime</u>

<u>RMCRT</u>
We will address known performance issues in the RMCRT:Radiometer and utilize it in production size simulations. We will continue developing and testing the CUDA:GPU implementations of RMCRT:single level and RMCRT:Adaptive mesh algorithms. We seek to demonstrate strong and weak scaling characteristics, and examine accuracy and cost of RMCRT with its different modes of operation versus Discrete Ordinates Method. Plans are to integrate Kokkos into key portions of Uintah infrastructure. We will continue to work with the Sandia team on performance and potential GPU tasks within the MiniAero component.

During year two, the goal is to leverage the insights gained from these experiments to improve upon MIC-based performance of this algorithm. Areas of particular interest to be addressed are effective memory management and vectorization. These are believed to be two key areas hindering performance of the current algorithm. An additional goal is to conduct initial scalability studies using TACC's Stampede to better understand how this algorithm scales on the Intel MIC Architecture.

Software Engineering

We will release two more software updates (one every six months). Plans include to transition the software development environment for source management to the GitHub platform. Debugging will be tackled for the interaction between Uintah's Pthread and hypre's OpenMP threads for the threaded scheduler. (At larger thread counts, we are experiencing stalling in the threaded scheduler.) The team will investigate patch coalescing when handing off the linear system to Hypre, since better performance is achieved with larger and fewer patches per node when using OpenMP threads. We will investigate the sweep method as an alternative solver implementations for the Discrete Ordinates method. (This should allow us to eliminate an expensive call to the hypre library.) The results will be used to pursue better scalability of the ARCHES component when solving the Clean Coal Problem.

We will improve performance on multithreaded architectures and on GPU. We will re-examine performance on Xeon Phi when Knights Landing hardware becomes generally available. (The DSL will also continue adding functionality as required.) This summer, variable density pressure treatment for low-mach solvers will be examined using projection schemes as part of our interaction with SNL. Work will continue to support ARCHES transition to the DSL and continue to forward-deploy technologies in Wasatch to harden them prior to ARCHES adoption. The team will seek to demonstrate scalable GPU performance on a fluid-flow problem.

In Situ Visualization

Plans are to fully integrate Uintah and VisIt so that it scales appropriately. Design work will allow all simulation variables to be exposed to VisIt (not just the variables that would be saved to disk) and allow for all variables to be exposed regardless of the level or patch and allow the same for run time data. Plans also include integration of Volume Rendering of the compositing library from the TOD-Tree paper into VisIt. We will work to improve the GUI for Volume rendering in VisIt - especially the transfer function editor to make it more user friendly and address the application scientists' needs.

Physics: Radiation, Particle Combustion, Multiphase Flow, LES Environment

In the coming year additional simple devolatilization models will be examined, including the distributed activation energy model and possible improvements using the RF and BT models. This will primarily encompass a better agreement between the initial rate predictions and the long-time "equilibrium" predictions. Data sets will be identified to perform additional uncertainty quantification. Particle swelling during devolatilization will also be addressed during this year to incorporate either in the devolatilization kinetics or the char kinetics. A journal paper is currently being written by the BYU team to discuss this research effort.

The char conversion codes will be subjected to a thorough sensitivity analysis. The analysis will indicate where the models need the most attention. We anticipate that coal-specific kinetics, coal swelling, and particle mode-of-burning will be high priorities. The most important parameters will be updated for a more comprehensive and accurate treatment of the involved physics, and the char conversion model, as a whole, will be modified for uncertainty quantification. Experiments will also be completed for the Utah Sufco coal and they will be incorporated with a collected body of relevant literature data for use in uncertainty and consistency analysis.

Models for secondary coal tar pyrolysis and soot formation will be improved, making these models compatible with mass and energy balance closure in Arches simulations. Additional aspects will be added to the code to simulate soot gasification by H_2O and CO_2 and soot formation from light gases. A semi-implicit capability will be created for calculation of soot formation in small subgrid time-steps, therefore allowing larger time-steps in the fluid dynamics simulation. Simulations of coal including soot formation in which tar, soot number density, and soot mass are transported currently are being performed in the year-two effort. We expect to test the sensitivity of simulations to the level of detail included in the soot model description.

The soot model now implemented in the one-dimensional turbulence code will be refined and models for devolatilization, tar and soot into Wasatch will be inserted. We plan to develop high-fidelity wall-particle interaction models for Wasatch and investigate particle fragmentation effects.

Plans are to address known performance issues in the RMCRT:Radiometer and utilize it in production size simulations. Work will continue on developing and testing the CUDA:GPU implementations of RMCRT:single level and RMCRT:Adaptive mesh algorithms. We will demonstrate strong and weak scaling characteristics. We will examine accuracy and cost of RMCRT with it's different modes of operation versus DOM.

Validation, Verification and Uncertainty Quantification

Our general objectives are the pursuit of algorithmic, theoretical, and computational developments of the mathematical framework that forms the foundation of the deterministic, Bound-to-Bound DataCollaboration (B2B-DC), and application of B2B-DC to the UQ/V&V needs of the PSAAP II project. The year-two, specific objectives of the UCB Team are:

1. To work on reduced-order models for coal devolatilization with the Utah and BYU teams, developing relationships for uncertainties of the detailed and reduced models.
2. To explore a hybrid approach to UQ by extending the B2B-DC perspective to calculate posterior distributions for predictions (parameters, key quantities of the overarching boiler performance problem). This will be accomplished by uniform sampling of sets described by nonconvex inequalities.

3. To develop new consistency measure, which mathematically aims to directly address (through an L1 cost) the question of which (and possibly only a few) of the dataset units are most responsible for inconsistency.
4. To work with the Utah team on collection of experimental data on coal devolalitalization. We will be leveraging the technology developed at UCB called Process Informatics Modeling Environment (PrIME) for data organization, archival, and seamless linking to UQ analysis tools, such as B2B-DC.
5. To release annual update to B2B-DC analysis toolbox, with improved functionality geared specifically towards the Utah PSAAP problem. This will include testing using the new Matlab unit test environment.

We will continue V&V/UQ evaluations of L1500 experimental and simulation data consistency from year-one and identify modifications required to experimental methods, data collected, equipment needed, etc. to enhance quality of data for target QOI in planned L1500 experiments for year-two. We will pursue additional experimental work at smaller scale, with particular emphasis on study of deposition and wall deposit properties.

Lessons learned from our validation with uncertainty quantification experience include:

- Move modeling of deposition / ash transformation
- Less emphasis on reaction chemistry
- More sensitivity studies (I/U map)
- Seek more computational resources
- Speed up production code
- Improve memory and I/O management

IFRF work is planned to:

- Perform additional sensitivity analysis to determine one model and one scenario parameter for V/UQ analysis. Scenario parameter is most likely the wall thermal resistance.
- Complete first cycle of two-parameter V/UQ analysis using Rate-Controlled Constrained Equilibrium (RCCE) as the reaction/combustion model
- Develop CO model, implement in ARCHES

L1500 work is planned to:

- Refine instrument model for radiometer to include a correction for wall temperature
- Develop instrument model for measuring gas composition
- Complete first cycle of V/UQ analysis for 0% swirl case - this goal was started during year one
- Perform first cycle of V/UQ analysis for 100 % swirl case
- Develop improved char oxidation model, implement in ARCHES

PUBLICATIONS AND PRESENTATIONS

Bremer, Peer-Timo, Bernd Mohr, Valerio Pascucci and Martin Schulz. *VPA: First Workshop on Visual Performance Analysis.* Supercomputing 2014. New Orleans, LA. November 16-21, 2014.

Cyr, E. C., E. Phipps, M. A. Heroux, J. Brown, E. T. Coon, R.C. Kirby, V. Tzanio, J.C. Sutherland and R, R. Trott. *Algorithms and Abstractions for Assembly in PDE Codes: Workshop Report* - SAND2015-1379. Albuquerque, NM. February 2015.

Diaz-Ibarra, Oscar H., Jennifer Spinti, Andrew Fry, Jeremy N. Thornock, Benjamin Issac, Derek Harris, Michal Hradisky, Sean Smith, Eric Eddings and Philip J. Smith. *A validation/uncertainty quantification (V/UQ) analysis for a 1.5 MW L1500 furnace*. Paper to be presented at the AFRC 2015 Industrial Combustion Symposium. September 2015.

Draper, Teri, Pal Toth, Terry A. Ring and Eric G. Eddings. *Optical heat flux and temperature measurements on a 100 kW oxy-fuel combustor*. 7th European Combustion Meeting (ECM 2015), Budapest, Hungary. March 30 - April 2, 2015.

Earl, Christopher, Matthew Might, Abhishek Bagusetty, and James C. Sutherland. *Nebo: An efficient, parallel, and portable domain-specific language for numerically solving partial differential equations.* Journal of Systems and Software. Accepted 2015.

Frenklach, Michael, Andy Packard, Jerome Sacks, R. Paulo and G. Garcia-Donato. *Comparison of statistic and deterministic frameworks of uncertainty quantification*. Journal of Uncertainty Quantification. In review 2015.

Frenklach, Michael and Andy Packard. *Bound-to-Bound Data Collaboration*. Columbia University, Mechanical Engineering, invited seminar. October 10, 2014.

Frenklach, Michael, Andy Packard, Jerome Sacks, R. Paulo and G. Garcia-Donato. *Mini-Symposium on Uncertainty Quantification in Computational Combustion*. 15th International Conference on Numerical Combustion. Avignon, France. April 19-22, 2015.

Grosset, Pascal, Manasa Prasad, Cameron Christensen, Aaron Knoll and Charles Hansen. *TOD-Tree: Task-Overlapped Direct send Tree Image Compositing for Hybrid MPI Parallelism.* EuroGraphics Parallel Graphics and Visualization Symposium. May 2015.

Holland, Troy M. and Thomas H. Fletcher. *Coal Particle Combustion*. Poster presented at the Stewardship Science Academic Programs (SSAP) Symposium. Santa Fe, New Mexico. March 11-12, 2015.

Holmen, John, Alan Humphrey, and Martin Berzins. *Beginnings of a Scalability Centric Xeon Phi Port. High Performance Parallelism Pearls Vol 2.* (accepted). Ed J. Reinders et al. 2015.

Humphrey, Alan, Todd Harman, Martin Berzins and Philip J. Smith, *Scalable Algorithm for*

Radiative Heat Transfer Using Reverse Monte Carlo Ray Tracing, ISC High Performance. Frankfurt, Germany. 2015 (accepted).

Humphrey, Alan. *Radiation Modeling Using Reverse Monte Carlo Ray Tracing Within the Uintah Framework.* SIAM Conference on Computational Science and Engineering. March 16, 2015.

Josephson, Alexander J,. Benjamin J. Isaac, David O. Lignell and Thomas H. Fletcher.Kim. *Large eddy simulation of an oxy-coal combustor.* 9th U.S. National Combustion Meeting. Cincinnati, OH. May 17-20, 2015.

Mark and Charles Hansen. *Surface Flow Visualization using the Closest Point Embedding.* 2015 IEEE Pacific Visualization Symposium. April 2015.

Kim, Mark and Charles Hansen. *GPU Surface Extraction with the Closest Point Embedding*. Proceedings of IS&T/SPIE Visualization and Data Analysis. February 2015.

Kumar, Sidharth, John Edwards, Peer-Timo Bremmer, Aaron Knoll, Cameron Christensen, Venkatram Vishwanath, Philip Carns, John A. Schmidt and Valerio Pascucci. *Efficient I/O and Storage of Adaptive Resolution Data*. Supercomputing 2014. New Orleans, LA. November 16-21, 2014.

Landge, Aaditya G., Valerio Pascucci, Attila Gyulassy, Janin C. Bennett, Hermanth Kolla, Jacqueline Chen and Peer-Timo Bremer. *In-situ Feature Extraction of Large Scale Combustion Simulations Using Segmented Merge Trees.* Supercomputing 2014. New Orleans, LA. November 16-21, 2014

Lewis, A. D., Troy M. Holland, N. R. Marchant, Eric G. Fletcher, D. J. Henley, E. G. Fuller, and Thomas H. Fletcher. *Steam Gasification Rates of Three Bituminous Coal Chars in an Entrained-Flow Reactor at Pressurized Conditions*. Energy and Fuels, **29**, 1479−1493 (2015).

Lignell, David, Alex Josephson, Benjamin Isaac, and Thomas H. Fletcher. *Large eddy simulation of soot formation in an oxy-coal combustor.* 15th International Conference on Numerical Combustion. Avignon, France. April 19-22, 2015.

Richards, Andrew and Thomas H. Fletcher. *A modified two-step model of devolatilization.* 15th International Conference on Numerical Combustion. Avignon, France. April 19-22, 2015.

Saad, Tony and James C. Sutherland. *A High-Order Explicit Projection Method for Low-Mach Reacting Flows.* 15th International Conference on Numerical Combustion. Avignon, France. April 19-22, 2015.

Saad, Tony, Abhishek Bagusetty, and James C. Sutherland. *Wasatch: A CPU/GPU-Ready Multiphysics Code Using a Domain Specific Language*. SIAM Computational Science and Engineering Conference. Salt Lake City, Utah. March 14-18, 2015.

Saad, Tony, Christopher Earl, Abhishek Bagusetty, Mathew Might and James C. Sutherland.

Uintah/Wasatch: Addressing Multiphsyics Complexity in a High-Performance Computing Environment. SIAM Computational Science and Engineering Conference. Salt Lake City, UT. March 14-18, 2015.

Sutherland, James C. *Efficient Abstractions for Exascale Software Design.* SIAM Conference on Computational Science and Engineering. Salt Lake City, UT. March 14-18, 2015.

Sutherland, James C. and Abhishek Bagusetty. *DAG-Based Efficient Scalable and Portable PDE Software.* SIAM Computational Science and Engineering Conference. Salt Lake City, UT. March 14-18, 2015.

Sutherland, James C. *Multiscale Modeling on Multiscale Computers*. 5th International Conference on Coupled Thermo-Hydro-Mechanical-Chemical (THMC) Processes in Geosystems. Salt Lake City, UT: Keynote Address. February 25-27, 2015.

Sutherland, James C., Matthew Might, Christopher Earl, Tony Saad and & Abhishek Bagusetty. *Flexible, Efficient Abstractions for High Performance Computation on Current and Emerging Architectures*. Algorithms and Abstractions for Assembly in PDE Codes Workshop. Albuquerque, NM. June 2014