

Quarterly Progress Report
Covering Period: 1/1/2012 - 3/31/2012
Date of Report: 4/15/2012

Award Number: DE-EE0004449

Project Title: Unconventional and Renewable Energy Research Utilizing Advanced Computer Simulations (UT)

Project Period: 09/23/2010 - 03/22/2013

Recipient Organization: SCI Institute, University of Utah
72 south Campus Drive, Room 3750
Salt Lake City, Utah 84112

Partners: Chris R. Johnson, PhD
Martin Berzins, PhD
Ross Whitaker, PhD
Mike Kirby, PhD
Cost Sharing Partner: University of Utah

Technical Contact: Chris R. Johnson, PhD
72 south Campus Drive, Room 3750
Salt Lake City, Utah 84112
(801) 585-1867
crj@sci.utah.edu

Business Contact: Edward Black
72 south Campus Drive, Room 3750
Salt Lake City, Utah 84112
(801) 585-1867
ed@sci.utah.edu

DOE Project Officer: Charles Alsup
304-285-5432
calsup@netl.doe.gov

DOE HQ Contact: Bob Gemmer
202-586-5885
Bob.Gemmer@ee.doe.gov

DOE Contract Specialist: Jacquelyn Wilson
304-284-4135
Jacquelyn.Wilson@netl.doe.gov

Project Objective:

The ability to develop science-based and validated computational tools to simulate and facilitate the development of clean, highly efficient energy systems of the future requires innovation in several key computational science technologies, including scientific data management, scientific visualization, scientific software environments, and scientific computing. The overall objective of this work is to leverage our expertise and experience in both scientific visualization and complex science-based simulations toward the accurate and robust simulation of science-based phenomena in the area of unconventional and renewable energy research. This work is aimed at garnering a better understanding of science-based phenomena in energy research and also the advancement of the Uintah software system. The Uintah software system accommodates the massive amounts of data and advanced algorithmic, software, and hardware technologies required to deal with the enormity and complexity of the simulation data in this area of research. To accomplish these goals, we are creating new numerical and visualization techniques needed to assess the uncertainty of the simulation, extend the Uintah scientific problem-solving environment for large-scale simulation of science-based systems, and integrate and extend the data provenance infrastructure of Uintah to systematically capture provenance information and track simulation parameter studies.

Background:

Science-based development of clean and efficient energy systems often involves modeling and simulations of fluid flows, chemical reactions and mechanical properties within heterogeneous media. As part of our DOE-funded (1997-2009) Center for Simulation of Accidental Fires and Explosions (C-SAFE), we created the Uintah scientific problem-solving environment. Uintah is a parallel software environment for solving large-scale computational mechanics and fluid dynamics systems, and has particular strengths when dealing with systems that require large deformations, fire simulation, and fluid-structure interactions. Uintah, general-purpose, fluid-structure interaction code has been used to characterize a wide array of physical systems and processes encompassing a wide range of time and length scales - from microseconds and microns to minutes and meters. Complex simulations require both immense computational power and complex software. Typical

simulations include solvers for structural mechanics, fluids, chemical reactions, and material models, which are efficiently integrated to achieve the scalability required to perform the simulations. Uintah scales to cores by using a novel asynchronous task-based approach for challenging AMR applications. Novel parallel computing algorithms, on both CPUs and GPUs, are needed when simulating large-scale complex science-based energy systems. In moving beyond petascale, it will be necessary to make use of GPU-like architectures as the ongoing convergence between GPUs and multi-core CPUs continues. Task-based codes like Uintah are very well placed to exploit such architectures.

The challenge for finite element type simulations is that the memory access patterns are not well suited for the cache coherency required for efficient operations on streaming architectures. The problem becomes worse for sparse systems associated with large simulations. Thus, the performance improvements over CPU implementations have been limited. An alternative is to take advantage of the geometric configuration of unstructured meshes, and to invent compact, efficient data structures that allow SIMD processing of individual cells and subsequent SIMD assembly of cell computations and mapping onto global degrees of freedom in the solution. The problem becomes more challenging for algorithms that are effective on multi-GPU clusters, such as the NVIDIA cluster at the SCI Institute. We anticipate the need for hierarchical domain decompositions that provide sufficient computational density and efficient communication. This work will pursue GPU and GPU-cluster based algorithms for numerical simulations of combustion using both generic linear solvers and specialized solutions that directly map unstructured and structured domains onto streaming architectures.

The system must also provide data visualization capabilities that allow interaction and analysis of the simulated data. The SCI Institute is an international leader in scientific visualization research. Additionally we are exploring the use of higher fidelity visualization with methods based on the use of high-order mesh elements.

With large computational simulations there is substantial uncertainty inherent in any prediction of science-based systems. A number of factors contribute to uncertainty, including experimental measurements, mathematical formulation, and the way different processes are coupled together in the numerical approach for simulation. Tracking of and analysis of this uncertainty is critical to any work that will truly impact the creation of future energy systems.

Exploration of large-scale scientific systems using computational simulations produces massive amounts of data that must be managed and analyzed. Because of the volume of data manipulated, and the complexity of the simulations and analysis workflows, which are iteratively adjusted as users generate and evaluate hypotheses, it is crucial to maintain detailed provenance (*i.e.*, audit trails or histories) of the derived results. Provenance is necessary to ensure reproducibility as well as enable verification and validation of the simulation codes and results. In order to manage large-scale simulations and the analysis of their results, we will use systems such as the VisTrails software (<http://www.vistrails.org>), an open-source provenance management and scientific workflow system that was designed to support the scientific discovery process, to guide us in building "hooks" into Uintah for provenance systems.

Accomplishments

Uintah Software System

Alstom Boiler Problem This simulation emerges from the work of the Combustion and Reaction Simulations (CRSim) research group led by Dr. Philip J. Smith group at the University of Utah on advanced ultra-supercritical electric power generation boilers. The purpose of these boilers is to transfer thermal energy, generated by fuel oxidation, into steam which is then passed to a turbine that generates electricity at the thermodynamic efficiency of a Rankine steam cycle. The efficiency of the overall process of converting the chemical energy of coal into electricity increases as the temperature and pressure of the superheated steam entering the turbine increases. Creating super-critical steam at 760C has been called “advanced ultra-super critical (AUSC)”. These AUSC boilers can be more than 35% more efficient than today’s U.S. fleet of coal-fired power plants. Oxy-fuel combustion, in which an O₂/CO₂ mixture replaces air, is a promising solution for reducing greenhouse gases and pollutant emissions. Together with recycled flue gases, burning coal with relatively pure oxygen can produce a highly concentrated flue gas stream (up to 95% CO₂), which makes carbon sequestration more economical. However, the presently available cryogenic air separation process required for oxy-combustion results in a reduction of efficiency of 6.4% . Thus, the option of combining AUSC with oxy-combustion will make up for this efficiency loss (due to O₂ separation) and potentially allow for a more efficacious carbon capture, compared with existing coal-fired power plants.

Switching from conventional to oxy-fired coal combustion raises significant challenges. For instance, in employing a strategy to retrofit existing air-fired boilers, we must first arrive at a clearer understanding of how an O₂/CO₂ mixture will affect the kinetics, aerodynamics, and heat transfer of the flame. However, the combustion of pulverized coal is a more complex phenomenon than that of gaseous or liquid fuels. Moreover, performing experimental measurements of pulverized coal flame characteristics is extremely difficult. In this, the development of new combustion furnaces and burners is still empirical and requires a significant amount of time and money.

Dr. Smith’s group uses CFD simulations to model the turbulent flow of combustion in the particulate nature of pulverized coal. Large Eddy Simulation (LES) is used to accurately predict oxy-coal flame characteristics and has proven useful in understanding and engineering the retrofitting process. In predicting turbulent mixing and combustion dynamics, LES has shown excellent results when compared to the common Reynolds-Averaged Navier-Stokes (RANS) method. Although the computational requirements of LES are greater than those of RANS, its superior predictive abilities make it the clear choice for practical combustion applications.

In this project, LES is applied to a semi-industrial scale, tangentially-fired coal boiler (50MW) in oxy-combustion conditions. Arches, a LES tool developed at the University of Utah, is a massively parallel code (shown to scale upwards to 260K core on Titan/XK6) which couples its capabilities with the Direct Quadrature Method of Moments (DQMOM), to model coal particles in turbulent flows in a novel way.

Heat is transferred in the firebox of a coal-fired boiler by radiation. The design of new fireboxes for an AUSC oxy-combustion technology will require accurate radiative heat transfer simulations for environments that have not yet been studied (increased CO₂ concentrations, higher temperatures, different radiative properties for new metal alloys, etc.). Our simulation experience shows that radiative heat-transfer computations take 1/3 to 2/3 of the computational load of a boiler

simulation. By using spectral calculations with complex spectral properties, we can increase the accuracy of this computational load even further.

Professor Smith's group proposes to take advantage of the parallel-computing hardware and the advances in Monte Carlo ray-tracing technology, pioneered in computer visualization tools to use Reverse Monte Carlo Ray Tracing (RMCRT) algorithms to solve for the radiative-flux divergence while accounting, in an efficient manner, for the effects of spectral participating media. RMCRT lends itself to scalable parallelism because the intensities of each ray are mutually exclusive; therefore, multiple rays can be traced simultaneously at any given time step. However, the "all-to-all" nature of radiation requires information about the entire computation to be stored on each computational patch. We will explore new methods for taking advantages of multilevel methods to develop an algorithm that will overcome the detrimental scaling effects of the "all-to-all" nature of radiation.

Carbon Capture Modeling

This effort involves a partnership with a Silicon valley startup, where the University of Utah performs detailed, large-eddy simulations (LES) of crystallization in an industrial reactor, using Jaguar XT5 and the new Titan XK6. The physical properties of interest are the resulting particle-size distribution of the crystals and the relative abundance amongst the multiple polymorphs. Figure 1 shows an example of this carbon-capture simulation. The partnership will simulate and model data provided by the industrial partner in a validation/uncertainty-quantification (V/UQ) analysis. This effort has resulted in a joint identification of exactly how high-performance computing and V/UQ efforts at the university can accelerate deployment of the industrial technology. Although the geometric complexities of these reactors exceed all previous simulations with the university (Uintah/Arches) code, we were able to complete the simulations and demonstrate their utility to the industrial partner. These simulations required coupling between fast time-scale acid-base reactions and slow time-scale precipitation processes.

In addition, this effort has incorporated modeling tools for the calculation of acid/base equilibrium under a variety of constraints in the aqueous phase. The implemented model focused on two components: a tabulation of equilibrium states constrained to a mixture fraction and a precipitation extent of reaction calculated in the LES. The data for the LES tabulation is performed using a commercial acid/base equilibrium code called OLI.

Generic Fluid-Structure Interaction Problems

Solving fluid-structure interactions has long been a feature of Uintah. In doing so, making efficient use of large-scale parallel architectures has proven to be more challenging than we anticipated, in terms of software engineering. We were using Uintah on Jaguar to do very large scale fluid structure calculations and have demonstrated scaling up to 198K cores on Jaguar (XT5), using both an mpi-only and a hybrid threads/mpi scheduler. We moved to the hybrid threads/mpi to help reduce the memory footprint of the global data structures, and to improve its load-balancing and scheduling. The hybrid/threads model has worked well and we have seen improvements in our scaling results for our AMR Fluid-Structure simulations on the XT5 system.

We attribute the performance improvements to the improved communication infrastructure of the XK6.

Algorithm Development

ElVis: A System for the Accurate and Interactive Visualization of High-Order Finite Element Solutions This work presents the Element Visualizer (ElVis), a new scientific visualization system for use with high-order finite element solutions to PDEs in three dimensions. This system is designed to minimize visualization errors of these types of fields by querying the underlying finite element basis functions (*e.g.*, high-order polynomials) directly, leading to pixel-exact representations of solutions and geometry. The system interacts with simulation data through runtime plugins, which only require users to implement a handful of operations fundamental to finite element solvers. The data in turn can be visualized through the use of cut surfaces, contours, isosurfaces, and volume rendering. These visualization algorithms are implemented using the NVIDIA OptiX GPU-based ray-tracing engine, which provides accelerated ray traversal of the high-order geometry, and CUDA, which allows for effective parallel evaluation of the visualization algorithms. The direct interface between ElVis and the underlying data differentiates it from existing visualization tools. Current tools assume the underlying data is composed of linear primitives; high-order data must be interpolated with linear functions as a result. In this work, examples drawn from aerodynamic simulations's high-order discontinuous Galerkin finite element solutions of aerodynamic flows in particular will demonstrate the superiority of ElVis' pixel-exact approach when compared with traditional linear-interpolation methods. Such methods can introduce a number of inaccuracies in the resulting visualization, making it unclear if visual artifacts are genuine to the solution data or if these artifacts are the result of interpolation errors. Linear methods additionally cannot properly visualize curved geometries (elements or boundaries) which can greatly inhibit developer's debugging efforts. Pixel-exact visualization exhibits none of these issues, removing the visualization scheme as a source of uncertainty for engineers using ElVis.

This work was submitted to the IEEE Visualization Conference.

Scientific Visualization

Data Analysis with the Morse-Smale Complex

In many areas, scientists deal with increasingly high-dimensional data sets. An important aspect for these scientists is to gain a qualitative understanding of the process or system from which the data is gathered. Often, both input variables and an outcome are observed and the data can be characterized as a sample from a high-dimensional scalar function. This work presents the R package *msr* for exploratory data analysis of multivariate scalar functions based on the Morse-Smale complex. The Morse-Smale complex provides a topologically meaningful decomposition of the domain. The *msr* package implements a discrete approximation of the Morse-Smale complex for data sets. In previous work this approximation has been exploited for visualization and partition-based regression, which are both supported in the *msr* package. The visualization combines the Morse-Smale complex with dimension-reduction techniques for a visual summary representation that serves as a guide for interactive exploration of the high-dimensional function. In a similar fashion, the regression employs a combination of linear models based on the Morse-Smale decomposition of the domain. This regression approach yields topologically accurate estimates and facilitates interpretation of general trends and statistical comparisons between partitions. In this manner, the *msr* package supports high-dimensional data understanding and exploration through the Morse-Smale complex.

This work was accepted in the Journal of Statistical Software.

Uncertainty Visualization

Entropy for Uncertainty Visualization of Categorical Data Standard deviation and quantiles are the de-facto standard to encode uncertainty in visualizations. However, these statistics are not always appropriate or even feasible. This work discusses the use of entropy, from the field of information theory, as a method for encoding the uncertainty within a categorical data set. For categorical data, statistical moments and quantiles are not well defined. In this setting, entropy can be used to quantify the spread of probabilities over the categories. This research discusses the meaning of entropy in this context and demonstrates its efficiency as a summary encoding to convey uncertainty in brain magnetic resonance images.

This work was submitted to the Eurographics Conference on Visualization (EuroVis).

Progress and Status:

- Personnel - No changes during this reporting period.
- Equipment - No equipment was purchased during this reporting period.

Scope issues:

There are no scope changes.

Budget and Schedule Status:

There are no budget nor schedule status changes.

Patents:

There are no patent applications attached to this award.

Publications / Presentations:

Publications

I. Altrogge, T. Preusser, T. Kröger, S. Haase, T. Pätz, and R.M. Kirby. *Sensitivity Analysis for the Optimization of Radiofrequency Ablation in the Presence of Material Parameter Uncertainty*, In International Journal on Uncertainty Quantification .

B. Nelson, R. Haimes, R.M. Kirby. *GPU-Based Interactive Cut-Surface Extraction From High-Order Finite Element Fields*, In IEEE Transactions on Visualization and Computer Graphics .

C. Yang, D. Xiu and R.M. Kirby. *Visualization of Covariance and Cross-covariance Fields*, In Int. J. Uncertainty Quantification.

Presentations

Q.Meng, *Scaling Uintah's AMR to 200K cores and Beyond*, AMR session SIAM Parallel Processing for Scientific Computing Conference Feb 2012. (Invited talk).

A.Humphrey, *Migrating the Uintah Framework to CPU/GPU Architecture* SIAM Parallel Processing, Scientific Computing Conference, Feb 2012 .

M.Berzins, *Scalable Parallel Solution of Fluid & Structure Interaction Problems - Uintah and Wasatch*, Leeds University, School of Computing, UK Jan 2012 (Invited talk) .

M.Berzins, *Uintah a Scalable DAG-Based Framework For Multi-Scale Multi-Physics Problems*, NSF Computation workshop, APS Meeting Boston Feb 2012 (Invited talk).

Chris Johnson, *Visual Computing: Making Sense of a Complex World*, Rutgers University, New Brunswick, March 2012 (Distinguished Lecture) .

Chris Johnson, *The Art of Science*, Kimball Art Center, Park City, March 2012

Mike Kirby, *Quantifying Uncertainty in an Attempt to Answer Biomedical Questions: From the Computer Lab to the Clinic*, Imperial College London (UK) March 2012

Plans for Next Quarter:**Uintah Software System**

This section of the project is proceeding well. We expect to have the energy applications described above running on both NSF and DOE machines on larger core counts very soon. The results obtained for the radiation GPU code are promising and we hope that these will transfer to a full-scale simulation. Our long term goal is to be able to run all these simulations on the full Titan Architecture when it comes online towards the end of 2012.

Algorithm Development

Ross Whitaker, Zhisong Fu and Mike Kirby are working on a manuscript concerning accomplishing FEM simulations on the GPU. They anticipate submitting the manuscript to the Journal of Computational and Applied Mathematics in the next two months. In this work, they present a full finite element method simulation pipeline for solving bilinear problems on unstructured meshes using the GPU. Their proposed GPU assembly step needs very few preprocessing and shows significant performance boost compared to an optimized CPU implementation. For the linear system solving, they use a variant of multigrid method and propose novel data and computing mapping strategies and corresponding data structures to suit GPU architecture.