Quarterly Progress Report
Covering Period: 10/1/2010 - 12/31/2010
Date of Report: 01/31/2011

Award Number:	DE-EE0004449
Project Title:	Unconventional and Renewable Energy Research Utilizing Advanced Computer Simulations (UT)
Project Period:	09/23/2011 - 03/22/2012
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: Jacqulyn Wilson
304-284-4135
Jacqulyn.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 the sparse systems associated with large simulations, and 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. The PI, Chris Johnson, co-leads the DOE Visualization and Analytics Center for Enabling Technology (DOE-VACET). In this work we are leveraging our expertise in large-scale visualization research and development toward the seamless integration of high-end visualization techniques with simulation results of science-based energy systems. 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 uncertainly 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 Unintah for provenance systems.

Accomplishments:

Uintah Software System:

Dynamic Task Scheduling for Scalable Parallel AMR

Over the past months, we have continued work on the Uintah Software. The software has been ported to the Jaguar computer at Oak Ridge as part of the INCITE allocation we have there. in order to make it possible for Uintah to run on multi-core and eventually GPU machines we are rewriting the core data warehouse in Uintah to make use of cases in which many cores share memory and there is less memory per individual core.

This also requieres us to rethink the task schedular for Uintah. The original task scheduler for Uintah ran computational tasks in a predefined order. To improve the performance of Uintah for petascale architecture we have designed a new dynamic task scheduler that allows better overlapping of the communications and computations. The new scheduler supports asynchronous, out of order scheduling of computational tasks by putting them in a distributed directed acyclic graph(DAG) and isolating task memory. In this work we have shown the effectiveness of this new approach on large scale fluid-structure examples. We now need to start to extend the scheduler still further to address the case when tasks must be able to run on a large number of cores with a relatively small data bandwidth to memory. This work is currently underway.

Algorithm Development:

Solving PDEs on GPUs

Over the first quarter, we have continued work on alternatives for solving PDEs on GPUs. One such alternative is to exploit the topology of the domain, which determines the sparsity of the associated matrices (*e.g.* stiffness matrix) for the linear formulation. In prior work we have done this successfully for finite-difference formulations on regular grids. In that work nonlinear diffusion equations are solved on the GPU using a semi-implicit scheme with a conjugate gradient solver implemented on the data structure of the domain, without the use of an intermediate matrix representation or a generic solution for the linear system. In this case we saw GPU efficiencies of 2550%*e.g.* 50100 advantage over a CPU implementation. This strategy takes advantage of the fact that virtually all of the algebraic operations on the associated linear system have a geometric interpretationthey are local operations on neighborhoods of grid points or reductions, accumulations of data across the domain. The investigators have demonstrated various types of solvers for nonlinear equations on regular grids using this strategy.

Scientific Visualization:

Large-Scale Ray Tracing Visualization:

We plan on investigating ray-tracing for the analysis of large-scale data. In particularly, ray-tracing is efficient for datasets best represented by a large number of spherical glyphs such as MPM simulations such as those in Uintah as well as moleculate simulations. We will investigate the use of state of the art NUMA architectures for interactive ray-tracing of massive datasets. We will investigate the optimal acceleration structure by studying the scaling of both grid-based methods and bounding volume hierarchies.

Uncertainty Visualization:
Over the first quarter we hired a postdoctoral research fellow, Dr. Joel Daniels, who will spend .5 FTE of his time working on creating new uncertainty visualization techniques. We also created an Uncertainty Visualization Seminar Series that features presentations on uncertainty visualization on a twice monthly basis. Technical progress in the first quarter focused on meeting with application scientists and engineers and project Co-PIs to assess uncertainty analysis and visualization needs.

Progress and Status:

- Personnel - In the past quarter we have identified and added personnel for the proposed work.
- Equipment - The SCI Institute has purchased and installed a 200+ core SGI UV shared memory computing system that will benefit this work.

Scope issues:

There are no scope changes.

Budget and Schedule Status:

Not Included in Web Version

Patents:

There are no patent applications attached to this award.

Publications / Presentations:

Since the award is new, there are no publications resulting from this award yet.

Chris Johnson gave an invited panel presentation on Uncertainty Large-Scale Visualization at the *Extreme Scale Visual Analytics Workshop*, Salt Lake City, October 2010.

Chris Johnson gave an invited presentation on Uncertainty Visualization at the *Uncertainty Quantification in Scientific Inference Conference*, Santa Fe, November 2010.

Chris Johnson received the 2010 IEEE Visualization Career Award at the *IEEE Visualization 2010 Conference*, Salt Lake City, October 2010.

Ross Whitaker Co-Chaired the *IEEE Visualization 2010 Conference*, Salt Lake City, October 2010.

Plans for Next Quarter:

Algorithm Development

In the area of GPU processing, over the next several quarters we will work toward extending the 2D fast iterative method (an unstructured eikonal solver) to 3D tetrahedral meshes and will also exploring unstructured multigrid solvers on the GPU for solving FEM-constructed linear systems.

Provence Enabling Uintah

Over the next two quarters we will generate a high level strategy for identifying the logical points of data/state extraction for the Uintah system.

Visualization

During the next two quarters we will work toward identifying datasets to use in the examination of the use of high order elements for visualization.