Quarterly Progress Report
Covering Period: 7/1/2011 - 9/30/2011
Date of Report: 10/15/2011

Award Number: DE-EE0004449

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

Project Period: 09/23/2010 - 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 Uintah for provenance systems.

Accomplishments:

Uintah Software System:

Uintah uses domain decomposition with a task-graph approach for asynchronous communication and automatic message generation. Uintah was originally designed to run and perform well on around 1000 processors using a static task scheduler, running tasks in a predefined order. Through significant development and advancements in its load balancer and by adding the ability to dynamically execute tasks out-of-order to its scheduler, Uintah evolved to demonstrate weak and strong scalability up to 98,304 cores. With Uintah, we have now shown that with an MPI-only approach, memory use associated with ghost cells and global meta-data becomes a barrier to scalability beyond O(100K) cores. This limitation has recently been overcome in Uintah through the use of hybrid parallelism, significantly reducing the memory footprint of the code per core. This hybrid memory approach (a combination of Pthreads and MPI) has been incorporated into a Threaded-MPI scheduler (Figure **??**) that has a single MPI process per socket and a control thread that assigns tasks to available worker threads (CPU cores). With the hybrid scheduler, Uintah now demonstrates scalability up to 196k cores on the DOE Jaguar system (Figure 2).

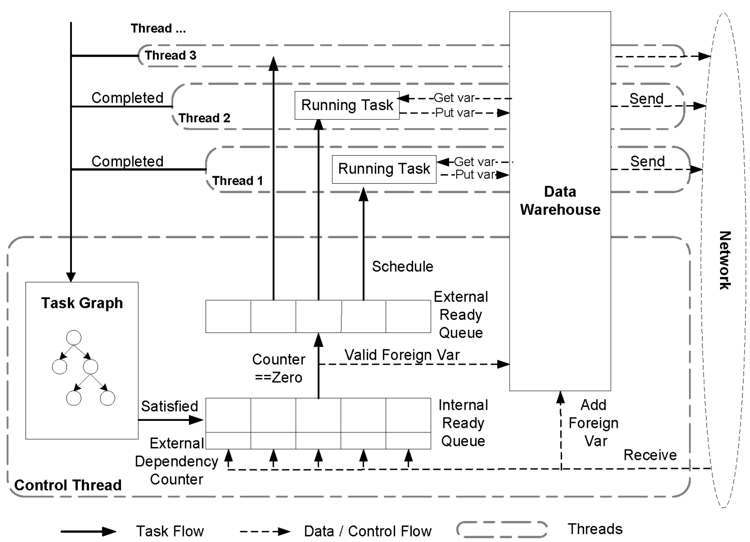

Figure 1: Threaded-MPI Task Scheduler.

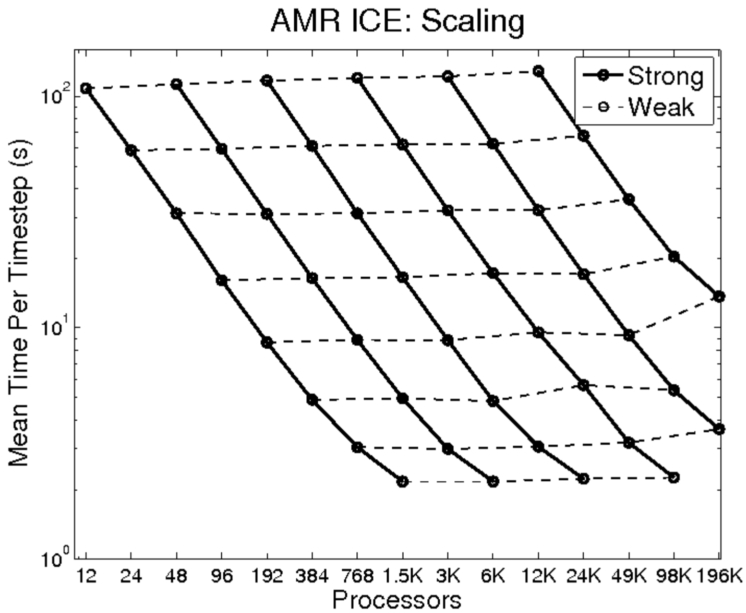

Figure 2: Uintah Scalability using Threaded-MPI Scheduler.

With advancements in scalability and decreased global memory usage due to Uintah's new hybrid scheduler, the next step has been to extend this scheduler and Uintah's asynchronous task-based approach to fit emerging CPU/GPU architectures. The focus of our work during this quarter has been to demonstrate and analyze the performance of key portions of Uintah on mixed CPU/GPU architectures using Nvidia's CUDA C language extensions. OpenCL will also be a longer-term consideration. This is partly in response to industry need to accelerate existing codes by running on ever-more prevalent GPUs, but also in preparation for hybrid architectures such as the proposed DOE Titan system.

The Uintah Threaded-MPI task scheduler provides the basis for our newest work. We are extending the hybrid scheduler to include dispatching tasks to other types of processors. Our current focus is on adding the infrastructure support necessary to additionally schedule and execute GPU tasks. This would mean that per node, Uintah's scheduler would have the ability to execute tasks on both CPU cores as well as all available GPUs. This will speedup key, computationally intensive regions of specific Uintah tasks, and will allow Uintah to make use of all available execution hardware on emerging architectures. Our design will remain broad enough to accommodate the use of other accelerator designs such as the Intel MIC chip. Our work this quarter began by first adding full CUDA support to the existing Uintah build system. This portion of our work is complete and has been proven successful on several local GPU clusters. Initial analysis was then done to determine what tasks would readily lend themselves to SIMD parallelism and exhibited properties that would yield a performance increase. For highly regular Uintah grid variables, we have been able to expose the underlying representation (raw memory), essentially a flat array representation of an abstract three-dimensional data structure. This exposed, raw memory is easily mapped onto a GPU where individual cells are then SIMD processed. Methods have been explored and tested to make use of CUDA pitched memory, essentially hardware managed, padded memory to handle alignment restrictions, avoid non-coalesced memory transactions and to automatically offer the best performance while future-proofing Uintah GPU code. Our design includes offering the ability to have GPU-able tasks maintain a CPU version as well for further flexibility.

Initially, we have focused on migrating portions of the ICE component (the multi-material CFD code) to GPUs by profiling and analyzing several key ICE algorithms to estimate the theoretical peak performance we might see on Nvidia Tesla C20 series GPUs. These comparisons were based on a CPU similar to that found in Cray systems such as Kraken and Jaguar; a dual core AMD Opteron Processor 285, running at 2.6 MHz. Based on this analysis, it was determined that the first and second order advection operators within the ICE explicit pressure solve would see a 2-3x speedup. These operations comprised nearly 50% of the overall runtime for the ICE component.

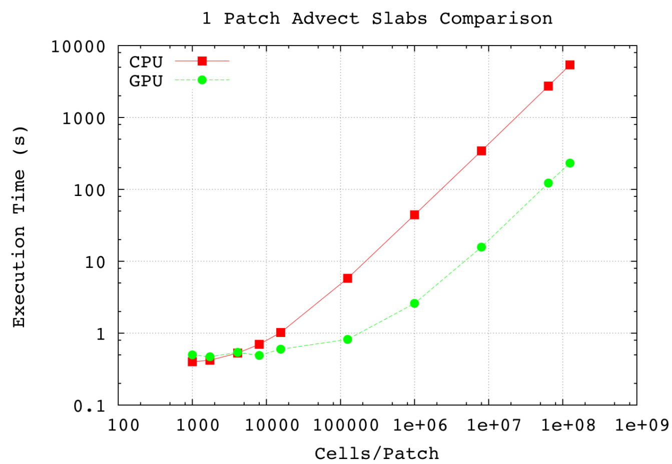

Figure 3: GPU vs CPU ICE AdvectSlabs Task Execution Times.

GPU versions of particular ICE tasks have been implemented and tested against CPU versions of the same task. Confirming our analysis, Figure 3 shows the executions times for the ICE::AdvectSlabs task on a 8-core, Intel(R) Xeon(R) CPU X5570 running at 2.93GHz compared to a Tesla C2050 / Tesla C2070 GPU, running 515 GFLOPs/sec of double precision processing performance, with 3 GB of GDDR5 memory offering 144 GB/s bandwidth.

Initial modifications to the Uintah infrastructure are currently underway to extend the ability of the existing hybrid Uintah scheduler to execute GPU tasks. A significant challenge to overcome is that much of the computation done within Uintah tasks will be bandwidth bound, in that data must be copied to and from the device between time-steps. These data copies are across the PCIe bus (PCI Express Gen2 16 for the C20 series cards), which has a maximum theoretical bandwidth of 8.0GB/s. In practice, we have found the actual maximum PCIe bandwidth for these copies is more on the order of 3.3GB/s using paged memory and 5.3GB/s using pinned host memory. We are investigating ways to keep certain data on the GPU.

Based on this challenge, the design being pursued will provide Uintah tasks the ability to do a memory preparation step prior to actual task execution. This functionality would be handled by the infrastructure, as it has full knowledge of task requirements up front. This ability will truly overlap communication and computation and will also serve to shield the component developer from the details of the explicit GPU memory allocation and copies. This idea might also be used for other accelerators, such as the Intel MIC chip, which will also have the PCIe bus as a bottleneck.

Algorithm Development: Solving PDEs on GPUs

GPU-Based Interactive Cut-Surface Extraction From High-Order Finite Element Fields

We present a GPU-based ray-tracing system for the accurate and interactive visualization of cut-surfaces through 3D simulations of physical processes created from spectral/hp high-order finite element methods. When used by the numerical analyst to debug the solver, the ability for the imagery to precisely reflect the data is critical. In practice, the investigator interactively selects from a palette of visualization tools to construct a scene that can answer a query of the data. This is effective as long as the implicit contract of image quality between the individual and the visualization system is upheld. OpenGL rendering of scientific visualizations has worked remarkably well for exploratory visualization for most solver results. This is due to the consistency between the use of first-order representations in the simulation and the linear assumptions inherent in OpenGL (planar fragments and color-space interpolation). Unfortunately, the contract is broken when the solver discretization is of higher-order. There have been attempts to mitigate this through the use of spatial adaptation and/or texture mapping. These methods do a better job of approximating what the imagery should be but are not exact and tend to be view-dependent. This research introduces new rendering mechanisms that specifically deal with the kinds of native data generated by high-order finite element solvers. The exploratory visualization tools are reassessed and cast in this system with the focus on image accuracy. This is accomplished in a GPU setting to ensure interactivity.

Uncertainty Visualization

Uncertainty Workshop at IEEE VisWeek 2011

Chris Johnson has been organizing a full day workshop : *Working with Uncertainty Workshop: Representation, Quantification, Propagation, Visualization, and Communication of Uncertainty* for the IEEE VisWeek Conference in October in Providence. Kristi Potter is working on the following two projects submitted to the workshop.

Uncertainty in the Development and Use of Equation of State Models The research objective of this project is to determine effective visualization techniques for uncertainty of equation of state material models. The ultimate goal is to develop a software tool to be used within the material modeling stakeholder community for understanding material model behavior, including material model uncertainty. Focus groups were used to identify and understand the needs of potential users (the focus group participants) regarding material model visualization and analysis, with the intent of ensuring the usability, utility, and adoptability of the software eventually developed. Four prototype visualizations were developed for displaying the uncertainty of equation of state models; these prototypes are discussed on a companion project on *Uncertainty Visualization Prototypes for Materials Modeling*. Within the focus groups, the prototypes provide a concrete object for participants to discuss. While comments about the prototypes were expected, focus groups revealed a number of technical and organizational issues related to material modeling and material model uncertainty, as well as opportunities for visualization and analysis software to address these issues.

Uncertainty Visualization Prototypes for Materials Modeling Material models describe the behavior of a specific material or class of materials and are used as inputs to multiphysics numerical simulations. Because the models are based on theory, they often require empirical information to calibrate or specify free parameters and results from the models may define ranges of possible values or a collection of valid scenarios. Sources of uncertainty within simulations are abundant; this work focuses on the uncertainty arising from variability in the material models and aims at understanding how users comprehend, incorporate, and utilize this qualitative information and how to enhance understanding through visual representations. Our approach begins with understanding who the users of the material models are and how they are working with uncertainty information. To this end, we have conducted focus groups to engage modelers, analysts and code developers from the Sandia material modeling community in discussions on how visualization can help them understand the impact of material uncertainties in their workflow. The focus groups provided us with detailed insights into the challenge of developing usable and useful representations of material models and associated uncertainties for a user community with a wide range of interests and applications for such visualizations. To facilitate discussion, we developed four visualization prototypes, each of which present uncertainty within a material model in a unique way. Participants evaluated specific features of each prototype and described scenarios in which different elements could prove helpful. We present each of the prototypes used in the focus group, using results from a simplified equation of state simulation which produced seven realizations of a material surface.

Provenance

The VisTrails provenance management system has been extended to capture provenance during a user session with the Uintah problem-solving environment. Currently, a small number of VisTrails user modules are in place to allow the system to execute Uintah-based simulations from within the VisTrails environment. Additionally, VisTrails modules have been created to allow scientists to change the inputs to their simulations. By enabling this functionality within VisTrails, robust provenance-enabled parameter studies are possible using the native properties of VisTrails.

Progress and Status:

- Personnel

  Sergey Yakovlev, PhD. He is currently working as a postdoc on extending my Hybridized Discontinuous Galerkin work (previously done in 2D) to 3D within the Nektar++ framework. He is working with the DOE funded postdoc Linh Ha to also extend the 2D HDG work to the GPU

  Erik Anderson, PhD. He started working as a postdoc on the Provenance Project during this current 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

K. Potter, R.M. Kirby, D. Xiu, and C.R. Johnson. *Interactive visualization of probability and cumulative density functions.* International Journal of Uncertainty Quantification, 2011 (submitted).

F. Jiao, Y. Guv, J. Phillips, and C.R. Johnson. *Uncertainty Visualization in HARDI based on Ensembles of ODFs.* Proceedings of the IEEE Pacific Visualization Conference 2012, (submitted).

Z. Fu, R.M. Kirby, R. Whitaker. *A Fast Iterative Method for Solving the Eikonal Equation on Tetrahedral Domains.* Journal of Computational Physics (accepted pending revisions).

B. Nelson, R. Haimes, R.M. Kirby. *GPU-Based Interactive Cut-Surface Extraction From High-Order Finite Element Fields.* IEEE Transactions on Visualization and Computer Graphics, October 2011 (accepted).

Presentations

Kristi Potter, *Stochastic Modeling and Uncertainty Quantification*, Uncertainty Visualization State of the Art USA and South America Symposium, August 3, 2011

Mike Kirby, *Sensitivity Analysis for the Optimization of Radiofrequency Ablation in the Presence of Material Parameter Uncertainty*, Uncertainty Visualization State of the Art USA/South America Symposium, August 3, 2011

Chris Johnson,*Visualization of Error and Uncertainty*, IFIP Working Conference on Uncertainty Quantification in Scientific Computing, Boulder, August 2011 (Invited Talk).

Chris Johnson, *Image-Based Biomedical Modeling, Simulation, and Visualization*, 24th Biennial Conference on Numerical Analysis, Glasgow, Scotland, July 2011.

Chris Johnson, *Visualizing Uncertainty*, IFIP Conference on Uncertainty Quantification in Scientific Computing, Boulder, August 2011.

Plans for Next Quarter:

Uintah Software System:

The next stage in Uintah GPU development will be to complete a working prototype of the Threaded-MPI-GPU scheduler, that is able to ultimately run a larger-scale, distributed version of ICE that executes both CPU and GPU tasks. An allocation on the Keeneland system is being pursued for this work during the next quarter. Keeneland is a Linux cluster based on commodity next-generation processors and GPUs, and is managed by the Georgia Institute of Technology, Oak Ridge National Lab, University of Tennessee-Knoxville, and the National Institute for Computational Sciences. Keeneland will be added to the XSEDE resource allocation base and would allow our work to progress toward running Uintah on the proposed DOE Titan system.

A local system here at the SCI Institute has been identified and upgraded with Nvidia Fermi-based GPU hardware and will be utilized in our Uintah GPU development for nightly regression testing and a Buildbot to automate the compile/test cycle required to validate changes. This machine is currently built and will be configured and put into commission during the first part of the next quarter.

Analysis similar to that which was done for the ICE component will then be performed for the Material Point Method (MPM) code. Here, objects are discretized into particles, or material points, each of which contains all state data for the small region of material that it represents. This includes the position, mass, volume, velocity, stress and state of deformation of that material. Particles do not interact with each other directly, rather the particle information is accumulated to the grid, where the equations of motion are integrated forward in time. This time advanced solution is then used to update the particle state.

Our goal beyond the next quarter will be to run large-scale, distributed MPM and ultimately MPMICE simulations on hybrid CPU/GPU architectures.

Provenance Enabling Uintah

Although recording some provenance of Uintah-based simulation runs is now possible, more work is required to extend VisTrails. The current implementation of Uintah in the VisTrails environment does not support the execution of simulations on remote computing platforms. Additional work must be completed to properly communicate with remote job schedulers and the job submission portal currently in development.

Uncertainty Visualization

QuizLens Kristi Potter will be submitting the QuizLens project to Eurovis 2012 on Dec 4th. This work is evolving into a uncertainty visualization system to aid in the understanding of probabilistic volumetric data.

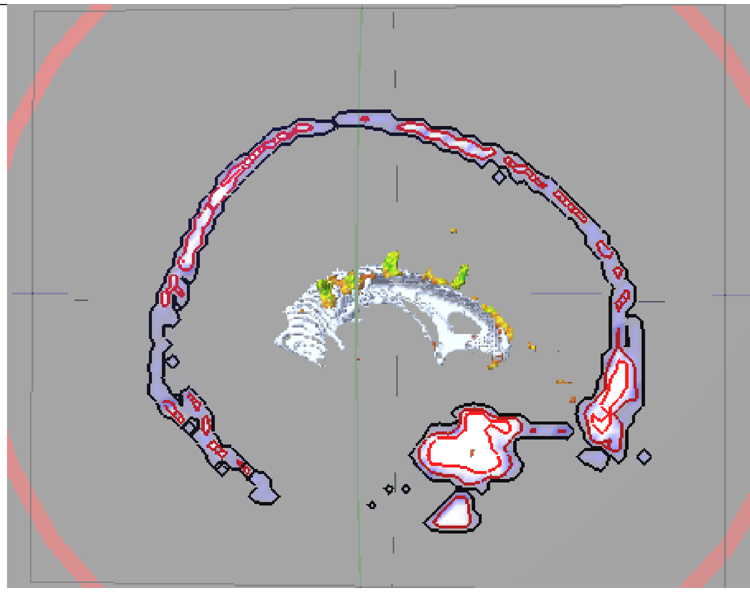

Figure 4: Probabilistic boundaries between tissues types which are displayed both in 2D screen space and as 3D volumetric.