Quarterly Progress Report
Covering Period: 4/1/2012 - 6/30/2012
Date of Report: 7/15/2012

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: 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 the 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 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

Linear Solver Interface

Uintah's linear solver interface has been significantly improved and can now be used to solve energy related problems up to 256K cores on the Jaguar XK6 system.

For both the Alstom Boiler Problem and Carbon Capture Modeling, the Uintah ARCHES Combustion Simulation Component is used to solve the discretized version of the Navier-Stokes equations. ARCHES is a stencil-based PDE code and so achieves scalability through its use of the Uintah infrastructure. The low-Mach, pressure-approach of ARCHES requires a solution of a pressure-projection set of equations at every time step. Typically both the PETSc and the hypre packages have been used to solve these systems in parallel, in what can be the most computationally intensive part of the simulation.

In our technical report "Large Scale Parallel Solution of Incompressible Flow Problems using Uintah and hypre", we observed that the overall time to solution for the linear solve could be modeled using a power law relationship dependent on the number of cores. Using this power law relationship, we can predict the time to solution for core counts that exceed the current machine capabilities. We used the raw scaling data and linear least squares fit of the power law relationship (from the log-log plot of the data). By looking at the overall scalability of the linear solver and the individual components, we can understand the direction that improvements in linear solver technology must take in order to solve the challenging problems that we are investigating in our current work.

We also looked at several test problems that embodied the complicated physics underlying both the Alstom Boiler Problem and the Carbon Capture Problem. The Taylor Green flow and an example taken from the modeling of Helium plumes were used to look at the individual multi-grid solver steps that make up the bulk of the calculation time for the ARCHES calculation (i.e. the solution of the Pressure Poisson equation). The Helium plume examples under study is challenging as it requires both the full solution of the incompressible Navier Stokes equations and uses sub-grid models to account for any unresolved turbulence scales, and thus is a an import test problem for ARCHES and Uintah.

Using Uintah on Hybrid Architectures A hierarchical, GPU-based radiation transport model has been developed, showing speedups of 30-60 times over the CPU equivalent.

We have now developed a hierarchical ray-tracing, radiation-transport model for the GPU that provides Uintah with additional capabilities for heat transfer and electromagnetic wave propagation in general. This GPU-based radiation transport model has shown significant speedups over the CPU counterpart and has been run at scale on the heterogeneous TitanDev (ORNL) and Keeneland systems (NICS, ORNL). This GPU-based radiation-transfer mechanism runs in unison with the CPU portion of a simulation on heterogeneous architectures using Uintah's hybrid CPU-GPU task scheduler. This hybrid task scheduler is capable of assigning tasks to both CPU cores and available GPUs on-node, achieving a high degree of node-level parallelism. To achieve this capability, Uintah's runtime system has been extended by an added multi-stage GPU queuing architecture and GPU task controller for efficient scheduling of both CPU and GPU tasks.

This new runtime system automatically handles the details of asynchronous memory copies to and from the GPU and introduces a novel method of pre-fetching and preparing GPU memory prior to GPU task execution. Our design remains broad enough to make use of other accelerator designs such as the Intel MIC (Xeon Phi) chip.

Algorithm Development

A High-Performance Multi-Element Processing Framework on GPUs Many computational engineering problems, ranging from finite element methods to image processing, involve the batch-processing of an elemental operation on a large number of data items. While multi-element processing has the potential to harness the computational power of streaming massively SIMD systems, current optimization techniques often concentrate on maximizing the performance of the elemental operation. Frameworks that take this greedy optimization approach often fail to extract the maximum processing power of the system. One can exploit the knowledge that the same operation will be accomplished on a large number of data items to organize the computation to maximize the use of the available streaming hardware. We present a high-performance, multi-element processing framework that harnesses the processing power of GPUs. To demonstrate the efficacy of our approach, we reported the results of using batch processing of commonly used linear algebra operations

This work was submitted to IEEE SuperComputing 2012.

Efficient Coordination of Linear Algebra Operations for h and p Finite Element Methods

Finite element methods (FEM) are applicable across a range of application areas and are ubiquitously used throughout many scientific and industrial communities. In low-order FEM, contributions from each elemental subdomain are often assembled into a single global matrix, while high-order methods operations are typically performed per subdomain. There is, therefore, a rich structure arising from this elemental construction, permitting a range of algorithm strategies to target processor cache boundaries and limit function-call overheads. In our research project, we demonstrate how a hybridized approach to assembly, using a preprocessing library to selectively coalesce elemental submatrices, can lead to optimal processor throughput during FEM operations and reduced solution times. We highlight our approach by first demonstrating and ameliorating the previously recognized inefficiency of BLAS operations for low-rank matrices. Performance improvements are then obtained by coalescing elemental sub-matrices to construct matrices of optimal throughput for a given BLAS implementation and system specifications.

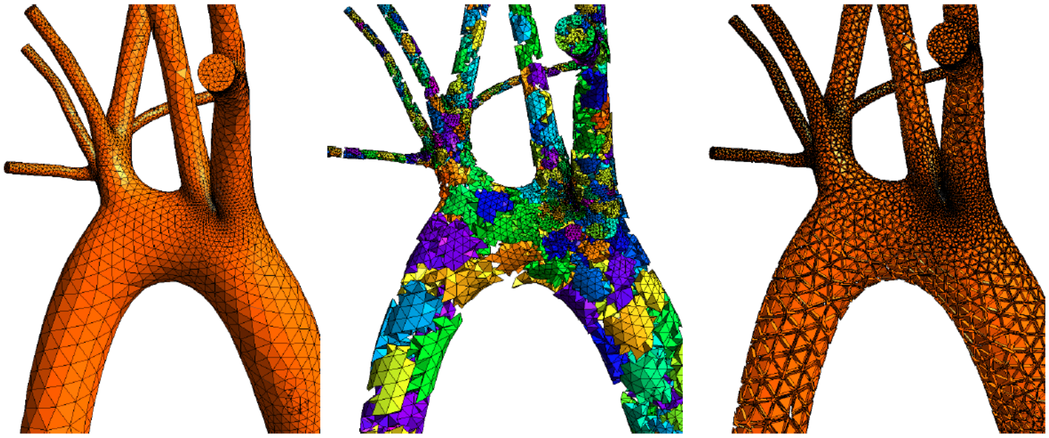

Figure 1: Graphical example of the three evaluation strategies in the context of an aorta simulation showing the globally assembled system (left), local elemental system (right) and a partitioning of the mesh to create a hybridized system (center).

This work, in collaboration with Imperial College in London, was submitted to IEEE Super-Computing 2012.

Scientific Visualization

High dimensional visual exploration of uncertainty

In our preliminary work, an ensemble of observations is placed into HD feature space and projected to 2D using PCA with basic interaction capabilities for data exploration by defining regions to select subsets of points. Figure 2 is an example of this type of data, where a volume of data is defined with each voxel containing 11 probabilities, each representing one of 11 material types (i.e. bone, white matter, skin, air, etc.). The first step of the process is to cast the data into a HD space by forming an ensemble of values which is simply done by assembling the 11 voxel probabilities into an 11D ensemble.

The next step of our approach is to reduce the dimensionality of the HD feature space for display on the computer screen. Many-dimensionality reduction techniques exist, each with unique capabilities for bringing out data features. For example, PCA will find strong linear components within the data while approaches such as Isomap or Local Linear Embedding might better cluster similar points. In our work, PCA is used to reduce the dimensionality of the ensembles down to 2D for direct display on the computer screen.

The final step of the process is interacting with the projected feature space by allowing users to select regions and refine their dimensionality reduction. Simultaneously, the linked-view of the spatial domain reacts to selection and refinement. In our example in Figure 2, interesting features are selected and refined as they are visually identified. The refinement process quickly leads to a meaningful segmentation of various material types without a significant need to understand the input data, the high-dimensional embedding, or the projected feature space.

Uncertainty Visualization

Uncertainty when there is no mean

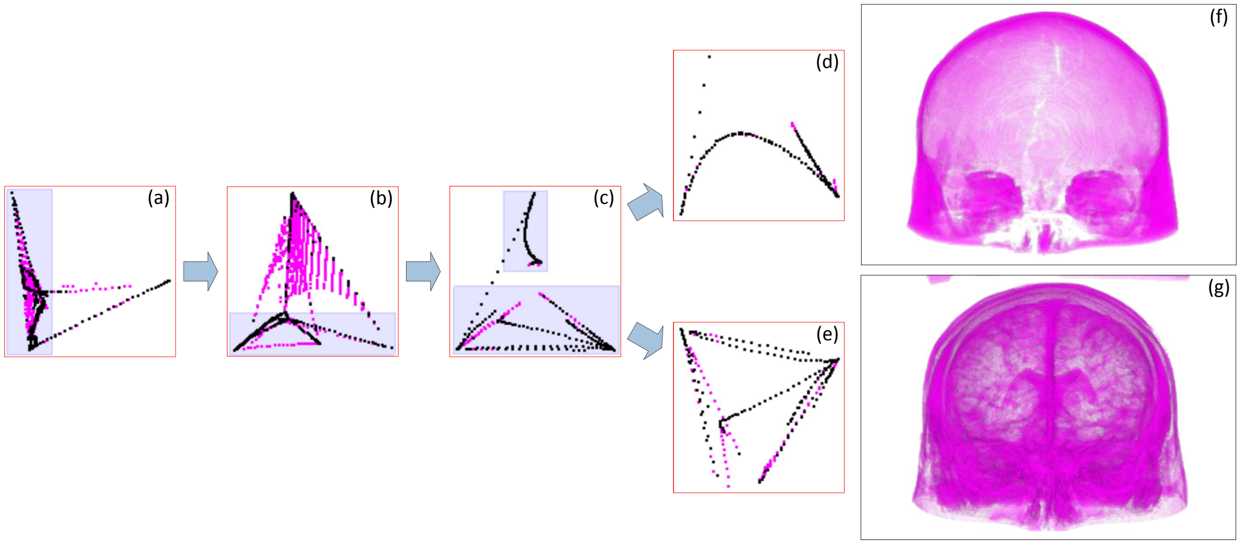

Figure 2: Demonstration of our high dimensional feature-detection and exploration approach. The head dataset used fuzzy segmentation to produce a percentage distribution function (PDF) for each of 11 material types. (a) These 11 probabilities are assembled into ensembles and placed high dimensional space. Finally, they are projected down to 2D using PCA. (b-c) The user is able to select subsets of points and refine the analysis. (e-f) Interesting structures, once selected, can be visualized (g-h).

Uncertainty is tricky business. The term is used to encompass all sorts of unknowns, including error, deviations, missing information, or confidence levels. There exist numerous methods for quantifying and expressing uncertainty, and its existence is persistent and accumulative throughout the visualization pipeline. For visualization researchers, this term is particularly complex due to the limitations of the visual display medium. As such, aggregation is a commonplace technique used to summarize uncertainty for visualization purposes, with two particular summary statistics standing out as the de-facto characterization: mean and standard deviation. These statistics reduce uncertainty down to an expected value and variation from that value, and are particularly effective in expressing normally distributed data. However, these statistics are not always appropriate or even feasible, particularly when the uncertainty within a data set cannot be described as a probability distribution function.

This work was submitted to Visualization Viewpoints in IEEE Computer Graphics & Applications Magazine.

Provenance

Davison de St. Germain has been added part time to the NETL project, replacing Erik Anderson on the provenance management system. He has reviewed the Uintah VisTrails prototype module and discussed the current functionality with Dr. Anderson, including examining possible methods to harden and extend the current functionality as well as approaches for increasing the codes ability to run both remotely and on large supercomputers. While the current prototype works in a development environment in a desktop situation, it will need to be greatly enhanced in order to provide a tool that Uintah end users would be comfortable using.

To begin the move from prototype to end user system, Mr. de St. Germain has installed VisTrails both locally on a desktop platform and on one of the University of Utah's High Performance

Computing (HPC) supercomputing platforms. This will allow for local development, and will provide a stepping stone for the installation of an upgraded Uintah VisTrail module on one of the supercomputers at the Texas Advanced Computing Center (TACC). (TACC will provide the hundreds of thousands or millions of computer hours necessary to run parameter tests using the Uintah system.)

At a high level, there are two main objectives for this effort:

1) Update the current VisTrails module to be user friendly and more robust. 2) Update the current VisTrails module to work in an HPC setting, including support for remote execution and the ability to interact with a batch scheduling system.

Progress and Status:

- Personnel - Davison de Saint Germain was added to the Provenance project as a Chief Software Developer. Prof. Paul Rosen was added to the Visualization project. Paul Rosen joined the SCI Institute as a Post Doc in 2010 and was promoted as an Assistant Research Professor in 2011.
- 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

V. G. Weirs, N. Fabian, K. Potter, L. McNamara, T. Otahal. *Uncertainty in the Development and Use of Equation of State Models*. International Journal for Uncertainty Quantification, 2012

M. Berzins, J. Schmidt, Q. Meng, A. Humphrey. *Past, Present, and Future Scalability of the Uintah Software*. Blue Waters Workshop, 2012

A. Humphrey, Q. Meng, M. Berzins, and T. Harman. *Radiation Modeling Using the Uintah Heterogeneous CPU/GPU Runtime System*. In Proc. XSEDE 2012

J.R. Peterson, J.C. Beckvermit, T. Harman, M. Berzins, and C.A. Wight. *Multiscale Modeling of High Explosives for Transportation Accidents*. In Proc. XSEDE 2012

J. Schmidt, M. Berzins, J. Thornock, T. Saad, J. Sutherland. *Large Scale Parallel Solution of Incompressible Flow Problems using Uintah and hypre*, Technical Report, No. UUSCI-2012-002, Scientific Computing and Imaging Institute, 2012

Q. Meng and M. Berzins. *Scalable Large-scale Fluid-structure Interaction Solvers in the Uintah Framework via Hybrid Task-based Parallelism Algorithms* Technical Report, No, UUSCI-2012-004, Scientific Computing and Imaging Institute, 2012.

Presentations *Visual Computing*, NSF OCI CyberBridges Workshop, DC, June 2012 (Keynote Presentation).

Image-Based Biomedical Computing and Visualization, 3rd International Emerging Computational Methods for the Life Sciences Workshop, Delft, June 2012 (Keynote Presentation).

Visualization in Art and Humanities Digital Humanities Workshop, Oxford, June 2012 (Keynote Presentation).

Large-Scale Visual Data Analysis, International Parallel and Distributed Processing Symposium, Shanghai, May 2012 (Keynote Presentation).

Visual Computing, Peking University, Beijing, May 2012.

Image-Based Biomedical Modeling, Simulation, and Visualization, Symposium on Visualizing Complex Biomedical Systems, Experimental Biology Conference 2012, San Diego, April 2012.

The Challenges at the Interface of Life Sciences and Cyberinfrastructure, ECMLS Workshop, Delft, June 2012.

Plans for Next Quarter:

Uintah Software System

This section of the project is proceeding well. Uintah now scales to near full-machine capacity on Jaguar XK6, significant improvements have been made to Uintah's linear solver interface, allowing for large-scale, energy related problems to be solved, and Uintah is successfully running on hybrid architectures. A de-centralized CPU-GPU scheduler will also be implemented to replace the current master-worker model to GPU-based radiation transport model will also be done. Our long-term goal is to be able to run all of these simulations on the full Titan Architecture when it comes online towards the end of 2012.

Provenance

Mr. de St. Germain has also commenced direct and ongoing discussions with the scientists that use Uintah for their large-scale fluid and structural computational dynamics simulations. These preliminary discussions are aimed at more specifically specifying the functionality that will be necessary to provide automated provenance and parameter exploration in Uintah.