

Final Report
Workshop on CS/Math Institutes and
High Risk / High Payoff Technologies for Applications
<http://www.csm.ornl.gov/workshops/institutes/>

Al Geist, Phil Colella, Mike Heroux

A two and a half day workshop on CS/Math Institutes and High Risk / High Payoff Technologies for Applications was held October 7-9 at the Chicago O'Hare airport. There were 30 attendees from universities, DOE labs, and companies from across the country. All of the talks presented, the list of discussion questions asked of the attendees, and the summary presentations of the discussions are all available at the workshop website:

<http://www.csm.ornl.gov/workshops/institutes/>

The workshop was run in two parallel tracks. The High Risk technologies track had presentations by several application developers describing the potential high risk/high payoff areas in their application domain. These were followed by discussions by the attendees on how the math and CS communities could help make this high payoff a reality. The outcome of this track was a list of potential high risk endeavors and a few illustrative examples.

The charge to the Institutes track was to explore the community's thoughts on how to create, develop, and manage joint math and computer science institutes. There were four presentations in this track. Each presentation generated long discussions of the challenges and opportunities for effective math and CS collaboration. The outcome of this track was a list of challenging and important areas that could become themes for CS/Math institutes.

Joint CS/Math Institutes Track

The vision for a CS/Math institute is an organization either physical or virtual that is comprised of both computer science and math researchers working together on a set of projects under a single unifying theme. It is expected that the Math and CS researchers would be a mix of lab, industry, and university participants. The workshop attendees agreed that a successful project must be tightly coupled and synergistic, to the extent that the success of the effort relies on the combined success of both the Math and CS components. Given the multi-institutional and multi-disciplinary nature of the institutes, the participants discussed the management challenges and concluded that:

- There should be a single PI and approximately 10-20 members
- An Institute should have a single theme with multiple projects.
- Projects must clearly demonstrate a need for a combined Math/CS effort and milestones must depend on joint effort.
- Institutes must be long-lived (5-10 years) to realize their potential.
- Funding needs to be adequate to sustain the institute (\$1M/year was considered too small. \$3M/year was OK)

The primary metric for the success of an institute is getting the results of its projects in use by real applications. Institutes need to consider how they are going to accomplish this. There was agreement among the attendees that libraries are the typical software test beds where new programming models and execution models are proved out. These same numerical and

communication libraries provide a fast vehicle for getting the new concepts into use by the application developers.

While libraries are often the testbeds for new programming models; MPI is the portable programming model today. For any new programming model to succeed it must be, at a minimum, as portable as MPI across the key HPC systems, clusters, and development platforms. If there is a programming model shift, tools to assist in the code transformations to new models and new algorithms are going to be critical in transitioning the millions of lines of code affected.

Key Topics and Themes for CS/Math Institutes

The following topics were identified as key challenges where synergistic math and computer science efforts are required to make an impact. These are examples of potential themes that CS/Math institute proposals could be built around.

1. **Effective use of multicore/manycore and heterogeneous architectures:** The emergence of multicore/manycore and heterogeneous architectures is probably the single biggest change in scalable computing in the past decade. Effective use of these new node architectures is critical to reducing the performance gap between the peak performance of the hardware and the realized performance of the applications. There are many issues that need to be addressed, but the most important include determining how to design high-performance software for new node architectures. Traditionally scalable applications have explicitly partitioned work, data and communication with MPI and have otherwise used simple sequential programming within a node. This approach is not sufficient for new node architectures. A successfully designed application for next-generation scalable systems must expose vectorizable loops and a large volume of fine-grain functional parallelism, and must do so in a portable way. Self-adapting and auto-tuning software can be an effective tool for performance portability, but current efforts need to expand their scope to take into account additional factors for easy-to-use code.
2. **Portable programming model and execution model for extreme scale architectures:** A potential institute theme is fostering the development of standard, portable and high performance programming models for manycore nodes. Current portable multicore programming models will not scale well to manycore and the best manycore programming models are presently proprietary. Related to this effort is a need for simulators and a study of programming environments and models, and how to restructure software for increased memory hierarchies.
3. **New multicore-friendly and multicore-aware algorithms:** Scalable multicore systems bring new computation/communication ratios. Within a node data transfers between cores is relatively inexpensive, but temporal affinity is still important for effective cache use. Across nodes, the relative cost of data transfer is growing very large. The development of new algorithms that take these issues into account can often perform very well, as do communication-avoiding algorithms that increase the computation-communication ratio or algorithms that support simultaneous computation-communication, or algorithms that vectorize well and have a large volume of functional parallelism.
4. **Exploiting mixed precision:** Although single precision computations have always enjoyed a performance advantage over double precision, new scalable multicore systems will compel us to perform as much computations in single precision as possible. Furthermore, as problem sizes increase, high-fidelity computations can benefit from judicious use of double-double or similar high precision computations. The formulation

and development of efficient and accurate mixed precision computations could be a potential theme for a CS/Math institute.

5. **New tools for efficient development of optimized code:** Because small changes in code organization can have large impact on performance, we need tools that allow a program to express generic optimization principles, such as loop unrolling, memory prefetch, etc., that can then be translated to specific expressions on a given processor architecture.
6. **Application resilience:** Because scalable multicore systems have a rapidly increasing component count and the feature size within these components is getting very small, the probability of system faults is predicted to rapidly increase to the point where faults will be continuous. Many faults can be detected and eliminated without intervention from the application. However, increasingly applications will need to recover from faults, including “soft faults” that result in incorrect data and computations. A sophisticated solution to application resilience will require a fully integrated approach from all research and development areas in scalable computing. Hardware components, OS and runtime libraries and the application itself will need to interact in order to have a fully resilient application. Furthermore, algorithms will need to be fault-resilient, being able to recover from a variety of system failures. Programming model support will be necessary to facilitate advanced fault-tolerance and recovery algorithms. Ideas include, algorithm sanity checks, mechanisms to express the relative importance of getting the right answer and automation of sanity checks. Finally, we need to continue improving checkpoint/restart capabilities, especially reducing the footprint and overhead associated with the checkpoint and restart process and performing these operations in a diskless environment.
7. **Sensitivity analysis:** Most areas of modeling and simulation are still pushing to reach high-fidelity solutions to a given set of input conditions, the so-called “forward problem.” However, as fidelity for a given class of problems improves, it becomes possible and imperative to study the sensitivity of the problem to parameter variability and uncertainty, and to seek an optimal solution over a range of parameter values. These types of advanced capabilities bring with them a natural resource of parallelism that can easily exceed the computing capability of our largest systems. The most basic form of these methods utilize a black-box approach with respect to the forward problem and can simultaneously run many instances of the application, leading to an embarrassingly parallel execution model. More advanced methods require a tightly-coupled aggregation of forward models, but still generate very large simultaneous systems that can execute very well in parallel. This is an important area for growth as high-fidelity simulations increase in number and computing capabilities grow.
8. **Multiscale/multiphysics modeling:** In addition to sensitivity analysis, another natural direction of effort that will lead to increasing problem sizes involves coupling and tighter integration of multiple length and time scales and multiple physics models into a single simulation package. Many areas of science require accurate modeling and simulation of couple physics and scales to accurately predict complex phenomena. This is not a new area and many efforts are already in place. At the same, any opportunities to standardize or push forward in this direction will increase the ability to solve critical national problems.
9. **Fast implicit solves:** Carefully analyzing complex problems, and adapting preconditioners to the underlying problem physics is how most of the progress in this area is being made. However, it is typically the case that advanced preconditioners are composed of standard algebraic components such as advanced multigrid/multilevel

methods, incomplete factorizations and basic smoothers. Furthermore, we need to renew our focus on basic iterative methods in an attempt to address bottlenecks due to collective operations (e.g., dot-products) and poor kernel performance. Emphasis on block methods, recycling methods, s-step like methods and mixed precision formulations will be necessary to address the next generation of problems.

10. **Advanced transient algorithms:** In many transient simulations, time to solution is increased because increased spatial fidelity requires extremely small time steps, not for accuracy reasons but for computational stability. To overcome the time step constraint problem, algorithm advances such as time parallel algorithms, multi-grid and multi-grid-like time algorithms and magneto-compressive wave formulations should be considered. Parallel time methods have been considered in the past with limited success, but there are recent efforts that show some promise. Magneto-compressive wave formulations circumvent the time step constraint by proposing new formulations without the constraint. As we increase spatial fidelity, these formulations, and similar ones in other modeling areas, should receive attention.
11. **Effective use of new and emerging memory systems:** For many science applications, memory system performance is the primary factor in determining performance. New generations of multicore nodes will present even more challenges. In particular, without effective memory utilization, we will be unable to effectively use all of the cores on a node. Since scalable multicore systems rely on effective use of all cores for full scalability, effective memory use is required if we are to reduce the performance gap.
12. **Algorithms for strong scaling:** As we perform computations on very large systems and push the limits of scalability we consistently see that our asymptotic performance is primarily determined by degradation at scale due to load imbalance exposed by synchronization. Research and development of new algorithms that are more tolerant of load imbalance, by reducing the number of synchronizations or by tolerating increased asynchronous computations, is important and will also address the need for faster forward solves.
13. **Advanced debugging capabilities:** Although there are certainly debugging tools that are useful, they are not scalable to the size of Leadership class computers. Additional efforts are needed to aid programmers in debugging at large scale and performance debugging tools are needed to study how memory utilization and access patterns are affecting performance. Presently there is little help for a programmer to determine if performance of a given code section is optimal.
14. **Application performance tuning and motifs:** Large-scale applications are often unwieldy and impractical to use in the context of performance analysis, the studies of new languages and other performance-related activities. Furthermore, across many applications, there are common execution and memory access patterns that, if addressed specifically, can impact a broad spectrum of codes. As a result, proxies such as motifs, mini-applications and related efforts can greatly benefit the science community. Current efforts in motifs need to be extended and the study of interoperable motifs is required. In addition, development of key mini-applications that represent the performance-determining computations of key application areas is important. Such mini-applications can be used to predict the performance of real applications in many speculative situations. Furthermore, mini-applications can be rewritten quickly to test new programming environments and models, something that is an essential pre-requisite to rewriting a large-scale application.
15. **Scalable computer system resource management tools:** Computer system resource tools must manage scheduling of resources efficiently. However, we have observed that

advances in discrete optimization algorithms have not been fully utilized in our own systems resource management tools. As systems go to 1 million cores, new discrete optimization methods for computer system resource management are required.

High Risk/High Impact Technologies Track

The High Risk track took a case study approach. With energy and climate being important national issues to address, the four application talks were on global climate change, fusion energy, nuclear energy (fission), and combustion. The workshop had a talk scheduled on nanotechnology and materials research but that speaker had travel problems and was unable to attend. Each of the speakers was asked to identify potential high risk/high impact projects in their respective areas of science. For this workshop high-risk / high-impact projects were defined to be those where success could provide large increment in scientific capability but the risk is considered too high for the applications community to undertake. Three categories of such projects were identified.

- Type 1: Well-characterized application of a new technology – risk comes from the need for hardened implementations of the technology and the need for a bridge between the apps domain and the experts in the new technology. Example: implementing an existing model in a new programming language or programming framework.
- Type 2: Well-established technologies applied to a new problem area – risk comes from whether the methodology can be successfully modified to meet problem-specific needs. Example: AMR for climate.
- Type 3: Fundamental new approaches, particularly in domains where there is little prior art in modeling. Example: uncertainty quantification for multiphysics applications.

These represent equally risky positions from the point of view of the applications scientist, but for different reasons. For Type 1, the risk is to a large extent institutional – will the new technology have sufficient support so that porting the application will be worth the domain scientist's time? Type 2 and Type 3 represent increasing degrees of technical risk. In the case of Type 2, there is some prior art to indicate that the advanced technology could work for the new science domain, whereas for Type 3, it is clear at the outset that new ideas will be required. For each of the case studies, we attempted to categorize potential projects in terms of the type of risk involved.

(1) The U.S. Fusion Simulation Project (FSP) (presented by John Cary). FSP has integrated modeling and simulation as principal goals. Cary also discussed issues in accelerator design including RF cavities and wakefield accelerators that could be designed and optimized using simulation. The following high impact projects were identified.

- Autotuning of Electromagnetic Particle in Cell (EM-PIC) applications.(Type 2). This includes the discovery of and optimization over parameters for existing data structures, and automatic restructuring of data to permit effective use of multicore. Also included in this area are GPU or cell-based EM-PIC.
- New approaches to EM (Type 2). These include new ADI approaches, and high-order embedded boundary methods. The current embedded-boundary approaches have been quite successful in computing both time-domain and frequency-domain problems, but there is still room for improvement in both the performance and accuracy.

- Co-development of concurrently parallel, multi-component/multi-physics components or simulation methods, including coupling (Type 3). This area consists of two distinct, but interrelated, problems. One is the lack of mathematically systematic approaches to the design of multiphysics codes corresponding to the coupling of diverse time scales, spatial scales, and underlying mathematical and computational representations of the physical processes to be coupled. A second issue is the software engineering of such complex multicomponent codes so that they will get high performance on high-end parallel platforms, while still having the flexibility to permit the required experimentation to different approaches to coupling.
- Systematic use of design optimization tools (Type 2). Computation is used extensively for design of devices in both fusion and accelerator modeling but there has been relatively little penetration of modern automatic design optimization tools into these fields. Such tools have the potential of greatly improving design practice.

An issue that cut across all the problems in this topic was that successful high impact teams must consist of hybrids from both sides: physicists with math / CS skills and math / CS people with physics understanding.

(2) Global climate modeling (presented by John Drake). He presented a number of high-impact problems in climate including: decadal prediction, water/carbon cycle, climate extremes, and abrupt climate change. High-risk projects to support these areas include:

- New dynamical core: adaptive mesh refinement, greater implicitness (Type 2). There is a great deal of experience in the use of adaptive mesh refinement in other fluid dynamics problems, so that the classification as Type 2 is justified. The use of adaptive methods would provide a possible approach to improved simulations in the water cycle (particularly of convective storms in the tropics), for extreme events, such as tropical cyclones; and for regional climate prediction. The need for performing long-time simulations (a century or more) places a premium on the development of more implicit methods.
- New data assimilation methods for oceans (Type 2) and carbon (Type 3). Data assimilation is a technique by which observational data, constrained by the dynamics of the simulation code, is used to generate initial data for calculating the solution to an initial-value problem. In order to use simulation to assess impacts of climate change on decadal time scales, it will be necessary to solve initial-value problems, for which data assimilation is essential. Data assimilation is a well-established methodology in numerical weather prediction, and the techniques developed there should apply directly to the atmospheric models in climate. Although is also a fluid system ocean differs sufficiently from the atmosphere in both its dynamics and the kind of data available to make extending assimilation to that case a Type 2 undertaking, while the models for carbon cycle (particularly sources and sinks) are substantially different from the fluids case to make assimilation for that case a Type 3 problem.
- New data analysis tools for rare / extreme events. (Type 3) Most of the data analyses performed for climate modeling have been in quantifying the variation in large-scale, time-averaged data, such as temperature and precipitation on global and regional scales. In some cases the key issue is the variation in the frequency of relatively rare events, such as the strongest (type 4 or type 5) tropical cyclones, the analysis of which would require new tools.

The long term goal and milestones are driven by the Intergovernmental Panel on Climate Change (IPCC) AR5 beginning in 2013. This imposes a deadline for successful completion of any high-risk project in this area. Also, successful high impact teams would need to be dedicated / motivated individuals eager to participate in climate community activities.

(3) Nuclear energy (presented by Andrew Siegel). Siegel's talk focused mostly on the technical issues surrounding simulation for the fast liquid metal breeder reactor. There are two overarching problems to be solved: improved operational efficiency / tolerances, and the design of passive safety features. Simulation could contribute to improvements in these areas in two ways: by lowering rule-of-thumb design margins for existing designs, or by uncovering design innovations with much better economics and / or safety.

Possible projects include:

- **Uncertainty quantification for Reynolds-averaged Navier-Stokes (RANS) CFD (Type 3).** RANS is a relatively low-fidelity model for turbulent flow that represents the time-varying solution as a time-averaged steady state. However, RANS is widely used in engineering design calculations for nuclear reactors, due to its fast time to solution. Being able to quantify the sensitivities of RANS simulations to variations in model parameters would greatly increase the reliability of such calculations.
- **Adaptive methods for neutron transport (space, energy, angle) (Type 2).** The complex dependence of energy and angle resolution required as a function of the material suggests that the use of adaptive methods would lead to substantial cost savings in neutronics calculations. Such work would represent a substantial extension of adaptive methods, which have focused mainly on adaptivity in physical space for such problems.
- **Component framework to unify safety and design codes in a single infrastructure. (Type 2).** There is a huge range of problems in the safety and design areas. However, each of these requires a subset of a small number of simulation components: complex geometries, neutron transport, conjugate heat transfer, structural deformation, and material properties. It would be highly desirable to refactor the safety and design code space so that the various simulation capabilities could be assembled from a common library implementing these components. This would greatly improve the economics of porting codes to new platforms and new architectures, and of developing new capabilities in response to new requirements.
- **Alternative methods for CFD in complex geometries. (Type 2).** One of the principal bottlenecks in simulation of nuclear reactors is grid generation for complex multimaterial configurations that requires considerable human intervention, with the process often taking more time than the simulation itself. New discretization approaches, such as generalizations of cut-cell methods or of overlapping grid methods that permit accurate representation of material interfaces, have the potential of reducing the grid-generation time to a small fraction of what it currently takes, as well as making it a completely automated process.
- **Time accurate coupling methods for fast transients. (Type 2)** While steady RANS calculations have been the mainstay of fluids modeling for nuclear reactors, there are a number of problems, such as transient heat transfer issues related to safety, that require a high-resolution, time-dependent computational capability. The development of capabilities that can handle the range of flow and heat transfer regimes that arise in these problems is in its infancy, and more experimentation with different approaches to solving these problems is needed. In addition, such capabilities could also be used to calibrate RANS calculations, or could be used in a hybrid mode, in which time-accurate methods and RANS are used in different locations depending of conditions.

(4) Combustion (presented by John Bell). Bell identified a large number of high-impact problems including: turbines for stationary power generation, and I/C engines for transportation. Research in new fuels (hydrogen, syngas; biodeisel) require fundamental changes in the science base with large-scale, high-fidelity simulations of turbulent combustion at the center. Successful high risk teams would be focused on combustion applications and active in collaborations with experimentalists, chemical kinetics experts and transport experts. Possible high risk projects include:

- Integrate more realistic laboratory scale simulation / experiments with kinetic mechanisms development (Type 3). For hydrocarbon fuels, turbulent flames are sufficiently close to being locally one-dimensional that the development of chemical kinetics mechanisms can be performed based on simple one-dimensional experimental configurations (e.g. laminar flames), with the match of kinetics data to experimental flame data mediated through the corresponding one-dimensional simulations. These are then used to derive an engineering model for turbulent combustion For more heavily hydrogenated fuels, turbulent flames are no longer locally one-dimensional, and a feedback loop involving high-fidelity multidimensional turbulent combustion calculations, multidimensional turbulent combustion experiments, and kinetics and transport mechanism development, is required. To carry this out, new computational sensitivity and optimization techniques will be required to develop hierarchy of kinetics / transport models for applications with linkage back to mechanism development.
- Gas-phase simulations with high-fidelity kinetics for engineering scale devices. This require a combining several technologies that are already under development or are quite mature: computational methods for complex geometries and appropriate adaptivity; and capabilities for low-Mach number reacting flows for lean premixed combustion that interact with acoustics.
- Predict emissions from bio-diesel fuels in I/C engines (Type 3). This requires CFD in moving geometries; high-fidelity multiphase reacting flow capabilities; and a hierarchy of tractable high-pressure kinetics models. The high risk comes from the need for appropriate multiphase flow and kinetics models, which do not yet exist.
- High-order methods for multi-physics combustion problems (Type 2). The goal here is to combine the high-order approach that has been used very successfully for fully compressible flow on uniform grids, with the low-Mach number adaptive methods for which offer great advantages in computational efficiency. Such an approach would require appropriate versions of high-order spatial and temporal discretizations, robust discretization methods for highly convective flows and coupled solvers, and high-order extensions of operator splitting and semi-implicit methods. All of these techniques are currently under development.

Industrial impact an essential long-term goal for combustion (and also for the other energy topics).

Cross-Cutting Issues

The High Risk track identified several cross-cutting dependencies, some of which are potential CS/Math Institutes themes. First the workshop attendees recognized that the DOE SciDAC program has been an important source of success in computational science and that any high risk efforts will assume and depend on the SciDAC Centers and Institutes continuing to support their existing tools, and to provide new tools and techniques. Some other cross-cutting requirements were:

- Petascale data analysis infrastructure and tools for analysis of large data sets with complex phase space representations.
- Robust and fast parallel I/O
- Program language support and kernel library support for multi-core/NUMA nodes
- Rapid prototyping tools; e.g. matlab for HPC
- Load balancing for large machines, new architectures

Size and duration of projects were a source of concern to the workshop participants. Project sizes were typically scoped at 3-5 FTE, with five years duration required to obtain results with real applications impact. Finally, we wish to emphasize that the choice of four fields were meant to be case studies, not an exhaustive list of potential high-risk projects. It was the view of the workshop participants that there are many other applications and potential projects that would fit into the template that emerged from our discussions.

Attached as appendices are the agenda and list of workshop attendees.

Agenda

Workshop on CS/Math Institutes and High Risk / High Payoff Technologies for Applications

Tuesday, October 7, 2008

- 8:00 a.m. Breakfast Hotel Restaurant
- 9:00 a.m. Welcome and thanks Al Geist
- 9:10 a.m. Workshop Purpose and Expected Results..... Dan Hitchcock
- 9:40 a.m. Questions and Answers about Scope and Goals of workshop

10:00 a.m. **Break and split into two tracks**
 Concord A&B Rooms United Airlines Room 3rd
 floor

	Math/CS Institutes Chair: Mike Heroux		High Risk – High Payoff Chair: Phil Colella
10:15 a.m.	Math Libs – Jack Dongarra		Fusion Opportunities – John Cary
11:00 a.m.	Answer PPT questions Relative to this topic		Answer PPT questions relative to these opportunities
12:30	Lunch (provided)		
1:30 p.m.	App Requirements – Trey White		Climate Opportunities – John Drake
2:15 p.m.	Answer PPT questions Relative to this topic		Answer PPT questions relative to these opportunities
3:45 p.m.	Break		
4:00 p.m.	Continue working on questions from morning and afternoon		Continue working on questions from morning and afternoon
5:00 p.m.	Adjourn for day		
	Dinner on own – (go somewhere as group)		

Wednesday, October 8, 2008

8:00 a.m. Breakfast Hotel Restaurant

9:00 a.m. **Dual Tracks Resume with new topics and application areas**

	Math/CS Institutes Chair: Mike Heroux		High Risk – High Payoff Chair: Phil Colella
9:15 a.m.	AMR – Brian van Straalen		Combustion – John Bell
10:00 a.m.	Answer PPT questions Relative to this topic		Answer PPT questions relative to these opportunities
11:30	Lunch (provided)		
12:30 p.m.	PDE Libs – Barry Smith		Nuclear – Andrew Siegel
1:15 p.m.	Answer PPT questions Relative to this topic		Answer PPT questions relative to these opportunities
2:45 p.m.	Break		
3:00 p.m.	Create summary slides/ write-up of Institutes track		Create summary slides/ write-up of High Risk track
5:00 p.m.	Adjourn for day		
	Dinner on own – (go somewhere as group)		

Thursday, October 9, 2008

8:00 a.m. Breakfast Hotel Restaurant

9:00 a.m. Summary of CS/Institutes track tbd

9:30 a.m. Discussion of findings by all attendees

10:00 a.m. **Break**

10:30 a.m. Summary of High Risk- High Payoff technologies tbd

11:00 a.m. Discussion of findings by all attendees

11:30 a.m. **Workshop ends**

Attendees

Dan Hitchcock	ASCR
Steve Lee	LLNL
Stephen Scott	ORNL
Sandy Landsberg	ASCR
Al Geist	ORNL
Phil Colella	LBL
Mike Heroux	SNL
Jack Dongarra	UTK/ORNL
Trey White	ORNL
Barry Smith	ANL
Brian van Straalen	LBL
Kathy Yelick	LBL
Bill Gropp	U. Ill
Rich Graham	ORNL
Rusty Lusk	ANL
Ron Brightwell	SNL
Bronis de Supinski	LLNL
David Bailey	LBL
Paul Fischer	ANL
John Drake	ORNL
Andrew Siegel	ANL
John Cary	U. Colorado
Lori Diachin	LLNL
John Bell	LBL
David Bernholdt	ORNL
Erik Boman	SNL
Eric Phipps	SNL
John Shadid	SNL
Jeff Vetter	ORNL