SAND2007-4518P

CIS External Review 2006
Overview of Algorithms and Enabling Technology

1. Introduction: Algorithms and enabling technology (A&ET) is a key part of the CIS efforts to develop and support a high performance modeling and simulation (M&S) capability at Sandia National Laboratories. This area complements the other areas of CIS investment, Applications and Systems, and these three areas together form a coordinated, strategic approach towards developing an effective M&S capability.

There are several factors that determine the specific investment areas within A&ET areas and levels of investment. The most important of these is the Laboratories' mission. The specific needs are determined based on close interactions and collaborations with M&S applications. Other factors include the interests and capabilities of the staff and the level of funding support that is available. CIS staff play a large role in identifying funding sources and obtaining this funding through proposals or other mechanisms. Funding areas for A&ET are covered in section 3.

Even though specific needs are determined based on interactions and collaborations, especially with the code teams, there is both a "pull" and a "push" aspect to our algorithms and enabling technology work. A pull occurs when a user (e.g., an analyst) requests a specific capability or improvement. For example, we get many requests for improved preconditioners for solvers or dynamic load balancing capabilities. Much of this work is funded through the Advanced Simulation and Computing (ASC) program of the DOE/NNSA and is delivered through libraries to users.

There is also a strong "push" aspect to our work. Most of the algorithms and enabling technology researchers work very closely with users and, based on their understanding of the applications and the needs, are able to determine promising research areas and directions. An historical example of this would be multigrid preconditioners, which after more than ten years of investment are having a huge impact on current M&S capabilities.

Another area of "push" that is now having a big impact is discrete mathematics, graph theory and informatics. Sandia has maintained a research effort in these areas for many years, funded by MICS, CSRF and LDRD (see Section 3, below). This research effort has been somewhat larger than in other DOE laboratories, and has found applications in such areas as infrastructure analysis, secure transportation, and manufacturing. It has also positioned us to contribute in the areas of intelligence and homeland security.

2. Technical Investment areas. The investment areas for A&ET are listed below. Descriptions are brief and focus on future investments. A few accomplishments for the past year are listed; however, details of these accomplishments are left to the individual technical presentations. The following discussion covers all A&ET investment areas, not just the areas presented in the review this year.

2.1 Integrated Capabilities: The group is making a strong effort to integrate its capabilities in solvers and optimization and to provide this capability to applications.

Within the ASC program A&ET has committed to an FY07 milestone that will design, implement, document, deploy and test an integrated solver/optimization capability based on the Belos (block linear solvers), Anasazi (block eigensolvers), Rythmos (time integrators) and MOOCHO (optimization) packages using the Thyra abstract interfaces in Trilinos. These packages provide state-of-the-art algorithm implementations using a state-of-the-art software design and implementation. Along with existing Trilinos packages, this work will represent a fully integrated solution capability from basic linear algebra, linear solver, nonlinear solvers, time integrators to simulation-constrained optimization, where all packages are fully interoperable and a full suite of internal and external solvers can be brought to bear on our most challenging simulation problems and improve our design through optimization capabilities.

2.2. Solvers. This has traditionally been a strong investment area for Sandia. Solvers are the key to most scientific modeling and simulation efforts and often more than one half of the computational time is spent in the solvers. Most of the solver research is driven by finite element simulation codes which are being developed within the ASC program, although applications such as optimization and circuit simulation are also being addressed.

In the area of linear solvers, we have investments in traditional algebraic preconditioners. These preconditioners are typically robust and can be used in a black box fashion by many applications. However, performance is often problem dependent and we maintain an effort to improve the efficiency of these solvers as codes and models become more complex, for example by using segregated preconditioners. We are also investing in block solvers for systems with multiple right hand sides. Block solvers have the potential to improve the single-processor performance, the parallel efficiency and the algorithmic efficiency.

We also continue to invest in multilevel preconditioners, particularly for problems arising from unstructured grid, finite element simulations. These solvers provide the best hope for algorithmic scalability, although performance can be very problem-dependent. Multilevel solvers have been successfully used in many applications, including high-energy density physics, compressible fluid dynamics, and device modeling.

We also conduct research into nonlinear solvers and time integrators and are particularly interested in the interactions and potential synergy between linear and nonlinear solvers and time integrators. We are also investigating the effects of operator splitting and the potential loss of accuracy if the time-step and the spatial resolution are not consistent.

Our ongoing research in time integration methods includes the development of a second-order operator splitting advancement algorithm in the Alegra code for HEDP simulations. A significant software development effort in time integration is being designed to include sophisticated features to enable adjoint calculations, and includes a research component in global error estimation.

Our nonlinear solver effort is focusing on algorithms and software for coupling distinct application codes (such as those in the Sierra suite) so that the robustness and accuracy are much better then one gets from loose-coupling approaches. Our stability and bifurcation analysis capability for large-scale applications is still unique in the world.

Our solvers are delivered to codes through the Trilinos framework. Trilinos provides common interfaces both the users and to other algorithm developers as well as optimized code for common operations, such as matrix-vector multiplies. It also provides a structure that adheres to standards for software quality engineering, such as requirements planning, unit testing, regression testing, version control and bug tracking. Trilinos is widely used both within the DOE laboratories and in the external computing community.

2.3. Finite element methods and numerical analysis. High performance computer simulations based on PDE models play fundamental role in key Sandia applications. Among the many components of this technology, discretization establishes the essential link between continuum models and their algebraic representations. The algebraic models must provide accurate and physically correct solutions so as to enable validated computer simulations of the actual process. Design of such models requires careful analysis of the attending mathematical structure of the PDE model. This structure encodes intrinsic physical properties of the process that is being modeled, such as conservation laws, solution symmetries, positivity, and maximum principles, to name just few.

Our research is driven by the emerging need for computer simulation of processes involving multiple physics and multiple scales. Coexistence of different physics and/or scales in the simulation will necessarily require couplings between different types of FE models. Even with the advent of petaflops computing it is not realistic to expect that complete resolution of all relevant scales will be possible, neither is it safe to assume that a completely compatible discretization will be feasible in all cases. Thus, our main goal is to provide mathematical analysis in support of finite element modeling that addresses these needs by focusing on both compatible and regularized FE models.

We have begun an effort to provide a package within the Trilinos framework that will provide consistent discretizations for accurate simulations. This work will be presented as part of this year's CIS review.

2.3. Optimization. Optimization is a major thrust area for the algorithms and enabling technology effort. The Design and Optimization project focuses on three interrelated areas: (1) the development of fundamental optimization algorithms, (2) the development of frameworks for optimization, and (3) the development of tools and techniques for the quantification of uncertainty.

Optimization methods apply to any setting where one must choose a best solution from among a set of feasible solutions that satisfy a set of system constraints. The scientific community is increasingly recognizing optimization technology as a fundamental technique for exploring and predicting the behavior of complex natural and engineering systems. Optimization methods interface with and complement computational models to provide a powerful mechanism to explore "what-if" scenarios.

Optimization has already been applied in a wide range of DOE problems, including stockpile stewardship, nonproliferation and treaty verification, environment and waste management, and general policy decision support. Optimization is a key technology for logistics and resource management within DOE where methods for large-scale resource-constrained project scheduling have been effectively developed for the Pantex facility and the office of secure transportation. Optimization has also become a major tool for the ASC program's evaluation of weapons components. A new focus of our discrete optimization is logistics, and we have recently begun a significant program in this area, and we will most likely be presenting this area next year.

The CIS maintains a significant effort in the area of large-scale optimization. Areas of work include surrogate-based optimization, simultaneous analysis and design, and optimization under uncertainty (OUU). Surrogate-based optimization uses computationally inexpensive models of the objective and constraint functions to guide the optimization process. We are also beginning research efforts in the area of reduced-order modeling.

OUU involves a computationally expensive nesting of optimization and uncertainty quantification techniques. For OUU to be practical with ASC-scale simulations, this expense must be mitigated, for example through the use of surrogate-based approaches or through the exploitation of structure in analytic reliability uncertainty quantification methods. The algorithmic work in this area is delivered through the DAKOTA optimization and UQ framework.

Sandia also has unique capabilities in combinatorial and global optimization within DOE. Problems with explicit combinatorial structure often require discrete analysis to develop efficient solutions. We are developing techniques for specific applications like sensor placement in water networks and node allocation in parallel clusters. Additionally, we are developing general techniques like graph algorithms, online algorithms and polyhedral combinatorics. We are also developing global search methods for applications in which (near-)optimal solutions cannot be efficiently identified. We are developing parallel global optimizers that rigorously guarantee that a globally optimal solution is generated (e.g. branch and bound), as well as practical heuristics that balance global searching with efficient generation of near-optimal solutions

Our staff working in optimization are well-positioned to contribute to the new thrust in homeland security. We have begun work with the EPA in the placement of sensors in water networks to quickly identify and control the flow of introduced chemical or biological agents. We have also initiated homeland security projects in chemical and biological source inversion.

2.4. Informatics: The field of informatics has recently become a key part of the work done in the CIS. There are many reasons for this rapid rise. The primary driver has been the recent complete adoption of electronic databases for essentially all information storage, and the associated technological advances that have come in making this information easily accessible. At the same time, advances in automation and miniaturization have allowed high throughput experiments to be undertaken that generate enormous amounts of information, e.g., in biology and in the intelligence community. Finally, there have been many recent practical societal concerns that have driven a need to find subtle patterns hidden in large databases of existing knowledge. All of these different drivers are helping to lead Sandia towards more extensive investments in the area of informatics.

One important theme in our informatics accomplishments is the implementation of efficient, scaleable algorithms, leveraging our experience of high-performance computing. We have developed different scaleable algorithms to predict protein-protein interaction using only amino acid sequence for the DOE Genomes-to-Life project. We have also developed techniques to sort through tens of thousands of genes for groups of genes related to various clinical questions in cancer research, including distinguishing between two difference types of childhood leukemia. The concept of "clustering" data is a very important to informatics, and we have borrowed techniques from graph drawing to cluster data sets having many million points.

We continue to develop graph-based algorithms and tools to "process" large amounts of intelligence data, including static and dynamic clustering algorithms. The use of these tools is also being evaluated on different computer architectures.

2.5. Dynamic load balancing. Dynamic load balancing remains a critical area of research at Sandia. Work in the Zoltan project includes research to support emerging non-traditional applications as well as ASC customers. In addition, meta-partitioning research for unstructured adaptive mesh refinement (AMR) applications has begun, using Zoltan as a source of multiple partitioning algorithms.

Performance improvements and ease-of-use in ASC adaptive finite element codes and contact simulations motivated much of Zoltan's ASC development. Zoltan's degenerate geometry detection allows more natural decompositions for nearly 2D geometries (e.g., silicon wafers), regardless of their spatial orientation. Multicriteria geometric partitioning generates a single decomposition that is balanced with respect to more than one weighting factor (e.g., for multiphase simulations). And Zoltan's new C++ interface simplifies use of Zoltan in C++ applications.

Our current plans include linking the Zoltan load balancing package with the Trilinos framework.

2.6. Miscellaneous. There are some smaller efforts that do not fit into one of the categories above, but yet are a critical part of the CIS capabilities. These include work in graph theory, which support a broad range of project such as dynamic load balancing, optimization and solvers, classical applied mathematics, which includes work in mathematical modeling, asymptotics and stability and stability analysis which is a critical capability for Sandia.

3. Programs. Many programs provide funding to our algorithms and enabling technologies efforts. Each has its own objectives, constraints and funding mechanisms. These programs are summarized briefly below.

3.1. ASCR. Sandia receives significant research funding from the Office of Advanced Scientific Computing Research (ASCR) program in DOE's office of Science in several research areas. The funding levels given below are for all of Sandia, not just for the CIS.

- The mathematics base program is the flagship effort in this program, which focuses on mathematics and scientific computing research. Projects in this program are reviewed/selected internally (with ASCR concurrence) and peer-reviewed externally once every three years.
- The Scientific Discovery through Advance Computing (SciDAC) initiative funds multi-year, multi-institution projects in a range of areas. Sandia receives funding for projects in meshing, common component architectures and operating systems.
- The ASCR office funds peer-reviewed visualization projects for one to three years through this program. Sandia currently has one project funded.
- Sandia has projects focusing on light-weight kernel operating systems and run-time systems. They currently receive a total of $1.1M from the base program and SciDAC. This project will be presented in the Systems Session in this review.

The ASCR funding growth over the past couple of years has been in the SciDAC and SciDAC 2 program, and there has been some growth in the Sandia funding from these programs. We expect that there will continue to be growth in ASCR funding, although this growth will be tied to new initiatives in multiscale mathematics, petascale data and complex systems.

3.2. ASC. The Advanced Simulation and Computing Program is DOE's program to develop and apply advanced modeling and simulation capabilities in support of the nation's nuclear weapons stockpile. ASC includes an $7M Algorithms effort to develop and deploy advanced algorithms and enabling technologies to codes being developed within ASC. Projects are selected by program managers within the ASC applications program based on prioritized strategic needs and reviewed annually. We have reorganized the ASC algorithms program into two main areas

- Solvers and parallel algorithms. This subproject develops linear and nonlinear solvers, time integrators, finite elements methods, and interface tracking algorithms. This project also develops the Trilinos/Petra solver framework, which is the delivery vehicle for all codes developed within this subproject. This will also most likely include work in dynamic load balancing, adaptivity and contact algorithms.
- Optimization and uncertainty quantification. This will include our efforts in Verification and Validation, and uncertainty quantification. This subproject will also develop optimization tools as well as the DAKOTA optimization framework.

DAKOTA is currently used as a framework for Sandia's uncertainty quantification efforts.

ASC funding for algorithms decreased slightly in FY07; however, we are expecting significant cuts in FY08. This budget is ultimately tied to nuclear weapons (NW) funding levels for the laboratories.

3.3. CSRF. The Computer Science Research Foundation program funds research into computer science, computational science, and enabling technologies in support of Sandia's high performance computing effort. This was recently reorganized from approximately 30 small projects into 4 larger projects. These four projects are focused around the following areas with the approximate (yet to be determined) levels of funding

- Simulation Computing ($3.5M)
- Next-Generation Systems ($2.5M)
- Information Sciences ($1M)
- Disruptive Technologies ($1M)

CSRF is formally part of the ASC program and supports research that will impact the nuclear weapons program. Consequently, these projects are typically more applied than LDRD projects. CSRF funding decreased in FY07 and continues to be under pressure with the expected decline in NW funding levels. We are making a strong effort to protect this program and the ASC Algorithms program as NW budgets decline.

3.4. LDRD. Sandia's Laboratory-Directed Research and Development program is funded through a 6% tax on laboratory direct spending. Sandia invested approximately $3.5M in its Computer and Information Sciences Investment Area in FY06, but there will be a major laboratory reorganization of the LDRD program in FY07. The CIS investment area has been combined with the Engineering Sciences investment area for a larger area entitled "Enabling Predictive Sciences". So far, this has not caused large changes in the overall investment in Computing. Projects are selected based on an internal review and are reviewed annually. The level of funding in the Computer Science investment area is set by Sandia's vice presidents. Funding seems to have leveled off.

3.5. CSRI. The Computer Science Research Institute funds external collaborations in mathematics and computer science. External collaborations are important for the health of a research program for many reasons. For example, researchers are more aware of releated research in the areas of interest, research by external collaborators is highly leveraged and focused directly on Sandia problems, and external collaborations are an important aspect of recruiting. The CSRI funds on-site projects, short-term visits, long-term visits, including sabbaticals, summer students and faculty, workshops, and fellowship through NPSC and HPCS. The CSRI does not fund Sandia staff. The CSRI received $1.5M in FY07, but is not expected to receive any funding in FY08.

3.6. Miscellaneous funding sources. The programs summarized above provide most of the funding for the A&ET work in the CIS; however, there are some programs that provide smaller amounts of funding. These sources include the DoD (Department of Defense), DHS (Department of Homeland Security), NIH (National Institutes of Health), EPA (Environmental Protection Agency) and DARPA (Defense Advanced Research

Project Agency). Also, some projects receive funding directly from other customers within DOE. We have been successful in getting new money from DHS, MICS, and other government agencies.

4. Funding. The following table shows the approximate amount of funding in each of the technical investment areas and funding areas for FY07. All funding is for staff labor. The approximate conversion is $275K/staff member.

	ASCR	ASC	CSRF	LDRD	Misc.	Total
Solvers	$650K	$1,700K	$100K	$450K	$100K	$3,000K
FEM & NA	$450K	$1,800K	$1,300K	$300K	$300K	$4,150K
Optimization, V&V and UQ	$400K	$2,900K	$1,650K	$800K	$450K	$6,200K
Informatics	$150K	$1,100K	$950K	$1,500K	$700K	$4,400K
Discrete math and load balancing	$1,100K	$500K	$200K	$200K	$1,550K	$3,550K
Miscellaneous		$500K	$200K	$900K	$1,000K	$2,600K
Total	$2,750K	$8,500K	$4,400K	$4,150K	$4,100K	$23,900K