

Lawrence Berkeley National Laboratory

Lawrence Berkeley National Laboratory

Title

An expanded framework for the advanced computational testing and simulation toolkit

Permalink

<https://escholarship.org/uc/item/72c8p2p2>

Authors

Marques, Osni A.
Drummond, Leroy A.

Publication Date

2003-11-09

**An Expanded Framework for the Advanced Computational
Testing and Simulation Toolkit**

Osni A. Marques and Leroy A. Drummond
Lawrence Berkeley National Laboratory
One Cyclotron Road, MS 50F-1650
Berkeley, CA 94720-8139
[oamarques,ladrummond]@lbl.gov

Table of Contents

Preface	v
Abstract	vii
1. Introduction	1
2. Accomplishments	5
3. Lessons Learned	6
4. Related Activities	7
5. Agenda	9
5.1. <i>A Solid Base for the ACTS Toolkit</i>	9
5.2. <i>ACTS High Quality Software Certification</i>	10
5.3. <i>Inclusion of New Solutions to Computational Problems</i>	11
5.4. <i>Interoperability and Software Distribution</i>	11
5.5. <i>Active Collaboration with DOE Computer Facilities</i>	12
5.6. <i>Working in Collaboration with Other Software Initiatives</i>	13
5.7. <i>Encourage Feedback from Users</i>	13
5.8. <i>Further Development of the Tools</i>	14
5.9. <i>Expertise Tracking</i>	14
5.10. <i>Increasing the Visibility of the ACTS Information Center</i>	14
5.11. <i>Dissemination Plans</i>	15
6. Synergism of the Proposed Work	16
7. Benefits and Merits of this Proposed Work	17
8. Deliverables	19
Appendix A - Tool Summaries	29
A.1. <i>Numerical Tools</i>	29
A.2. <i>Tools That Support Code Development</i>	31

<i>A.3. Tools That Support Code Execution</i>	33
<i>A.4. Tools That Support Library Development</i>	35
Appendix B - Accomplishments (2000 and 2001)	37
<i>B.1. Software Installation and/or Updating on NERSC Platforms</i>	37
<i>B.2. Presentations Related to the ACTS Toolkit</i>	37
<i>B.3. Workshops on the ACTS Toolkit</i>	38
<i>B.4. Synergetic Activities</i>	38
<i>B.5. Activities Spurred at Other Centers</i>	39
Appendix C - Data from ACTS Toolkit Workshops	40
<i>Workshop I - How Can ACTS Work for You?</i>	40
<i>Workshop II - Solving Problems in Science and Engineering</i>	44
<i>Workshop III - High Performance Numerical Libraries for Science and Engineering</i>	48
Appendix D - Early Interactions with DOE Projects and Others	51
<i>D.1. DOE Projects</i>	51
<i>D.2. DOE Institutions</i>	54
<i>D.3. Other Projects</i>	54
<i>D.4. Other Institutions</i>	55
<i>D.5. Other Supporters of the Proposed Work</i>	57
<i>D.6. Interactions with Computer Vendors</i>	58
References	59

Preface

This document is an abridged version of a proposal submitted to the US Department of Energy's Mathematical, Information, and Computational Sciences (MICS) Program.

The *Advanced Computational Testing Software* (ACTS) Toolkit, which was part of the DOE 2000 Program, has evolved into the DOE *Advanced CompuTational Software* (ACTS) Collection. This evolution follows a reformulation of strategies and goals of the original project, as a means to provide a prompt response to advances in hardware, software development practices, user requirements, and important problems in computational sciences. Furthermore, the proposal *An Expanded Framework for the Advanced Computational Software Collection* is a description of such evolution. The new name, DOE ACTS Collection, provides for a more dynamic view of the proposed software framework, with well-defined criteria for the inclusion and retirement of tools. The acronym ACTS was preserved to stress a sense of continuation and long term software support among the ACTS users and the Computational Science Community.

We invite the reader to consult the documents available in <http://acts.nersc.gov/documents> for up-to-date information on current developments in the project, accomplishments and adjustments in plans.

Abstract

During the past decades there has been a continuous growth in the number of physical and societal problems that have been successfully studied and solved by means of computational modeling and simulation. Distinctively, a number of these are important scientific problems ranging in scale from the atomic to the cosmic. For example, ionization is a phenomenon as ubiquitous in modern society as the glow of fluorescent lights and the etching on silicon computer chips; but it was not until 1999 that researchers finally achieved a complete numerical solution to the simplest example of ionization, the collision of a hydrogen atom with an electron. On the opposite scale, cosmologists have long wondered whether the expansion of the Universe, which began with the Big Bang, would ever reverse itself, ending the Universe in a Big Crunch. In 2000, analysis of new measurements of the cosmic microwave background radiation showed that the geometry of the Universe is flat, and thus the Universe will continue expanding forever.

Both of these discoveries depended on high performance computer simulations that utilized computational tools included in the Advanced Computational Testing and Simulation (ACTS) Toolkit. The ACTS Toolkit is an umbrella project that brought together a number of general-purpose computational tool development projects funded and supported by the U.S. Department of Energy (DOE). These tools, which have been developed independently, mainly at DOE laboratories, make it easier for scientific code developers to write high performance applications for parallel computers. They tackle a number of computational issues that are common to a large number of scientific applications, mainly implementation of numerical algorithms, and support for code development, execution and optimization. The ACTS Toolkit Project enables the use of these tools by a much wider community of computational scientists, and promotes code portability, reusability, reduction of duplicate efforts, and tool maturity.

By analogy with a commercial company, the ACTS Project has put the formerly uncoordinated research and development performed at various places into one distribution channel. ACTS complements the R&D effort by adding technical support, quality assurance, and marketing. Specifically, ACTS has provided technical support, tool information and evaluations. Consequently, it has increased the availability and stability of the tools; and has educated scientists about the Toolkit through presentations and workshops. These efforts have resulted in increased use of the tools in significant scientific applications. Feedback from users has resulted in improvements in the tools.

Along with these successes comes the awareness of how much ongoing work remains to be done. There is still a gap between tool developers and application developers which leads to duplication of efforts. Users demand long-term support of the tools, and the tools evolve or are superseded by other tools. There is a demand for tool interoperability and more uniformity in the documentation and user interfaces. And there is a need for an intelligent and dynamic repository of high performance tools. In short, there is a need for a collaboration framework for ongoing development and deployment of computational tools.

This proposal addresses these and other issues by outlining an agenda for expanding the scope of the ACTS Toolkit Project based on lessons learned from current activities. Highlights of this agenda include peer-reviewed certification of new tools; finding tools to solve problems that are not currently addressed by the Toolkit; working in collaboration with other software initiatives and DOE computer facilities; expanding outreach efforts; promoting interoperability, further development of the tools; and improving functionality of the ACTS Information Center, among other tasks. The ultimate goal is to make the ACTS tools more widely used and more effective in solving DOE's and the nation's scientific problems through the creation of a reliable software infrastructure for scientific computing.

1. Introduction

The ACTS Toolkit is a set of computational tools developed primarily at DOE laboratories and is aimed at simplifying the solution of common and important computational problems. The use of the tools reduces the development time for new codes and the tools provide functionality that might not otherwise be available. All this potential cannot be achieved, however, if the tools are not used effectively or not used at all.

The ACTS Project brings together software development from DOE laboratories, and in some cases in collaboration with universities, to the hands of a scientific community that is broader than the community in a single DOE laboratory. Consequently, ACTS has promoted reusability over duplication of efforts and tool interoperability over “hard-wired” and intrusive library interfaces that only serve a specific project. By analogy with a commercial company, ACTS has put the formerly uncoordinated research and development performed at various places into one distribution channel. ACTS complements DOE research and development efforts by adding technical support, quality assurance and marketing.

ACTS is bringing tools to higher acceptance levels among computational scientists and institutions. This provides the necessary means for individual tool projects to interact with more users, so the tools can gradually mature, becoming both more robust and portable to state-of-the-art high performance computing environments. The following are examples of remarkable scientific applications that have greatly benefited from the use of ACTS tools and that have achieved results that could not have been realized otherwise:

- Collaborators at the Lawrence Berkeley National Laboratory (LBNL), Lawrence Livermore National Laboratory (LLNL), and the University of California at Davis have obtained a complete solution of the ionization of a hydrogen atom by collision with an electron, the simplest nontrivial example of the long-standing unsolved problem of scattering in a quantum system of three charge particles [1]. In this application, SuperLU was used to construct preconditioners for the solution of unsymmetric linear systems of order up to 1.8 million. SuperLU is a general-purpose library available in the ACTS Toolkit that provides for the direct solution of large, sparse, nonsymmetric systems of linear equations on high performance machines.
- On April 26, 2000, the international BOOMERANG (Balloon Observations of Millimetric Extragalactic Radiation and Geophysics) consortium, led by Andrew Lange of the California Institute of Technology and Paolo de Bernardis of Università di Roma La Sapienza, announced results of a very detailed measurement of the cosmic microwave background radiation (CMB). These results revealed that the curvature of the Universe is not positive or negative but flat [2]. Much of the data analysis was performed by Julian Borrill (LBNL) employing a software called MADCAP (Microwave Anisotropy Dataset Computational Analysis Package [14]) which uses several routines from ScaLAPACK. ScaLAPACK is one of the tools available in the ACTS Toolkit intended for linear algebra calculations.

- The parallel version of M3D, a multi-level 3D plasma physics code developed at the Princeton Plasma Physics Laboratory (PPPL), makes extensive use of PETSc for parallelization and solution of an unstructured mesh problem. One of the members of the PPPL team working on the parallelization of M3D became familiarized with PETSc after attending the ACTS Toolkit workshop held at LBNL on September 27-30, 2000 (see Appendix C). Recently, this scientist stated that the parallelization of M3D “would have been very difficult without PETSc, and would have required several physicists to spend a significant amount of time reinventing numerical algorithms instead of doing physics.”

More details on these codes and additional examples of scientific applications using tools from the ACTS Toolkit can be found in the ACTS Information Center at <http://acts.nersc.gov>. The tools currently in the ACTS Toolkit and their functionalities are listed in Table 1 (for more information on individual tools, their current status and our comments on the tools, see Appendix A).

The DOE’s National Energy Research Scientific Computing Center (NERSC [21]) has developed an infrastructure to enable and facilitate scientific simulations on its high-end computers. Making ACTS tools more effective in solving DOE scientific problems matches NERSC's goals of making top quality computing services available to its users. ACTS tools fall into a category NERSC labels “experimental software” (ES); this is software from the research community that is not supported by vendors. NERSC support for ES is generally very limited in scope because of issues related to long-term tool maintenance and development. Our role in supporting the ACTS Toolkit has substantially increased ES support to NERSC users.

The unique scenario provided by the ACTS Project facilitates a number of opportunities to take a leap forward in the use and acceptance of ES. We have worked on improving the usability, accessibility, and acceptance of ACTS tools and, in particular, provided the following services:

- *Technical support.* Currently, only a few of the ACTS tool development projects have established an infrastructure to support their users. We have been working in creating a uniform infrastructure to provide comparable support levels for all the ACTS tools. We have installed, maintained, and supported a subset of the ACTS tools on NERSC platforms based on requests from users and availability (i.e., there is a version of the software available for the target platform). This support has been coordinated with the tool developers and also extended to users outside the NERSC community wanting to install tools on their own computer systems. Two workshops have been organized at LBNL, and tutorials on the tools have been presented at several venues around the theme of computational sciences (see Appendix B for detailed reports on these activities). The centerpiece of our support effort is the ACTS Information Center (<http://acts.nersc.gov>). This site contains overviews, introductory information about each tool, preliminary evaluations, and instructions on where to get software and technical assistance. This website has been continuously improved and updated.
- *Evaluations.* We have performed software evaluations of tools on NERSC systems. The evaluations cover aspects from installation to tool performance, robustness, and ease of use. The evaluation results are shared first with the developers and then posted in the ACTS Information Center. This information has provided feedback for developers so they can improve the tools if necessary. ACTS users are encouraged to submit reports and evaluations on the tools via an

electronic interface available in the ACTS Information Center. Feedback has also been gathered at workshops, conferences, and special presentations. Specific examples of these can be found in the ACTS Information Center inside the individual tool pages (<http://acts.nersc.gov/tools.html>).

- *Increased availability, stability and performance.* We have installed and maintained current versions of ACTS tools on NERSC platforms (IBM RS6000, *seaborg*; and Cray T3E, *mcurie*) and the newly installed LBNL experimental SMP Linux cluster (*alvarez*), porting the tools when necessary. In order to meet user needs, and when primary developers are not able to meet application needs, tuning or limited application-driven development has also been provided.

ACTS is currently addressing the needs of the user community through an aggressive effort to improve support for users of ACTS tools. Users are encouraged to send questions and bug reports to acts-support@nersc.gov. Moreover, the activities under ACTS have led to some direct benefits to the computational science community, for instance:

- Fewer users have unnecessarily spent resources developing codes instead of using them. ACTS tools are portable to many platforms. In most cases they have been tuned and optimized for different computing environments and operating systems. Therefore, users can expect a reduction in the time spent porting and optimizing codes. In the event of code parallelization, users can easily take advantage of the parallel data layouts, communication paradigms, and distributed data structures and mechanisms provided by the tools, hence reducing the number of problems encountered. The ACTS tool interoperability requirement allows users to utilize a wider spectrum of computational resources with minimal code changes.
- The tools have been exposed to a broad and competitive community and the resulting collected expertise from working with a variety of users has produced enough feedback for helping new users choose the ACTS tool that best meets their needs based on functionality, applicability to a specific kind of problem, and feedback from other users. In many instances, ACTS tools are used at the core of widely used scientific codes. This has motivated the continuous utilization of these tools in new developments. One of the examples of this class of codes has been the parallel implementation of the FLAPW (full-potential linearized-augmented plane-wave) method, which uses ScaLAPACK and is widely employed in material sciences high-end computer simulations [3].
- Increased use means increased feedback, ultimately resulting in higher quality tools. The user community of the ACTS Toolkit has provided the tool developers with measures of success for their tools through bug reports, features of interest and preferences. In all, this has resulted in an improvement in the quality of documentation, accumulation of domain expertise, and better parameters for tool research and development.

In the next three sections we present a list of accomplishments, lessons learned from our current activities and an overview of related activities. In Section 5, we outline an agenda for expanding the framework of the ACTS Toolkit Project. Highlights and deliverables of this agenda include a solid base for the ACTS Toolkit, peer-reviewed certification of new tools, finding tools to solve problems that are not currently addressed by the Toolkit, promoting software interoperability, working in

<i>Category</i>	<i>Tool</i>	<i>Functionalities</i>
Numerical	Aztec	Algorithms for the iterative solution of large sparse linear systems.
	Hypre	Algorithms for the iterative solution of large sparse linear systems, intuitive grid-centric interfaces, and dynamic configuration of parameters.
	PETSc	Tools for the solution of PDEs that require solving large-scale, sparse linear and nonlinear systems of equations.
	OPT++	Object-oriented nonlinear optimization package for serial architectures.
	PVODE	Solvers for the solution of systems of ordinary differential equations, nonlinear algebraic equations, and differential-algebraic equations.
	ScaLAPACK	Library of high performance dense linear algebra routines for distributed-memory message-passing.
	SuperLU	General-purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations.
Code Development	TAO	Large-scale optimization software, including nonlinear least squares, unconstrained minimization, bound constrained optimization, and general nonlinear optimization.
	Global Arrays	Library for writing parallel programs that use large arrays distributed across processing nodes and that offers a shared-memory view of distributed arrays.
	Overture	Object-Oriented tools for solving computational fluid dynamics and combustion problems in complex moving geometries.
	POET	Framework for encapsulating parallel algorithms.
Code Execution	POOMA	Collection of templated C++ classes for writing parallel PDE solvers using finite-difference and particle methods (see also Overture).
	CUMULVS	Framework that enables programmers to incorporate fault-tolerance, interactive visualization and computational steering into existing parallel programs.
	Globus	Bag of services for the creation of computational Grids and tools with which applications can be developed to access the Grid.
	PAWS	Framework for coupling parallel applications within a component-like model.
	SILOON	Tools and run-time support for building easy-to-use external interfaces to existing numerical codes.
	TAU	Set of tools for analyzing the performance of C, C++, Fortran and Java programs.
Library Development	Tulip	Implementation of the HPC++ interface specification for high performance machines.
	ATLAS and PHiPAC	Tools for the automatic generation of optimized numerical software for modern computer architectures and compilers.
	Nexus	Library that provides the multithreaded communication facilities required to implement advanced languages, libraries, and applications in heterogeneous parallel and distributed computing environments.
	PADRE	C++ layer for interfacing with libraries that distribute data on parallel computers.
	PETE	Extensible implementation of the expression template technique (C++ technique for passing expressions as function arguments).

Table 1. ACTS Tools and their functionalities (for more details see Appendix A).

collaboration with DOE computer facilities, working in collaboration with other software initiatives, encouraging feedback from users, further development of the tools, expertise tracking, improving the functionality of the ACTS Information Center, and expanding outreach efforts. The synergism of the proposed work, benefits and merits, and deliverables, are addressed in Sections 6, 7 and 8, respectively. Our ultimate goal is the creation of effective dissemination mechanisms, and a reliable software infrastructure for the long-term support and development of high quality software tools. In turn, this proposed project and related activities will directly benefit today's and future's research endeavors in computational sciences.

2. Accomplishments

In this section we summarize our accomplishments in the past two calendar years (2000 and 2001). A comprehensive description of our activities is given in Appendix B.

We have installed and updated versions of ACTS tools on NERSC platforms, IBM SP and Cray T3E. The subset of tools comprised Aztec, CUMULVS, Hypre, PETSc, PVODE, ScaLAPACK, SuperLU, TAO and TAU. The functionalities offered by these tools are described in Appendix A. Tool installations were based primarily on requests from users and also prompt availability of versions for the target machines. The ACTS Information Center contains updated information on the versions of tools installed, tool reviews and online tutorials for some tools.

We have participated and given presentations in various venues addressing computational sciences issues. The major events were the Supercomputing conferences organized by the ACM/IEEE, NERSC users group meetings, and the National Partnership for Advanced Computational Infrastructure (NPACI [64]) hands-on meetings. By participating in these and other events we had the opportunity to bring the tools outside the DOE community, discuss issues pertaining to the use of the tools and the ACTS Project. We also received feedback from users and became familiar with applications that could benefit from the use of ACTS tools. One of the results of these interactions has been the gathering of feedback and relevant data for the future agenda of the ACTS Project. This information has been shared with members of the computational science community. For instance, we encapsulated this information in a subsection of a document prepared by the NERSC User Group Executive Board (NUGEX [75]). The resulting document will be used for planning purposes by DOE's Office of Advanced Scientific Computing Research (OASCR).

For outreach and education, we have organized two workshops on the ACTS tools at LBNL. The workshops targeted academia and industry through a dissemination and instruction on the state-of-the-art tools for high performance computing environments, and simultaneously providing an umbrella for ACTS tool developers to receive the feedback from these communities. These workshops offered an introduction to the ACTS Toolkit for application scientists whose research demand includes either large amounts of computation, a large volume of data manipulation, the use of robust numerical algorithms, or combinations of these. The workshop programs included a range of tutorials on the tools, discussion sessions aimed to solve specific computational needs by the participants, and hands-on practices using NERSC's state-of-the-art computers.

A third workshop on a subset of the ACTS tools was held at the Second Los Alamos Computer Science Institute (LACSI [16]). This workshop consisted of talks on the development and application of state-of-the-art numerical tools that are currently used worldwide in science and engineering. Such mathematical tools for addressing more complex physical and societal phenomena, along with the growth of computing resources, is driving the continuous growth of the computational sciences community. This community includes scientists, designers and developers of high-end technology who require computerized modeling solutions.

Finally, we have also provided support to researchers that needed help prototyping their codes using ACTS tools or had computational problems that could be potentially addressed by these tools.

3. Lessons Learned

As a result of the activities summarized in the previous section, and also by interacting with users and tool developers, we have learned various lessons that must be taken into consideration for a successful continuation of the ACTS Project. Here we summarize the most important issues:

- *There is still a gap between tool developers and application developers which leads to duplication of efforts.* Without projects like ACTS, application developers will continue to design and implement codes using techniques that are already available from other sources. Quite often these new developments are far from optimal because of the application developers' inexperience with all the different issues that lead to optimal performance. In many cases, application developers only consult sources like Numerical Recipes [6] that do not address platform optimization and parallelism, algorithm robustness, and language specific optimization issues. Most of the items in this proposed roadmap (see Section 5) are designed to vitalize the current efforts under ACTS to promote reusability of quality and reliable software.
- *The tools currently included in the ACTS Toolkit should be seen as dynamically configurable toolkits and should be grouped into toolkits upon user/application demand.* Based on the particular needs of an application, users generally benefit from only a subset of the functionality available in ACTS; therefore, they may only need to install a subset of the ACTS tools in their computing environments. A future agenda for the ACTS Toolkit must include an automatic mechanism that will allow a friendlier and dynamic installation of the tools based on particular demands (See Subsection 5.4).
- *Users demand long-term support of the tools.* One of the main concerns that users have expressed to us is the longevity of support from tool developers and required evolution of the software as the hardware technology continues to evolve and the complexity of the scientific application continues to grow. Inasmuch as a solid base of reusable tools is utilized inside newer tool developments, users are guaranteed the evolution of the tools; thus long-term support and functionality are also guaranteed (see Subsection 5.1).
- *Applications and users play an important role in hardening tools.* The main parameters for maturity are portability, robustness, acceptance, and long-term support. It is particularly the interactions with real users and real applications that have made the software mature, portable, robust and better documented. In turn, mature software will be widely accepted inside a given scientific community. The current framework of the ACTS Toolkit has promoted the tools to a wider national and international audience, thus increasing not only the visibility of the tools worldwide but also the range of users and applications.
- *Tools evolve or are superseded by other tools.* As technology continues to advance, there are some tool functionalities that are either no longer needed or are improved as direct consequences of the user demands. An example of these changes in the ACTS Toolkit is the Aztec library being superseded by AztecOO, which is one of the components of the Trilinos solver framework [36]. An umbrella project like ACTS provides mechanisms to help users with the transition and adoption of the new tools.

- *There is a demand for tool interoperability and more uniformity in the documentation and user interfaces.* Users want to experiment with functionalities available in a subset of the ACTS tools, and finding similar user interfaces and comparable levels of support and documentation makes this task even simpler and risk free. Furthermore, the computational challenges at hand demand new software developments that interact with legacy code practices, data and computer languages. ACTS provides a natural infrastructure to put in practice all the code, data and language interoperability required by these new challenges (see Subsection 5.4).
- *There is a need for an intelligent and dynamic catalog/repository of high performance tools.* The need for a centralized software and reference repository is vital for preventing the duplication of efforts. Currently, the ACTS Information Center provides pointers to tools currently funded under ACTS with the accumulated expertise from tool users and scientific domains that are served by the tools. Additionally, we envision the inclusion of pointers to tools offering functionality not currently available in the ACTS Toolkit as well as fair comparisons with other tools that offer functionality that overlaps with the ACTS tools.

4. Related Activities

Noticeably, various initiatives have addressed the daunting task of providing solutions for and facilitating the development of high-end scientific computing codes, in particular on parallel computers. In the following we briefly list some of these initiatives.

Netlib [22] pioneered the task of collecting and distributing mathematical software, papers, and databases over the Internet. Although Netlib has provided useful services and been accessed over 150 million times, it does not answer all the needs of today's high-end computing simulations. One of the reasons is that Netlib serves mostly as a software repository and provides a rather rudimentary keyword-based search mechanism for the libraries available. Also, it often assumes that the interested user has some reasonable knowledge of linear algebra and in various cases the software documentation it provides is not consistent. The high performance branch of Netlib is HPC-Netlib [23], which provides some information about high performance mathematical software, both research and commercial, as well as some guidelines to software selection and performance issues. The National HPC Software Exchange (NHSE [24]) extended these services, by promoting software sharing and reusability within and across the HPC agency programs on a sustainable basis. NHSE facilitates the development of discipline-oriented software and document repositories and promotes contributions and use via the World Wide Web. Existing NHSE repositories include HPC-Netlib, the Parallel Tools Library (PTLib [30]), and the Chemistry Software and Information Resources (CSIR [25]). More recently, a forum for discussing scientific computing in object-oriented environments and a collection of object-oriented numerical libraries have been made available by OO Numerics [79].

Another online resource is the Guide to Available Mathematical Software (GAMS [26]), which provides a problem classification system and a cross-index to and virtual repository of mathematical and statistical software components that can be used in computational science and engineering applications. To some extent GAMS resources are, however, more limited than the ones offered by Netlib. A more specialized service is provided by MGNet [27], which is a repository for

information related to multigrid, multilevel, multiscale, aggregation, defect correction, and domain decomposition methods that are used primarily by scientists and engineers to solve partial differential equations on serial or parallel computers. Specialized services are also provided by NEOS [28], for the solution of optimization problems. NEOS also enables the solution of optimization problems remotely over the Internet.

Other initiatives address issues related mainly to portable timing routines and tools for debugging high performance applications, communication libraries, compiler technologies, parallel I/O, or distributed processing tools. Dealing with these issues are, for example, the research projects promoted by the Parallel Tools Consortium (Ptools [29]), PTLib [30], and the Center for Programming Models for Scalable Parallel Computing [31].

Educational opportunities are explored by the Computational Science Educational Project (CSEP [32]), an initiative sponsored by DOE. CSEP has filled important gaps by offering an introduction to high performance computing for research and studies in computational science and computational engineering. It targets students in science and engineering at the advanced undergraduate level and higher, and provides tutorials on public domain software that are of general interest. Therefore, it has attracted many users, ranging from high-school students to people from the commercial sector. Furthermore, CSEP offers electronic lectures and an electronic book covering various aspects of computational sciences that arise from the collaborative work by teams of mathematicians, computer scientists, and scientists, covering a wide range of disciplines. A wealth of information is also available in the form of on-line lectures: UC Berkeley's Applications of Parallel Computers [33], MIT's Applied Parallel Computing [34], University of Illinois at Urbana-Champaign's Parallel Numerical Algorithms [35], as well as the textbook Scientific Computing: An Introductory Survey [4]. Contributing to these education efforts, a dictionary of algorithms, algorithmic techniques, data structures, archetypical problems, and related definitions, is provided by [81].

The ACTS Toolkit Project is unique among these initiatives because it has provided not only pointers and documentation about software tools that address a wide number of problems in computational sciences, but it also accumulates the expertise and user feedback on the use of the software tools and scientific applications that used them. Under ACTS this expertise is captured through independent software evaluations, participation in the developer user groups e-mail list, presentation of a gallery of applications that use a particular tool, leverage between tool developers and tool users, workshops and tutorials, tool classification and general ACTS support (*acts-support@nersc.gov*).

Experience shows that in many cases software developed by scientific groups may not be suitable for distribution and use by peer scientists. Among the causes, we find that these software developments often neglect robustness, reliability and interoperability. This comes as a result of priorities that typically lie on the normal professional incentives of peer-reviewed publications and departmental appointments, rather than on the quality of the software used to produce the scientific advancements. A solution to this problem is to use readily available and reliable software technology in the development of scientific and also engineering codes. It turns out that ACTS is a vehicle that makes these technologies suitable for distribution and use among a much wider and less expert user population.

Therefore, the ACTS Information Center and its inherited accumulated expertise are allowing developers of algorithms, tools, languages, and standards to reach out other scientific communities with more robust and state-of-the-art tools. Very important to this accumulation of knowledge and expertise is the transfer of software technology that not only helps users to understand good programming practices and instances of utilization of the tools but also demonstrates why other less engineered practices would not be suitable for code optimization, portability, transfer, and extension to larger software frameworks or coupled complex scientific codes. These latter tasks are now greatly in demand by the computational science community (see [5], for example).

In contrast to all the aforementioned on-line software initiatives, ACTS has targeted the computational science community at large, where a significant percentage of the efforts focus on the development of complex parallel codes and their optimization. This expensive process often requires specialized support and information about the availability of the tools. For this reason, two workshops on the ACTS Toolkit have helped to match the tools with scientific applications. The tools and potential users that participated in these workshops are listed in Appendix C.

5. Agenda

The ACTS Toolkit Project has begun to serve as a mechanism for the development, support, and promotion of quality high performance software tools. The goals and first successes of this project have been enhanced by the increasing demand for complex high performing computer simulations, closer interactions between computer scientists and other domain scientists, less duplication of efforts, and interoperability to pick-and-play with kernels from a variety of computational services.

First of all, we propose to expand the scope of the ACTS Toolkit Project based on the lessons learned and the activities being now carried out by us under *acts-support*. The continuation and enhancement of the ACTS Project towards the creation of a reliable software infrastructure for scientific computing should also take into consideration the 11 items outlined in the following subsections.

5.1. A Solid Base for the ACTS Toolkit

Currently, the ACTS Toolkit comprises 22 tools, of which 3 have been phased out (see Appendix A), that provide implementations of numerical algorithms, support for code development and run time environments, and library development. While some of the tools have matured and tool developers have acquired more expertise by interacting with more users, there are some tools for which development and support have been discontinued, and better tools have superseded others. The matured tools now form a solid base that has contributed to the acceptance of the ACTS tools in computation sciences.

We propose a mechanism for the inclusion of new tools into the toolkit with a yearly peer-reviewed process that certifies the tools as part of the ACTS Toolkit. All tool development projects will be encouraged and technically supported by *acts-support* to reach a fully ready and matured state. At

the end of a fiscal year, tools should present progress reports, and a certification panel would judge which tools to include and which need further development. We will work with tool developers and leading computational scientists to define the criteria to be used in the certification process.

The matured and certified tools in the ACTS Toolkit will continue to set higher standards for high performance software development and continue to build the scientific community's confidence in the tools. The accumulated expertise will also provide assistance to other tool developers to cope with these advancements (see Subsection 5.4). This centralized approach also encourages a natural evolution for the basic set of tools to adapt to arising scientific challenges and new technology, thus promoting long-term and continuous support to users in the computational sciences community.

Currently, we have provided three levels of assistance, depending on the tool developers and tool readiness. At the *basic level*, *acts-support* has only basic knowledge of the tools and familiarity with their functionality, and provides help with tool installation. The *intermediate level* comprises the basic level along with staff trained to provide a higher level of support to users of a particular tool. Lastly, the *high level* support comprises the intermediate level plus staff acquiring higher levels of familiarity with the tools (for instance through the prototyping of codes using a particular tool) and conducting tutorial on the tools. In this proposed roadmap all three levels will continue to be supported and in addition we will work closely with staff at the developers' sites to increase their participation and contributions to the ACTS Information Center.

5.2. ACTS High Quality Software Certification

In 1999, the PITAC Report [74] recommended the creation of a national library of certified domain-specific software in order to reduce the labor required for software development, testing and evolution (some of those recommendations were later implemented by the Networking and Information Technology Research and Development Program [80]). ACTS has the potential of becoming a major player in such an activity. Computer vendors and a great part of the scientific computing community are already familiar with at least a couple of the ACTS tools. The implementation of the peer-reviewed certification (see Subsection 5.1) will push the frontiers of the tools forward, and the tools or subsets of the tools will be eventually integrated in vendor-supported software libraries (for example, ScaLAPACK, one of the numerical tools in the ACTS Toolkit, has already been incorporated in IBM's and Cray's scientific libraries).

We propose the creation of software certification for tools in the ACTS Toolkit that should be defined in terms of software robustness, functionality, portability, documentation, availability, and interoperability. The goal is not to preclude new software paradigms but to promote existing and new paradigms based on what works best for the computational science community. Further, this certification process defines only a minimal set of requirements that a tool must satisfy in order to be classified as matured within the Toolkit. We will provide support to tool developers via the ACTS Information Center and *acts-support* to meet these requirements and shared information from the panel reviews. In addition, every time that a new version of a tool is made available, we will perform an evaluation on the applicable platforms and write the corresponding report that will be made available via the ACTS Information Center (see Subsection 5.10).

5.3. Inclusion of New Solutions to Computational Problems

As mentioned before, the development of good quality complex parallel codes is usually very expensive. Most scientists would prefer spending their time using codes and producing scientific breakthroughs rather than programming. Therefore, the ACTS Toolkit Project should continue promoting code reusability for the solution of common and important computational problems as a means of accelerating scientific discoveries.

While the ACTS Toolkit has already provided a set of tools for the solution of a variety of problems, it is still incomplete. For example, it lacks algorithms for the solution of sparse eigenvalue problems, direct methods for the solution of indefinite symmetric linear systems, and software for visualization and data mining. However, researchers outside the DOE labs have already developed some implementations of these techniques. The ACTS Toolkit could therefore fully attend the computational needs of the DOE's Office of Science (SC) community by also providing a catalog of high performance tools for the solution of important and common problems, fostering the development of tools that are not available and recommending good quality tools developed by parties not necessarily funded by DOE.

ACTS will continue to be in a leading role in promoting high quality performance computing tools and working closely with other communities to enhance the current content of the toolkit. The metric for inclusion will be tool maturity, applicability, readiness and interoperability. The selection and certification of tools will be administered by the same peer-reviewed mechanism introduced in Subsection 5.1.

5.4. Interoperability and Software Distribution

The current ACTS funding has encouraged software interoperability and has spawned a few collaborations between projects under ACTS. Interoperability may affect performance in some cases but it potentially reduces time to solution and, most important, it assures longevity of the software. We have been working on a more general solution that can be addressed by the Common Component Architecture (CCA [41]) specifications funded by the DOE Scientific Discovery through Advanced Computing (SciDAC) Program [15]. We foresee the use of these specifications as a requirement for tools to be certified by the ACTS Toolkit Project and this will facilitate the interaction between the tools and migration from one tool to the other when necessary.

Moreover, language choices made by the tool developers must not dictate the choice of language used by the application developers. Thus, we propose to deal with the language interoperability issue by exploring and exploiting interface language techniques like the ones used by the Babel project [44].

The ultimate goal of the software certification (Subsection 5.2) and interoperability will be the distribution of the ACTS tools on user demand. For now, this has remained a difficult task given the different compiler requirements, makefiles, distributions, and installations posed by the different tool developers. By composing all the requirements and specifications inside centralized mechanisms for software distribution inside the ACTS Information Center, we will be able to

provide a better service to the scientific community and alleviate early difficulties related to tool installation and testing. To realize these tasks, we propose the creation of dynamically configurable scripts that will allow users to install subsets of the interoperable tools under different computing platforms using scripting languages.

The ACTS Software repository will be a collection of actual tar (tape archive) and compressed files that will contain the distribution files of supported versions of the tools and corresponding licensing agreements. For tools that require special licensing agreements, as it is currently the case for some tools in ACTS, the repository will contain actual links to the software download sites where potential users can sign and agree to special licensing requirements. The ACTS repository will provide the aforementioned automatic installation scripts for all the tools, whether the distribution files are in the local ACTS repository or in a remote location. For the latter, users will be instructed to complete the required downloads before running the installation scripts.

5.5. Active Collaboration with DOE Computer Facilities

The spectrum of users in the DOE computer facilities consists of scientific challenge teams (multidisciplinary and multi-institutional teams engaged in research, development, and deployment of scientific codes, mathematical models, and computational methods to exploit the capabilities of terascale computers), high-end capability users (single PI teams and their groups of collaborators or students), and new users transitioning from midrange computers. ACTS can be instrumental for users in the transition from midrange computing applications to high-end ones.

Our participation in the ACTS Toolkit Project has built the necessary expertise to support an infrastructure for dissemination of software tools, helping users prototype the codes, and leverage between developers and the users of a tool. As a matter of fact, the relevance of the ACTS Project for NERSC users has been acknowledged in the “Green Book” prepared by the NUGEX [75]. The Green Book is a document written at regular intervals by a NUGEX task force detailing the hardware, software, infrastructure, and training needs of the NERSC user community. It also helps NERSC and the DOE Office of Science in the planning process of the NERSC facility.

Although the ACTS Toolkit Project has close relations with NERSC, we would like to stress that ACTS is not exclusive to NERSC. Therefore, we also propose to carry out the ACTS activities in close collaboration with other DOE computer facilities. In particular, we will coordinate these activities with the High Performance Computer Center of the Center for Computational Science [37] at the Oak Ridge National Laboratory (ORNL), the Pacific Northwest National Laboratory’s (PNNL) Environmental Molecular Sciences Laboratory (EMSL [38]), and Argonne National Laboratory’s (ANL) Mathematics and Computer Science Division [39]. These collaborations will enrich the ACTS repertoire of scientific applications and will provide for a better assessment on the performance of the tools on different computer platforms and configurations.

For more information about collaborations with DOE and non-DOE computer facilities see Subsections D.1 and D.4 in Appendix D.

5.6. Working in Collaboration with Other Software Initiatives

Currently, there are a number of high-end computing programs that could interact with the ACTS Toolkit Project, such as the SciDAC Program [15], NASA's High Performance Computing Cooperative Agreement Notice (CAN [19]) round III, NSF's Information Technology Research (ITR [20]) and NPACI [64]. All of these programs will develop software technology to address complex scientific problems in today's and the future's high-end hardware. In turn, there are several tools under the current ACTS umbrella that can potentially benefit projects funded by the aforementioned programs.

In particular, ACTS tools have a great potential for providing a jump start for the SciDAC science projects; some of the ACTS tools (such as Hypre, PETSc, ScaLAPACK and SuperLU) will also be used as base software for some SciDAC's Integrated Software Infrastructure Centers (ISICs). We have been interacting with various PIs of the SciDAC-funded ISICs, who have realized the need for a centralized tool deployment infrastructure that will facilitate their inter-ISIC interactions, collaborations with other DOE funded projects, and promotion of the resulting tools outside the SciDAC research teams. The roadmap proposed here will also take into consideration future needs that will arise in later stages of the ISICs. Currently, we have contacted several of the PIs of the SciDAC Projects. Although some of these projects are still at an early stage of development, they have all identified a need for deployment of their resulting software and have seen the benefits of interacting with the ACTS Project to profit from the existing ACTS infrastructure (more details can be found in Appendix A).

5.7. Encourage Feedback from Users

The ultimate measure of success for the high performance computing tools provided by ACTS is their role in the production of high quality scientific results. This can be assessed, for instance, by encouraging feedback from current and potential users as well from participants in the workshops and other activities related to ACTS, and by giving presentations and fomenting discussions in major events covering computational sciences topics.

We propose the continuation of independent software reviews to be made available in the ACTS Information Center. The frequency and updates of these reports should be increased as new versions of the software are released. Tool developers will continue to have an opportunity to examine the reviews and report any inaccuracy or suggest any changes when appropriate. Besides, we will collect feedback from current users of ACTS tools and use that information to motivate other potential users. In this respect, we propose the creation and investigation of other useful parameters to support the usage of the tools. This includes, but is not limited to, the time that an application team spent learning about the use of a tool, the time it took a development team to prototype an application with a particular tool, what makes a tool more successful or popular than others, and a set of metrics of the success of the ACTS Toolkit Project. Selected cases could be published in conference proceedings and eventually in a journal on computational sciences.

5.8. *Further Development of the Tools*

When appropriate and based on user specific needs, we propose to employ our expertise to improve the features provided by a particular tool. We foresee this kind of activity happening in collaboration with the tool developers or independently, when funding for the development of a certain tool has ceased. As an example, some users of ScaLAPACK have already expressed an interest in features that are not currently available in that library, such as tools for facilitating the block cyclic data distribution required by ScaLAPACK. Others have expressed interest in a generalized library for data representation and translation between the numerical software tools. A last example is a distributed and general-purpose coupling interface that preempts the communication bottlenecks caused by the current centralized practices.

We receive feedback from user needs and community requirements that have been passed to tool developers. Our interactions with other scientific groups at the national and international levels have also provided new directions for development (e.g., PyACTS [18]) and collaborations. The ACTS development venue will only create a more robust infrastructure and will provide guidance for future research endeavors and more parameters for next rounds of funding by the different agencies.

We propose to actively support tool developers in the adoption of interoperable software solutions like the CCA. We have already been participating in the CCA workshops, and this technology has already been introduced to participants of the ACTS Toolkit Workshop in October 2001. We will work with the researchers involved in the CCA activities in the preparation of documentation and future tutorials, as well as work with the CCA related projects and activities in the deployment of the CCA technology (See Appendix D).

5.9. *Expertise Tracking*

The ACTS Information Center attempts to collect and make available all information from tool developers, *acts-support* and users' experiences with the tools. In some cases, tool developers keep their own e-mail lists of users, by means of which questions and problems are posed and later either the developers or other experienced users propose answers and solutions. We propose to expand this capability in the ACTS Information Center by continuing to participate in these e-mail lists and using more specialized search tools to retrieve this information from feedback, evaluations, and reports collected on a particular tool. This capability will greatly enhance the functionality of the software catalog by providing self-guided information through the tool catalog, its documentation and reports. More details are given in Subsection 5.10.

5.10. *Increasing the Visibility of the ACTS Information Center*

Good quality software is usually labor-intensive to develop, test, maintain, and evolve. We envision ACTS as a mechanism that helps improve the life cycle of scientific computing software and also a delivery vehicle for software developed with DOE's Mathematical, Information, and Computational Sciences (MICS) Program funding. For this, we propose to continue improving the quality of the information that is now available in the ACTS Information Center and include new features like a

lexical search or knowledge-based engine [7][10] that can drive the users not only through the information center but also to the associated developers' information sites.

We propose the creation of a domain-expertise cross-indexing mechanism to facilitate the access to the ACTS Information Center. Users will find pointers to information categorized by scientific domain and tool utilization. Furthermore, a new semantic search engine that employs a concept-based information retrieval model will power this information center.

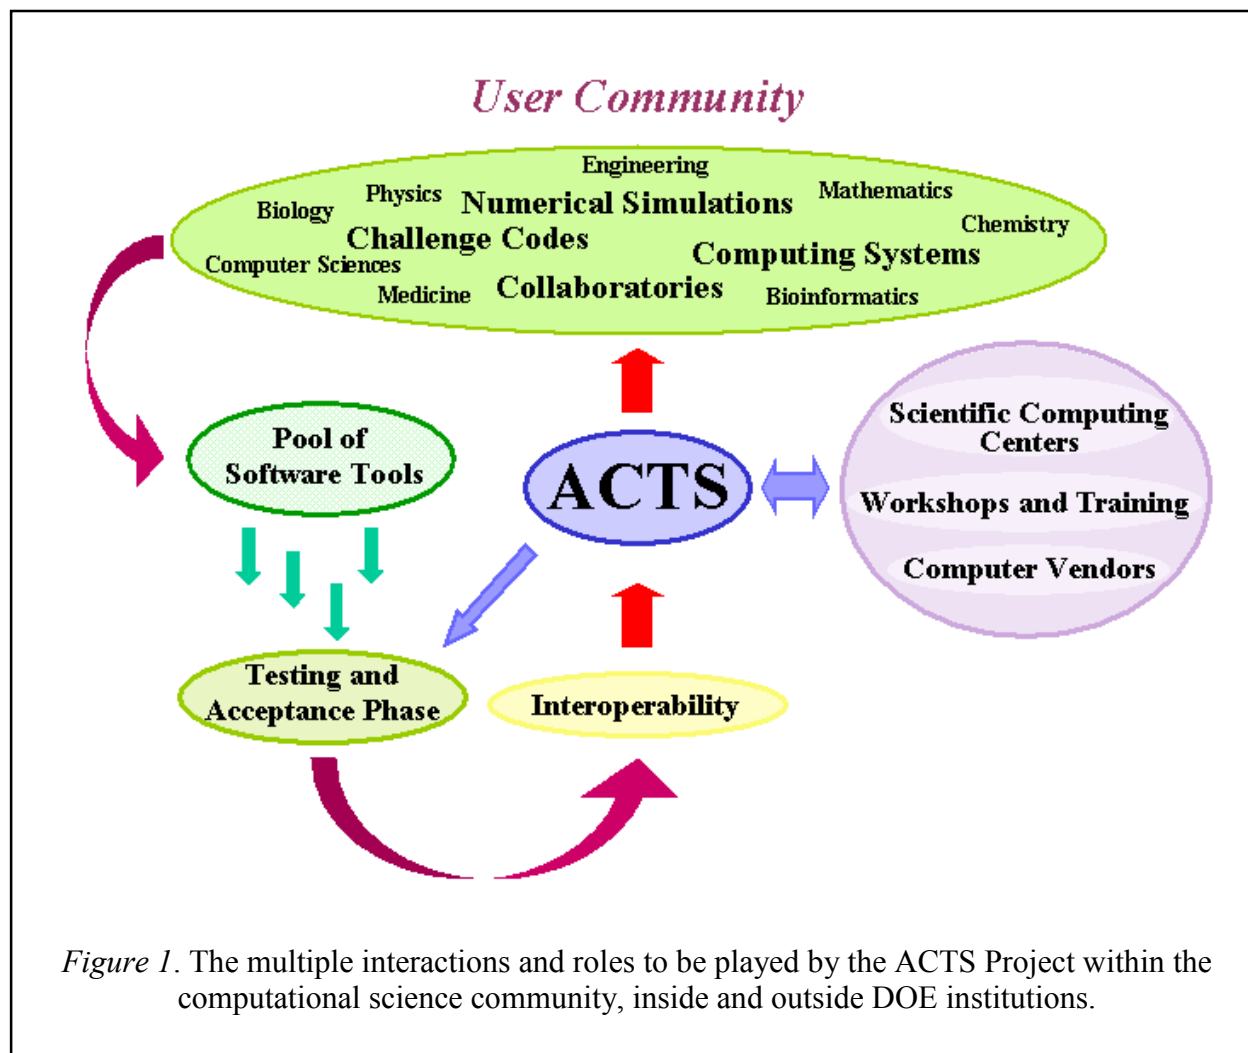
The ACTS Information Center should also provide different levels of interaction with the users depending on their level of familiarity with the tools. A beginner level may require more directions on tool selection and information on their functionalities as well as installation. Intermediate users may require more pointers to solutions to specific problems and questions about using the tools. Advanced users may need guidance in combining the functionality available in the tools as well as interaction with ACTS support and tool developers. All these levels will be included in the same ACTS website.

There are a few on-line tools in the DOE community that dexterously support dynamic web pages, distributed databases, search engines, and software repositories. An example of these tools is Alexandria [45], which is a research effort on the development, cataloging, and distribution of component-based software. We plan to interact with the developers of Alexandria and other similar tools to incorporate some of their functionality in the ACTS Information Center and accumulate their current expertise with the deployment of component software technology and large heterogeneous information. Clearly, the integration of these technologies into the ACTS Information Center will only broaden and improve our mechanisms to reach out to the computational science community.

5.11. Dissemination Plans

From our experience with the ACTS Toolkit Project, we have learned that education of graduate students and postdoctoral fellows provides a viable resource to build a bridge between computer scientists and domain scientists (see Appendix B for list of workshops and tutorials organized by ACTS). This need has also been identified in the review of the ACTS Project in January 2000 [17]. By educating other scientists on the use of a set of tools, they become familiar with the technology, accept this technology faster, are able to develop codes using state-of-the-art tools, and minimize the tool selection process. This outreach mechanism also provides a valuable technique for capturing expertise from the tool developers, users, and instructors. As a result, a database with expert information is now being built and made available to the computational sciences community via the ACTS Information Center.

As part of the dissemination plans, the ACTS Project will host one workshop per year and continue to reach out to other communities, such as other scientific computing centers, by participating in their various venues to reach their users (e.g., the NSF summer parallel institutes and scientific computing courses). The current infrastructure of the ACTS Project has also provided means for bringing the technology to institutions and researchers that are not yet playing leading roles in the scientific computing communities or do not have access to high-end computing. For instance, we



have participated in workshops organized by the DOE's Experimental Program to Stimulate Competitive Research (EPSCoR [40]) and have acted as point of contact for researchers that submitted proposals to this program (see Appendix B). ACTS will implement on-line tutorials and software demonstrations that will be accessible from the ACTS Information Center. As the World Wide Web and Grid technologies [12] have continued to provide improved mechanisms for distributed learning, we will exploit these mechanisms to improve the on-line tutorials and software demonstrations.

6. Synergism of the Proposed Work

Meeting the many needs of computational tool developers and users requires a virtual forum in which these communities can share expertise, questions, and concerns. The expanded ACTS project will fulfill this over-arching need by becoming a collaboration framework for on-going tool development and deployment. In Figure 1, we show how we envision the multiple interactions and roles to be played by the new phase of the ACTS Project within the computational science

community. ACTS will continue to work with a basic set of reliable tools and will pay close attention to other tools being developed by other initiatives in order to incorporate them into the Toolkit (see Subsections 5.1-5.4, 5.6 and 5.8). In the figure, the producers of tools to be tested and accepted into ACTS are tool developers working in the *User Community* and the *ACTS Project* (See Subsections 5.2 and 5.6).

We plan to continue exploiting the current ACTS infrastructure and to bring new software development under the umbrella. As presented in Subsections 5.5-5.11, the expanded infrastructure for the ACTS Project will be able to accumulate expertise from tool developers, users and application scientists and produce relevant feedback and knowledge to be disseminated inside the computational sciences community. Strong components of this infrastructure will be the coordination of support to users, organization of *Workshops* and other educational tools (such as on-line material) for education and *Training* purposes, and interaction with *Scientific Computer Centers and Computer Vendors* to ensure the visibility of the tools, to continue to promote software reusability, to set high standards for high performance computing software, and to expose the tools to a wide variety of applications and contexts. As depicted in Figure 1 and discussed in Subsections 5.1-5.2, these new software developments will be *Tested and Accepted* after meeting certain software criteria, in which *Interoperability* will play a major role to guarantee software reusability, incremental development, and a ready-to-use wide variety of services to the end users.

The hard problems at hand for the ACTS Toolkit, and high performance software in general, continue to be language and software interoperability, uniform software distribution and licensing, friendly tool interfaces and tool installation procedures, performance, and tool acceptance. As described in Sections 1, 3 and 5, we plan to implement the expansion of the ACTS Toolkit based on the valuable lessons learned from the outcomes of the current scope of the project, and work with the aforementioned institutions, projects, and DOE initiatives to formulate solutions to these hard problems.

7. Benefits and Merits of the Proposed Work

In the proposed work for an expanded ACTS framework, current and new users of the ACTS Toolkit will be able to take full advantage of a centralized software and reference repository. Additionally, users will have access to the accumulated expertise and knowledge on the different areas of computational sciences. This expansion has been described in Section 5 and the required levels of funding are justified in Section 9. Notice that all the details of the collaborations and expanded activities will be carried out in harmony with existing research institutions and already funded projects (like SciDAC), as it will be virtually impossible to deliver similar resources and services without their identified support and interest in this proposed work. Furthermore, most of the interactions and collaborations proposed in Section 5, and detailed with a few examples in Appendix D, are based in the principle of building an infrastructure that serves as an active instrument for research collaborations, as well as promotes long-term support and utilization of high quality interoperable software.

We will deliver an expanded ACTS infrastructure to host a solid base of computational tools with the appropriate levels of support and expertise, and to serve as a buffer between tool and application

development. We will produce reports, evaluations and coordinate peer reviews that will certify tools in the ACTS Toolkit, thus encouraging other tool development efforts to meet comparable or higher touchstones. In addition, we will provide a set of mechanisms for training and educating existing and potential users of the tools. We will continue current and establish new collaborations with high performance computing centers and computer vendors to guarantee the long-term support of the tools and their acceptance. We will reinforce the current ACTS Information Center with better didactical tools, lexical search engine, knowledge-based engine, web-based interfaces and centralized software and reference repository. The ACTS software repository will facilitate tool downloads and corresponding auto-build scripts that will guarantee a uniform and user-friendly installation procedures. This collection of deliverables will result in an aggressive program to continue supporting and reaching out the ever-expanding computational science communities and related fields.

The merit and novelty of this proposed expansion to the ACTS Project is that it addresses the needs of the computational science community at large, where significant efforts are focused on the development of complex parallel codes and their optimization. This expensive process often requires specialized support and information about the tools. The expanded ACTS Project will not only host a state-of-the-art software repository where tools and their ongoing developments are available and documented, but also a collaboration framework in which knowledge and expertise will be captured and shared for the benefit of the computational science community at large.

The benefits of our proposed work can be measured in many ways. A wide range of scientific code developers and users will benefit from information and education about state-of-the-art, high performance computational tools. They will benefit from the development and promotion of robust, effective, portable, usable, and durable software. They will benefit from the increased interoperability of tools, which promotes the evolution and adoption of current software development projects into future software technologies. They will benefit from multidisciplinary collaborations and the consequent accumulation of expertise. Perhaps they will benefit most from spending less time on code development and having more time to devote directly to scientific discovery. The resulting acceleration of discovery will benefit DOE's science mission, and our society, in such critical areas as fusion energy, accelerator design, combustion, and global climate change.

8. Deliverables

In this section, we present a list of core milestones, deliverables and related activities. Figure 2 shows a time chart with these milestones and deliverables for the current fiscal year and following three years (FY03-FY05). We have strategically categorized these milestones and deliverables into four groups. The first group, *Solid Base of Tools*, contains the activities necessary to maintain and deliver a robust set of quality tools. *Information Center* refers to the activities involved in the development, management and enhancement of the ACTS Information Center. *Outreach and Training* contains a minimum set of planned activities for reaching out users of the tools. Lastly, *Collaboration with Others* refers to the interactions with other projects and institutions for the successful delivery, port and acceptance of DOE's high performance software technology by other communities.

Inside square brackets we list the personnel to be involved in each milestone and related activities. An estimated length of time per activity is also included. This length of time corresponds to the width of the rectangles in Figure 2. The complete list of the personnel and their general roles in the project are described in Section 9. We plan to hire a Tool Specialist (TS) and an Information Center and Tool Repository Maintainer (ICTRM) upon full funding of this proposal. Also, some of the activities will involve collaborations with CCS/ORNL personnel (ORNL-POC)

For FY03 through FY05, we have only listed the minimum set of activities to be realized under *Outreach and Education* and *Collaborations With Others*. We expect that the expanded ACTS infrastructure will bring far more activities in the years FY03 to FY05 than in the first two years. Further, we plan to update this time chart as research initiatives, SciDAC deliverables, conferences, workshops, journals and other venues become available.




Figure 2. Milestones, deliverables and related activities.

FY02

Activity, Milestone or Deliverable	
Box-ID	
S1	Tool Support. Support for some of the tools currently in the ACTS Toolkit. Support levels are described in Subsection 5.1. [Marques, Drummond, TS, and ICTRM. Time: yearlong].
S2	Tool Installation. At NERSC, continue installation of tools on demand and availability on the IBM SP, CRAY T3E and Linux cluster. Also, perform software upgrades for tools already available on the NERSC systems (IBM SP and CRAY T3E), and the LBNL Linux cluster. [Marques and Drummond. Time: yearlong].
S3	Tool Evaluation. Perform tool evaluations and send feedback to tool developers. The deliverables are reports on the tools, which will be available later in the ACTS information center. [Marques and Drummond. Time: March to July].
S4	CCS/ORNL Collaboration. Setup and coordinate installation, support and training for ACTS tools at CCS/ORNL. [Marques, Drummond, and ORNL-POC. Time: February to December].
C1	ScIDAC Contacts. Contacted ScIDAC PIs. Contacted other collaborators to discuss the proposed expansion to the ACTS Project. Work with ScIDAC PIs to integrate tools from ISICs in the ACTS Toolkit. [Marques, and Drummond. Time: yearlong]
C2	CCA computer systems. Researched and worked on a solution to host the Common Component Architecture (CCA) repository, web servers and development systems at LBNL. The systems were moved to LBNL on Feb 7 and have been operational since. [Marques, Drummond, Allan and Armstrong. Time: October to mid-February].
C3	Collaboration. Prof. Jose Roman, from the Parallel Computing Group, Department of Information Systems and Computation of the Valencia University of Technology visited LBNL. Prof. Roman participated in the ACTS Toolkit Workshop in Berkeley and Santa Fe (see Appendix C). Prof. Vicente Hernandez and Prof. Vicente Vidal from the same department have been working with ACTS tools (ScalAPACK, PYODE and PETSc) in the development of their high performance applications for science and engineering. Collaborations with these groups started in June 2001 [Marques, and Drummond. Time: September to October].
C4	Education. Work with the San Diego Supercomputing Center's (SDSC) Scientific Computing Support group, the Education, Outreach, and Training Partnership for Advanced Computational Infrastructure (EOT-PACI), and the Center for Computational Science at Boston University developing an introductory course to High Performance Computing. The course will be presented in a Summer Computational Science institute that will target minorities and under represented groups in computational sciences. Co-organizers of this course are Roscoe Giles and Donald Frederick. [Drummond. Time: March-August].
C5	Collaboration. Resume collaborations with the Computer Sciences and Artificial Intelligence Department at the University of Alicante in Spain. In particular, work with two graduate students, aspiring for Doctorate degrees, who are interested in deploying tools in the ACTS Toolkit in their different scientific and engineering applications. [Drummond. Time: March to September].
C6	Collaboration. Possible collaboration with the National Center for Atmospheric Research (NCAR) Software Engineering Group in preparing the list of requirements for the Community Climate System Model (CCSM) and the NASA HPCC Earth System Modeling Framework. Interest on the ACTS tools was expressed in a meeting of the Software Engineering Working Group at NCAR, February 14-15. [Drummond. Time: next meeting scheduled in May]
C7	Visiting Scientist. Computational Science Graduate Fellow Practicum, Mary Ann Leung, will research and develop more accurate and high performing methods for material sciences using ACTS tools. [Drummond Marques and Canning. Time: June – September].

Collaboration With Others

Solid Base of Tools

FY02 (Cont.)

	Activity, Milestone or Deliverable
	Box-ID
	<p>11 ACTS On-line Information. Created the new ACTS Information Center, and updated tool information and evaluations [completed on Oct. 2001]. Continue incorporating feedback from previous ACTS Toolkit workshops, and update material from presentations and tutorials. [Marques and Drummond. Time: October to March].</p>
	<p>12 Software and Reference repository. Design software and reference repository (see Subsection 5.4). Research on available technology and licensing agreements between DOE labs and tool developers. [Marques and Drummond. Time: April to September].</p>
	<p>13 Tool Evaluations. Make tool evaluations available as they are produced and shared with tool developers (see S3). [Marques and Drummond. Time: April to September].</p>
	<p>14 Update Web Pages. Add presentations, tutorials, information and feedback from the NPACI All-hands Meeting (see O4), Copper Mountain Meeting (O5), workshop on Scaling to New Heights (see O7), and Institute for Mathematics and its Applications (see O9). Add material to be prepared for Introduction to High Performance Summer Institute (see C4 and O10). [Marques and Drummond. Time: May to September].</p>
	<p>15 New Website. In preparation for the ACTS Toolkit Workshop 2002 (see O8), design and create on-line forms, gather pointers to information related to ACTS for prospective workshop participants, and advertise the workshop. [Marques and Drummond. Time: July to September].</p>
	<p>Workshops:</p> <ul style="list-style-type: none"> • ACTS Toolkit Workshop II, <i>Solving Problems in Science and Engineering</i>, Berkeley, CA. [Marques and Drummond, October 10-13]. • Los Alamos Computer Science Institute (LACSI) Symposium, <i>High Performance Numerical Libraries for Science and Engineering</i>, Santa Fe, NM. [Marques and Drummond, October 16-18].
	<p>O2 SC2001. Denver, CO, talk in the NERSC booth, <i>What Can ACTS Do For You?</i> Information materials: demos, flyers and posters. NERSC Users Group: <i>Status of the ACTS Toolkit Project</i>. [Marques and Drummond, November 10-16]</p>
	<p>O3 CCA Meeting. Participated in the <i>CCA Meeting, Santa Fe, NM, with a presentation in the MxN discussion group. Also, we planned and discussed details of the migration of the CCA repositories and testbed systems to LBNL</i>. [Marques and Drummond, January 9-11].</p>
	<p>O4 Outreach and Education. Tutorials and presentations on numerical libraries for high performance Computing in the ACTS Toolkit. The venue will be the "All-Hands" meeting organized by NSF's National Partnership for Advanced Computing Infrastructure (NPACI), La Jolla, CA. [Marques and Drummond, March 6].</p>
	<p>O5 Outreach and Education. Seventh Copper Mountain Conference on Iterative Methods. We will present tutorials on Numerical Tools in the ACTS Toolkit in a workshop. [Marques and Drummond. Time: March 24-29].</p>
	<p>O6 Outreach. Participate in the Annual PT tools Meeting, Knoxville, TN. [Marques. Time: 2 days in May – to be determined].</p>
	<p>O7 Outreach. Workshop: Scaling to New Heights, Pittsburgh Supercomputing Center. We will present a talk on <i>Developing and Using Tools for High-Performance Computing Tools</i>. [Marques and Drummond, May 20-21].</p>
	<p>O8 Workshop. Organize and promote the ACTS Toolkit Workshop 2002. The workshop will take place either in the third or fourth week of September. [Marques and Drummond, July - September].</p>
	<p>O9 Outreach and Education. Participate in the Institute for Mathematics and its Applications' (IMA) summer graduate program, July 1-26 (http://www.ima.umn.edu/PI/2002summersgrad.html). [Marques and Drummond. Time: ACTS Tools one week].</p>
	<p>O10 Education. Participate in the High Performance Computing Summer Institute (see C4) that targets minority students in Science [Drummond, one week in August].</p>
	<p>O11 Outreach. Participate in SCICOMP6 to be held at LBNL. [Marques and Drummond. Time: August 19-23].</p>

FY03

		Activity, Milestone or Deliverable
	Box-ID	
	S5	Tool Support. Continue support as described in S1 . The <i>acts-support</i> will team up with CSS/ORNL. Provide support for tools that are in the ACTS Toolkit by the beginning of FY03. [Marques, Drummond, TS, ICTRM, and ORNL-POC. Time: yearlong].
	S6	Tool Installation. From S2 , continue tool installation on user demand, availability (i.e., new tools and new tool versions) at NERSC's and CSS/ORNL's systems. [ICTRM. Time: yearlong].
	S7	Integrate New ISIC Tools. Integrate available tool deliverables from ScIDAC ISICs to the ACTS Toolkit. The <i>acts-support</i> team becomes familiar with new tools. [Marques, Drummond, TS, ICTRM and ORNL-POC. Time: yearlong].
	S8	Tool Evaluation. Perform tool evaluations and deliver corresponding reports from evaluations. Share reports with tool developers. [Marques, Drummond, and TS. Time: January to May].
	S9	Peer-review Panel. Organize peer-review panels and distribute information regarding the tools to be reviewed by each panel. The peer-reviewed process will end with a 2-day meeting in which some tools will be certified and others retired. This process will deliver peer-review panels, reports and a solid based of tools from which some will be certified by ACTS. [Marques, Drummond, and TS. Time: May].
	S10	Tool Enhancement. For tools requiring only minor revisions, <i>acts-support</i> will work with tool developers to complete the recommendations from reviewers. The deliverables from these activities will result in better quality tools and potentially more certified tools. [Marques, Drummond, TS, and ICTRM. Time: mid-May to mid-July].
	C8	ScIDAC Collaborations. Follow up (from C1 and C2) and establish new ScIDAC and DOE collaborations. Continue working with ScIDAC ISICs defining mechanisms for long-term support and deployment of the tools. In addition, work with tool developers to meet certification requirements. [Marques, and Drummond. Time: yearlong].
	C9	Collaborations. Continue collaborations in C3 and C5 . Results from these collaborations will be published in international conferences and journals [Drummond and collaborators. Time: yearlong]. Continue Collaborations with Dr. Yves-Henri Sanejouand from the <i>Centre de Recherche Paul-Pascal</i> , Pessac. [Marques. Time: December to April]. Continue collaborations from early interactions with CITRIS (currently interested in PVODE), EOT-PACI (for outreach and education), NCAR (CCSM and ESMF, see C6), and the University of the Basque Country (PyACTS and PyClimate). [Marques and Drummond. Time: yearlong].
	C10	Tool Certification. Work with tool developers and other qualified computational scientists on the criteria for tool certification. The criteria will be delivered as a document to be shared with the computational science community. We will propose a Birds-of-a-Feather session at SC 2002 (see O12) to present and gather further feedback from the computational science community. [Marques, Drummond and Tool Developers. Time: October to mid-November].
	C11	Interoperability. Work with CCA and BABEL projects developing examples of component based software technology using ACTS tools. [Marques, Drummond, and TS. Time: February to March].
	C12	Visiting Scientist. Host another Computational Science Graduate fellow during her/his practicum. [Marques and Drummond. Time: June to September].
	I6	Distributed Information Center. Work with CCS at ORNL to manage and maintain the ACTS Information Center in a distributed manner. These process will require to setup policies and mechanisms to update and check versions of the ACTS Information Center, as well as, versions of the distributed software and reference repository. [Drummond, ICTRM, and ORNL-POC. Time: yearlong].
	I7	Search Engine. Design and develop a new semantic based search engine (See Subsection 5.10). This search engine will provide access to information available at other high performance computing and computational sciences sites as the ones described in Section 4. [Marques, Drummond, and ICTRM. Time: yearlong].
Information Center		
Collaboration With Others		

FY03 (Cont.)

Activity, Milestone or Deliverable	
Box-ID	
18	Update Web Pages. Create and update web pages for new ScidAC tools in the ACTS Information Center. [ICTRM. Time: yearlong].
19	Software and Reference Repository. Implement and setup software and reference repository (see 12). [Marques, Drummond and ICTRM. Time: October to December].
110	Outreach and Education Information. Add presentations, tutorials, information and feedback from the ACTS Toolkit Workshop 2002, and SCICOMP6 (see 08 and 011). [ICTRM. Time: October].
111	Tool Certification. Make tool certification document and process available in the ACTS Information Center (see 013 and C10). [ICTRM. Time: mid-December].
112	Tool Evaluation. Make tool evaluations available as they are produced and after shared with tool developers (see S8). [ICTRM. Time: February to mid-June].
113	Update Information Center and Repository. Update software and reference repositories with new tool information. Add presentations, tutorials and feedback from 014 , 015 and 016 . [ICTRM. Time: May to September].
114	Update Tool Certification and Retirement. Update ACTS Information Center with results from peer-reviewed panel certification (See S9). Thus, some tools will be certified and some will be retired. [Marques and ICTRM. Time: mid- June to Mid-July].
115	ACTS Toolkit Workshop 2003. In preparation for the ACTS Toolkit Workshop 2003 (see 018), design and create on-line forms, gather pointers to ACTS related information for prospective workshop participants, and heavily advertise the workshop. [Marques, Drummond, and ICTRM. Time: July to September].
012	Progress Report FY02. Prepare a FY02 progress report to DOE MICS office. [Marques and Drummond. Time: October].
013	SC2002. Baltimore, MD: we will prepare demos, presentations, and posters. We will propose a Birds-of-a-Feather on tool Certification and will present at least one talk on the status of the project. We will also be available at the ORNL and NERSC booth for consultations. [Marques, Drummond and TS. Time: November 16-22].
014	Outreach. <i>Participate (talk, paper presentation, and/or poster) in SIAM Conference on Computational Sciences and Engineering, San Diego.</i> [Marques and Drummond. Time: February 10-13].
015	Outreach and Education. Based on prior participations, we will present tutorials and presentations on numerical libraries for high performance Computing in the ACTS Toolkit. The venue will be the “All-Hands” meeting organized by NSF’s National Partnership for Advanced Computing Infrastructure (NPACI), La Jolla, CA. [Marques and Drummond. Time: one day between February and March].
016	Education. Based on prior discussions with various PI s from the ScidAC Center for Component Technology for Terascale Simulation Software (CCTSS), we will organize tutorials on interoperability and software reusability. [Drummond, Bernholdt, and McInes. Time: 2 days between March-April].
017	Reviewed Publication. Prepare and publish a paper on selected tool functionalities available in the ACTS Toolkit and their applications. [Drummond, Marques and TS. Time: Summer]
018	ACTS Toolkit Workshop 2003. Organize and promote the ACTS Toolkit Workshop 2003. The Workshop will take place either the third or fourth week of September. [Marques and Drummond, July - September].
019	Outreach. Participate and contribute to the SCICOMP 8, Minnesota. [Marques and Drummond. Time: sometime in August].
Outreach and Education	

FY04

Activity, Milestone or Deliverable	
Box-ID	
S11	Tool Support. Continue support as described in S5 and include support for tools included in FY03. We anticipate providing different levels of support for tool deliverables from SciDAC ISICs (S7). [Marques, Drummond, TS, ICTRM, and ORNL-POC. Time: yearlong].
S12	Tool Installation. From S6 , continue tool installation on user demand, availability at NERSC's and CSS/ORNL's systems. [ICTRM. Time: yearlong].
S13	Tool Evaluation. Perform tool evaluations and deliver corresponding reports from evaluations. Share reports with tool developers. [Marques, Drummond, and TS. Time: January to mid-May].
S14	Peer-review Panel. As in S9 , organize peer-reviews. [Marques, Drummond, and TS. Time: mid-May to mid-June].
S15	Tool Enhancement. For tools requiring only minor revisions, ACTS support will work with tool developers to complete the recommendations from reviewers. The deliverables from these activities will result in better quality tools and potentially more certified tools. [Marques, Drummond, TS, and ICTRM. Time: end-May to mid-July].
S16	ANL and PNNL Collaboration. Identify and work with point of contact at ANL and PNNL to extend the ACTS support to these centers. [Marques, Drummond, ORNL-POC, ANL-POC, and PNNL-POC Time: yearlong].
S17	Automatic Tool Installation. Design and implement automatic tool installation for tools that not provide similar service (<i>autoconf</i> -based, or similar). Standardized installation scripts for all of the supported tools, this process is required to support the dynamic configuration and installation of set of tools on user demand. Scripts will be supported on a range of computer platforms. [Marques, Drummond, TS, and ICTRM. Time: mid-January to September].
C13	SciDAC and Tool Developers Collaborations. Follow up collaborations from C8 . Continue work with SciDAC ISICs defining mechanisms for long-term support and deployment of the tools. Expand the current ACTS bug reporting mechanisms to include means for interacting with bug reporting mechanisms available at the tool developer sites. [Marques, Drummond, and TS. Time: yearlong].
C14	Collaborations. Continue collaborations in C9 . Results from these collaborations will be presented in conferences, and submitted for publication to journals. Continue collaborations from early interactions with CITRIS, EOT-PACI and possibly NCAR, Research and propose software enhancements to ScalAPACK (see Appendix D, and Subsection 5.8). [Marques and Drummond. Time: yearlong].
C15	Tool Certification. Work with tool developers and other qualified computational scientists on reviewing and updating the tool certification criteria. The work will deliver a document to be shared with the computational science community. We will propose a Birds-of-a-Feather session at SC 2003 (see O20) to present and gather further feedback from the computational science community. [Marques, Drummond and Tool Developers. Time: October to mid-November].
C16	DOE Science Grid. According to the DOE Science Grid development and implementation plan, this will be a suitable point in time for ACTS to start collaborations with the DOE Science Grid. This project will enable a further expansion of the ACTS infrastructure to cover other DOE High Performance centers (like PNNL and ANL, see S16). [Marques, Drummond, and TS. Time: March to September].
C17	Computer Vendor Forum. In collaboration with the SciDAC's PERC, we will organize a forum with computer vendors to promote the adoption of ACTS tools. In a joint effort we will deliver a set of scientific benchmarks and results that will be valuable to the research and computer vendor communities. Among the computer vendors that have expressed an interest in this collaboration are IBM, COMPAQ, SUN and Fujitsu. Points of contacts will be identified as the event nears. [Marques, Drummond, and TS. Time: June to September].
C18	Visiting Scientist Host another Computational Science Graduate fellow during her/his practicum. [Marques and Drummond. Time: June to September].
Collaboration With Others	

FY04 (Cont.)

Activity, Milestone or Deliverable	
Box-ID	
I16	Update Information Center. Build on top of mechanism resulting from I6 . Update web pages for new ScidAC deliverables. Expand Information Center to include ANL and PNNL (see S16 and C16). [Marques, Drummond, ICTRM, and ORNL-POC. Time: yearlong].
I17	Update Repository. Update software versions and available tools in the distributed software and reference repository. [ICTRM. Time: yearlong].
I18	Enhance Information Center and Search Engine. Build on top of the search engine resulting from I7 . Expand the search domain to look for information available in bug reporting mechanisms. The on-line Bug Report mechanism will interact with similar mechanisms installed at tool developers' sites (see C13). Design and implement a database system to support this work with corresponding on-line web-based interfaces (see Subsection 5.10). [Marques, Drummond, TS, and ICTRM. Time: yearlong].
I19	ACTS Toolkit Workshop 2003. Add presentations, tutorials, information and feedback from the ACTS Toolkit Workshop 2003. (see O17). [ICTRM. Time: October].
I20	Tool Certification. Make tool certification document and process available on-line (see O21 and C15). [ICTRM. Time: mid-December].
I21	Automatic Tool Installation. Facilitate on-line scripts for automatic tool installation (see S17), using web forms. Document these scripts with information on supported platforms, library requirements and compilers. [Drummond, TS, and ICTRM. Time: February to September]
I22	Tool Evaluation. Make tool evaluations available as they are produced and reviewed (see S13). [ICTRM. Time: February to mid-June].
I23	Update Tool Certification and Retirement. Update ACTS Information Center with results from peer-reviewed panel certification (See S14). Thus, some tools will be certified and some will be retired. [Marques and ICTRM. Time: mid-June to mid-July].
I24	ACTS Toolkit Workshop 2004. In preparation for the ACTS Toolkit Workshop 2003 (see O25), design and create on-line forms, gather pointers to ACTS related information for prospective workshop participants, and heavily advertise the workshop. [Marques, Drummond, and ICTRM. Time: July to September].
I25	Education and Outreach Information. Update new tool information, and include presentations, tutorials and talks from outreach and education activities like O22 . [ICTRM. Time: mid-July to September].
O20	Progress Report FY03. Prepare a FY03 progress report to DOE MICS office. [Marques and Drummond. Time: October].
O21	SC2003. We will prepare demos, presentations, and posters. We will propose a Birds-of-a-Feather on tool Certification and will present at least one talk on the status of the project. We will also be available at the ORNL and NERSC booth for consultations. [Marques, Drummond and TS. Time: one week in November].
O22	Education and Outreach. Based on prior participations, we will present tutorials and presentations on numerical libraries for high performance Computing in the ACTS Toolkit. The venue will be the "All-Hands" meeting organized by NSF's National Partnership for Advanced Computing Infrastructure (NPACI), La Jolla, CA. [Marques and Drummond. Time: one day between February and March].
O23	Publication. Prepare a paper on Scientific Accomplishments facilitated by ACTS Tools. [Drummond, Marques and TS. Time: Summer]
O24	Computer Vendors Forum. Organized a forum with computer vendors on the acceptance and adoption of research tools inside the vendor-supported software libraries. We will identify specific point of contacts and participating high performance computer vendors at a later time (see C17). [Drummond, Marques and TS. Time: mid-June]
O25	ACTS Toolkit Workshop 2004. Organize and promote the ACTS Toolkit Workshop 2004. The Workshop will take place either the third or fourth week of September. [Marques and Drummond, July - September].
O26	Outreach. Participate and contribute to the SCICOMP 10. [Marques and Drummond. Time: sometime in August].

FY05

Activity, Milestone or Deliverable	
Box-ID	
S18	Tool Support. Continue support as described in S11 and include support for tools included in FY04. We anticipate providing different levels of support for tools in the Toolkit at four DOE High Performance Computing Centers (more centers will also be identified as this time nears). [Marques, Drummond, TS, ICTRM, ORNL-POC, ANL-POC, and PNNL-POC. Time: yearlong].
S19	Tool Installation. From S12 , continue tool installation on user demand, availability at NERSC's, CSS/ORNL's, ANL, and PNNL systems (more centers will also be identified as this time nears). [ICTRM. Time: yearlong].
S20	Automatic Tool Installation. Build on S17 , and update automatic tool installation scripts when necessary. Work on dynamical configuration of sets of tools for user download and installation on a variety of supported platforms. [Marques, Drummond, TS, and ICTRM. Time: yearlong].
S21	DOE Science Grid. Build on results from collaborations under C16 and provide support for applicable tools from these collaborations. Design and implement web-portal (using tools like Grid Portal Development Toolkit, http://www-rtg.lbl.gov/Grid/projects/GPDK/) for the Information Center and Repositories. Develop tutorials and material to help ACTS users to deploy Grid technology. [Marques, Drummond, TS, and ICTRM. Time: yearlong].
S22	Tool Evaluation. Perform tool evaluations and deliver corresponding reports from evaluations. Share reports with tool developers. [Marques, Drummond, and TS. Time: January to mid-May].
S23	Peer-review Panel. As in S9 and S14 . Organize peer-reviews. [Marques, Drummond, and TS. Time: mid-May to mid-June].
S24	Tool Enhancement. For tools requiring only minor revisions, ACTS support will work with tool developers to complete the recommendations from reviewers. The deliverables from these activities will result in better quality tools and potential even certified tools by ACTS. [Marques, Drummond, TS, and ICTRM. Time: end-May to mid-July].
C19	SciDAC and Tool Developers Collaborations. Follow up collaborations from C13 . Continue work with SciDAC ISICs defining mechanisms for long-term support and deployment of the tools. [Marques, Drummond, and TS. Time: yearlong].
C20	Collaborations. Continue collaborations in C14 . Results from these collaborations will be presented in conferences and publish in journals. Propose to organize domain specific workshops and lectures for DOE programs, Academic Organizations (like SIAM), NSF ITR, and CTRIS to increase visibility of the ACTS Toolkit. Export available technology in the Toolkit to Problem Solving Environments (e.g., PMATLAB, NETSOLVE, and the Equation Solver Interface [ESI]). [Marques and Drummond. Time: yearlong].
C21	DOE Science Grid. Continue work and collaborations established in C16 . Design a web-portal for the Information Center using Grid technology (S21). A first prototype will be shown in SC04 and the system will be fully available by SC05. Seek participation in a Global Grid Forum Workshop. [Marques, Drummond, and TS. Time: yearlong].
C22	Computer Vendor Forum. Continue promoting ACTS tools among computer vendors. Results from these interactions will produce valuable feedback to the computational science community. [Marques, Drummond, and TS. Time: yearlong].
C23	Tool Certification. Work with tool developers and other qualified computational scientists on reviewing and updating the tool certification criteria (C15). The work will deliver a document to be shared with the computational science community. We will propose a Birds-of-a-Feather session at SC 2004 (see O28) to present and gather further feedback from the computational science community. [Marques, Drummond and Tool Developers. Time: October to mid-November].
C24	Future Directions. Integrate a forum of tool developers and other qualified computational scientists on future directions of the ACTS project and relevant areas of interest not yet covered at this point in time by the its expanded infrastructure. We will motivate forum discussions with project accomplishments, status and feedback from the user community. [Marques and Drummond. Time: mid-June]
C25	Host another Computational Science Graduate fellow during her/his practicum, [Marques and Drummond. Time: June to September]
Solid Base of Tools	
Collaboration With Others	

FY05 (Cont.)

Box-ID	Activity, Milestone or Deliverable
126	Update Information Center. Build on top of Information Center resulting from I16 . Update web pages for new SciDAC deliverables. [Marques, Drummond, ICTRM, ORNL-POC, ANL-POC and PNNL-POC. Time: yearlong].
127	Update Repository. Update software versions and available tools in the software and reference repository. [ICTRM. Time: yearlong].
128	Enhance Information Center. From (C16 and C21) collaborations, implement a web portal using Grid technology, A demo will be available in SC04. Cross-reference other computational sciences Grid portals. [Marques, Drummond, and ICTRM. Time: yearlong].
129	Enhanced On-Line Materials. Design and implement on-line and interactive tutorials targeting users with different levels of expertise. Some tutorials will be used in several outreach and education venues like SC05. [Drummond and ICTRM. Time: yearlong]
130	ACTS Toolkit Workshop 2004. Add presentations, tutorials, information and feedback from the ACTS Toolkit Workshop 2004. (see O25) [ICTRM. Time: October].
131	Tool Certification. Make tool certification document and process available on-line (see O28 and C23). [ICTRM. Time: mid-December]
132	Automatic Tool Installation. Facilitate on-line scripts for automatic tool installation (see S20), using web forms. Document these scripts with information on supported platforms, library requirements and compilers. [Drummond, TS, and ICTRM. Time: February to September]
133	Tool Evaluation. Make tool evaluations available as they are produced and reviewed (see S22). [ICTRM. Time: February to mid-June].
134	Update Tool Certification and Retirement. Update ACTS Information Center with results from peer-reviewed panel certification (See S14). Thus, some tools will be certified and some will be retired. [Marques and ICTRM. Time: mid-June to mid-July].
135	ACTS Toolkit Workshop 2005. In preparation for the ACTS Toolkit Workshop 2005 (see O32), design and create on-line forms, gather pointers to ACTS related information for prospective workshop participants, and heavily advertise the workshop. [Marques, Drummond, and ICTRM. Time: July to September].
136	Education and Outreach Information. Update new tool information, and include presentations, tutorials and talks from outreach and education activities like O29 and the ones proposed in C20 and C21 . [ICTRM. Time: mid-July to September].
O27	Progress Report FY04. Prepare a FY04 progress report to DOE MICS office. [Marques and Drummond. Time: October].
O28	SC2004. We will prepare demos, presentations, and posters. Participate in DOE Lab booths. Also will propose a Birds-of-a-Feather on tool Certification and will present at least one talk on the status of the project. [Marques, Drummond and TS. Time: one week in November].
O29	Education and Outreach. Based on prior participations, we will present tutorials and presentations on numerical libraries for high performance Computing in the ACTS Toolkit. The venue will be the “All-Hands” meeting organized by NSF’s National Partnership for Advanced Computing Infrastructure (NPACI), La Jolla, CA. [Marques and Drummond. Time: one day between February and March].
O30	Publication. Prepare and publish a paper on High Performance computational tools. [Drummond, Marques and TS. Time: Summer]
O31	Computational Science Forum. Organized a forum with tool developers and other qualified computational scientists on future directions of the ACTS project (see C24). [Drummond, Marques and TS. Time: mid-June]
O32	ACTS Toolkit Workshop 2005. Organize and promote the ACTS Toolkit Workshop 2005. The Workshop will take place either the third or fourth week of September. [Marques and Drummond. July - September].
O33	Outreach. Participate and contribute to the SCICOMP 10. [Marques and Drummond. Time: sometime in August].
O34	Final Report. Prepare a final progress report to DOE MICS office and provide feedback from Computational Science Forum (O31). [Marques and Drummond. Time: October].

Appendix A

Tool Summaries

The following is a list of tools currently available in the ACTS Toolkit, a brief description of each tool, its status, and participation in the ACTS Project. For evaluation reports of some of the tools please visit the individual tool pages under <http://acts.nersc.gov/tools.html>.

A.1. Numerical Tools

- **Aztec** (developed at Sandia National Laboratories, SNL) is a library that provides algorithms for the iterative solution of large sparse linear systems arising in scientific and engineering applications. It is a stand-alone package comprising a set of iterative solvers, preconditioners and matrix-vector multiplication routines. Users are not required to provide their own matrix-vector multiplication routines or preconditioners in order to solve a linear system. The Aztec library is written in C and is also callable from Fortran. It is portable to most parallel platforms since it uses MPI to perform data communication. Overall, the package was designed to be easy to use. The user may input the linear system in a simple format and Aztec will perform the necessary transformations for the matrix-vector multiplication and preconditioning. After the transformations, the iterative solvers can run efficiently. If the input matrix is suitably partitioned, the efficiency can be further enhanced.

Our comments: The developers of Aztec have participated in activities organized by the ACTS Toolkit Project, although Aztec is now in the process of being superseded by the Trilinos solver framework [36]. Aztec's main strength is that it is a small and stand-alone package that contains all the components needed for the solution of a real sparse linear system of equations without requiring matrix-vector multiplications or preconditioning routines from the user. The most commonly used schemes to solve a large linear system on a distributed parallel machine include four major components: a mechanism to express the linear system, a matrix-vector multiplication routine, a preconditioner routine, and an iterative method to compute the solution by using the matrix-vector multiplication routine and the preconditioner. Aztec has carefully implemented all these components and overall exhibits good parallel efficiencies on many test problems.

- **Hypre** (developed at LLNL) is a library for solving large, sparse linear systems of equations on massively parallel computers. The main features of this library are: scalable preconditioners, implementation of a suit of common iterative methods (these include Conjugate Gradient and GMRES for symmetric and unsymmetric matrices, respectively), intuitive grid-centric interfaces, and dynamic configuration of parameters. Hypre works for users with different levels of expertise and has user-defined interfaces for multiple languages.

Our comments: The developers of Hypre have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback for the ACTS Information Center. Hypre is also actively participating in the deployment and instrumentation of component technology and is one of the software tools to be used in the TOPS SciDAC-funded project.

- **OPT++** (developed at SNL) is an object-oriented nonlinear optimization package for serial architectures. It solves optimization problems of the form $\min_{x \in R^n} f(x)$, $h_i(x) = 0, i = 1, \dots, p$, $g_j(x) \geq 0, j = 1, \dots, m$, in which the user specifies the function f (and, when available, its first and second analytical derivatives) and the functions h and g . OPT++ provides four solution algorithms: a Newton method, a finite-difference Newton method, a Quasi-Newton method, and a nonlinear conjugate gradient method.

Our comments: Although originally included in the ACTS Toolkit, Opt++ is not a parallel tool per se and therefore has not been considered for installation, evaluation and support. Also, we are not aware of specific examples of its use. Currently in the ACTS Toolkit, TAO (see below) is providing high performance implementations for most of the algorithms in Opt++'s; therefore Opt++ has been phased out.

- **PETSc**, the Portable, Extensible Toolkit for Scientific computation (developed at Argonne National Laboratory), provides sets of tools for the parallel (as well as serial), numerical solution of PDEs that require solving large-scale, sparse linear and nonlinear systems of equations. PETSc includes nonlinear and linear equation solvers that employ a variety of Newton techniques and Krylov subspace methods. PETSc provides several parallel sparse matrix formats, including compressed row, block compressed row, and block diagonal storage.

Our comments: PETSc is one of the most popular and matured tools currently available in the ACTS Toolkit. The PETSc developers have eagerly participated in many of the activities organized by the ACTS Toolkit Project and have received very good reviews and feedback from their participation in the ACTS Toolkit workshops. PETSc has a growing number of application users and development projects that re-use some of the functionality available in PETSc (like TAO and SLEPc [76]). The PETSc team has long worked on the development of interoperable interfaces with other ACTS tools, and perhaps has implemented the most tool-to-tool interoperability interfaces. The PETSc team is also actively participating in a more general solution to interoperability issues like the CCA. An additional feature of PETSc is its exemplary information, education and documentation website. Members of the PETSc team have strongly supported the ACTS Project and have provided useful material for the outreach and support of users.

- **PVODE** (developed at LLNL) actually refers to a trio of closely related solvers: PVODE, for systems of ordinary differential equations, KINSOL, for systems of nonlinear algebraic equations, and IDA, for systems of differential-algebraic equations. These solvers have some modules in common, primarily a module of vector kernels, and a generic linear system solver based on a scaled preconditioned GMRES method. PVODE is a solver for large systems of ordinary differential equations on parallel machines. It contains methods for the solution of both stiff and non-stiff initial value problems. Integration methods include the variable coefficient forms of the Adams and backward differentiation formula methods. The linear systems that must be solved during the implicit time stepping are solved with iterative, preconditioned Krylov solvers. The user can either supply a preconditioner or use one that is included in the PVODE package. PVODE is an extension of the sequential package known as CVODE, which has been widely distributed and used [47].

Our comments: The developers of PVODE have participated in some of the activities organized by the ACTS Toolkit Project. The PVODE trio is most useful for solving large differential and nonlinear algebraic systems that arise in a variety of applications. Important DOE applications include chemical kinetics, atmospheric chemistry, semiconductors, and structural or mechanical systems. Participants in the ACTS Toolkit workshops have also shown interest in PVODE, but to the best of our knowledge PVODE has not yet been widely used outside of groups with direct ties to the developers.

- **ScaLAPACK** (developed at the University of Tennessee, Knoxville; University of California, Berkeley; and ORNL) is a library of high performance linear algebra routines for distributed-memory message-passing MIMD computers and networks of workstations supporting PVM or MPI. It is a continuation of the LAPACK project [48], which designed and produced analogous software for workstations, vector supercomputers, and shared-memory parallel computers. The goals of both projects are efficiency (to run as fast as possible), scalability (as the problem size and number of processors grow), reliability

(including error bounds), portability (across all important parallel machines), flexibility (so users can construct new routines from well-designed parts), and ease of use (by making the interface to LAPACK and ScaLAPACK look as similar as possible). The ScaLAPACK library contains routines for solving systems of linear equations, least squares, eigenvalue and singular value problems. They can also handle many associated computations such as matrix factorizations or estimation of condition numbers.

Our comments: ScaLAPACK is intended for use in large-scale applications that require numerical manipulation of large dense or band matrices. Within NERSC's user community, we have seen various routines from ScaLAPACK being used. The applications using ScaLAPACK include material sciences, computational chemistry and astrophysics (see Section 1, for example). The most cumbersome task for ScaLAPACK users is the data distribution required by the software, but we have plans to develop mechanisms to facilitate that step.

- **SuperLU** (developed at the University of California, Berkeley; and NERSC) is a general-purpose library for the direct solution of large, sparse, nonsymmetric systems of linear equations on high performance machines. The library is written in C and is callable from either C or Fortran. The library routines perform an LU decomposition with numerical pivoting and triangular system solves through forward and back substitution. The LU factorization routines can handle non-square matrices but the triangular solves are performed only for square matrices. The matrix columns may be preordered (before factorization) either through library or user-supplied routines. This reordering for sparsity is completely separate from the factorization. Working precision iterative refinement subroutines are provided for improved backward stability. Routines are also provided to equilibrate the system, estimate the condition number, calculate the relative backward error, and estimate error bounds for the refined solutions.

Our comments: The developers of SuperLU have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback for the ACTS Information Center. SuperLU is a general-purpose software with a wide range of potential end users (see Section 1, for example).

- **TAO**, the Toolkit for Advanced Optimization (developed at ANL), focuses on large-scale optimization software, including nonlinear least squares, unconstrained minimization, bound constrained optimization, and general nonlinear optimization. There are a variety of software tools for solving the aforementioned problems; however, only TAO offers an Object-Oriented solution that provides a flexible optimization toolkit capable of addressing issues of portability, versatility and scalability in many computational environments. The algorithms in TAO place strong emphasis on the reuse of external tools where appropriate. TAO's design enables bidirectional connection to lower-level linear algebra support (such as parallel sparse matrix data structures) that is available in toolkits like PETSc, as well as higher-level application frameworks like POOMA.

Our comments: The developers of TAO have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback for the ACTS Information Center. TAO is rapidly acquiring a growing community of users and serves as a good example of quality software being developed on top of existing reliable and robust software.

A.2. Tools That Support Code Development

- **Global Arrays** (GA, developed at PNNL) is a library for writing parallel programs that use large arrays distributed across processing nodes and that offers a shared-memory view of distributed arrays without destroying their NUMA characteristics. The library has both Fortran and C interfaces. Originally developed to support arrays as vectors and matrices (one or two dimensions), it currently supports up to seven dimensions in Fortran and even more in C. GA offers two types of operations: collective

operations (require participation and synchronization of all processes) and local operations (may be invoked independently by all processes). GA also comes with a visualizer that uses trace files to animate array access patterns. Its main purpose is to analyze the impact of distribution on performance.

Our comments: The developers of GA have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center. GA is meant to complement rather than replace the message-passing model, and is interoperable with MPI. GA was originally developed for use in the NWChem computational chemistry package [49] but has also been used in other chemistry packages such as GAMESS-UK, Columbus, Molpro, Molcas, and SOCI.

- **Overture** (developed at LLNL) is a set of Object-Oriented tools for solving computational fluid dynamics and combustion problems in complex moving geometries. It has been designed for solving problems on a structured grid or a collection of structured grids. It can use curvilinear grids, adaptive mesh refinement, and the composite overlapping grid method to represent problems involving complex domains with moving components. Overture programs are written at a very high level, using data-parallel array expressions in the style of High Performance Fortran (HPF). They can achieve high performance (comparable to Fortran) thanks to a preprocessor (C++ to C++) called ROSE. Effectively, ROSE is a replacement for the expression template technique of POOMA. Overture has aggregate array operations and tightly integrated graphical features based on OpenGL. AMR++, a package that directly supports adaptive mesh refinement methods, is built on top of Overture.

Our comments: The developers of Overture have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center. In October 2001 the Overture Project was named one of the 101 most important discoveries supported by the DOE Office of Science in the past 25 years. Participants of the ACTS Toolkit workshops have been motivated by the Overture presentations; however, Overture is not easy to install and targets mostly C++ programmers. We believe that some efforts, like Babel and SILOON, will provide the necessary language interoperability that will make Overture available to a wider audience.

- **POET** (Portable Object-oriented Environment and Toolkit, developed at SNL) is a framework for encapsulating parallel algorithms. In this context, a parallel algorithm is a set of user-written components that are inserted into the POET framework. POET provides services such as starting the parallel application, running components in a specified order, and distributing data among processors. POET also provides support for various sorts of stencil operations and a sparse conjugate-gradient solver. The POET framework is an object model, implemented in C++. Each user-written component must follow the POET template interface. Once this has been done, these components are transferable from one POET application to another. All components are run in a Single Program Multiple Data (SPMD) fashion.

Our comments: Development of POET has ceased. Follow-on projects like InDEPS [50] will provide similar functionalities and more.

- **POOMA** (Parallel Object-Oriented Methods and Applications, developed at Los Alamos National Laboratories, LANL) is a collection of templated C++ classes for writing parallel PDE solvers using finite-difference and particle methods (see also Overture). POOMA programs are written at a very high level, using data-parallel array expression in the style of HPF or at a serial level using iterators on each CPU. They can achieve high performance (comparable to Fortran) thanks to a clever compilation technique called expression templates [78]. Moreover, they achieve true portability across serial, parallel and distributed architectures. POOMA offers two main parameterized types, fields (multidimensional

arrays representing grids with definable centering) and particles (a set of particles, with user-definable characteristics).

Our Comments: POOMA can be used to write high performance PDE solvers using finite-difference methods on structured, unstructured, or adaptive grids, particle methods, or hybrid methods that combine those. A few years ago, POOMA received some attention from the scientific computing community. Unfortunately, it discourages potential users by requiring specific compilers. Also, the developers have never participated in the activities organized by the ACTS Toolkit Project. At this point is not clear whether there is still development and support for this tool.

A.3. Tools That Support Code Execution

- **CUMULVS** (Collaborative User Migration, User Library for Visualization and Steering, developed at ORNL) is a software framework that enables programmers to incorporate fault-tolerance, interactive visualization and computational steering into existing parallel programs. The CUMULVS software consists of two libraries, one for the application program, and one for the visualization and steering front-end (called the "viewer"). CUMULVS handles collecting and transferring distributed data fields to the viewers and oversees adjustments to steering parameters in the application. It also manages the dynamic attachment and detachment of multiple independent viewers to a running parallel application. In addition, CUMULVS provides a user-directed checkpoint/restart mechanism to enable users to integrate fault tolerance to a running parallel application.

Our comments: The developers of CUMULVS have eagerly participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center and distribution during workshops and training sessions. Participants of the ACTS Toolkit workshops have shown some interest in using CUMULVS and in a few cases participants expressed a minor concern about the existing dependencies on the PVM library currently used by CUMULVS. Among others, CUMULVS is very effective for monitoring and steering remote parallel execution and allowing the construction of attractive user interfaces that have greatly helped users that are not familiar with all the particulars of a given code and/or the complexities involved in parallel computing. CUMULVS has been showcased by the CCA forum and the SciDAC Center for Component Technology for Terascale Simulation Software (CCTSS) Project demos at SC2001. Its developers have taken the lead in the MxN component technology discussion group given their vast experience with parallel and remote steering applications. A suitable version of CUMULVS will soon be made available on the NERSC systems by the tool developers and us.

- **Globus** [51] provides a bag of services for the creation of computational Grids and tools with which applications can be developed to access the Grid. Globus itself provides the following core services: communication services, resource allocation and process management. It also provides a standard interface to local resource management systems, allowing computational Grid tools and applications to express resource allocation and process management requests in terms of a standard application programming interface (API) and protocol, without constraining individual sites to any specific resource management tool. A standard protocol (LDAP, the Lightweight Directory Access Protocol) allows for discovering, publishing, and accessing information about the configuration and status of the computational Grid.

Our comments: Globus developers were awarded the 1997 Global Information Infrastructure (GII) Next Generation Award for their work in advancing the technology and application of high performance distributed computing. The Globus technology is now used worldwide and it has provided a leading infrastructure for the deployment of and collaborations around Grid technology. Several scientific computing centers, like NERSC, are still researching and planning the actual implementation and

participation in the Grid. At NERSC, security has been a concern. The developers have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center and organized a number of workshops on the Grid technology. There are three main kinds of users of this technology: the developers of Grid services, Grid-enabled application developers and users of the applications. Nowadays, all three kinds of users have to rely on their institutional agreement and provisions to work with the Globus technology, therefore the current support to Globus users requires a larger orchestration of personnel and resources than any of the other ACTS tools.

- **PAWS** (Parallel Application Work Space, developed at LANL National Laboratories) provides a framework for coupling parallel applications within a component-like model. Central to the design of PAWS is the coupling of parallel applications using parallel communication channels. The coupled applications can be run on different machines, and the data structures in each coupled component can have different parallel distributions. PAWS is able to carry out the communication without having to resort to global gather/scatter operations. Instead, point-to-point transfers are performed in which each node sends segments of data to remote nodes directly and in parallel.

Our comments: The developers of PAWS have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center. With the abrupt urge to couple multi-disciplinary codes there are high and increasing demands for tools like PAWS. The tool is currently under consideration by several research groups and we foresee the integration of functionalities available in other coupling efforts and PAWS through the use of component technology. The PAWS development team is actively participating in the CCA forum and the SciDAC CCTTSS.

- **SILoon** (Scripting Interface Languages for Object-Oriented Numerics, developed at LANL National Laboratories) gives scientists the ability to rapidly prototype and solve problems on high performance parallel computers. SILOON provides tools and run-time support for building easy-to-use external interfaces to existing numerical codes. The developers hope to enable scientists and other application programmers to easily access existing object-oriented scientific frameworks and numerical libraries written in C, C++, and Fortran. SILOON is a middleware package that may be used with the most popular scripting languages: Perl, TCL, and Python. It generates “glue code” to bind external library interfaces and provide run-time support and a computational server that is controllable under the scripting languages of choice. The first versions of SILOON were targeted for the POOMA and Overture frameworks.

Our comments: The developers of SILOON have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center. The developers of SILOON have promised to make it available at NERSC for testing and further education of users and developers. The experience and knowledge produced by this tool development effort are actively used in the CCA Forum, particularly for the research of language and component interoperability.

- **TAU** (Tuning and Analysis Utilities, developed at the University of Oregon and LANL National Laboratories) is a set of tools for analyzing the performance of C, C++, Fortran and Java programs. It is normally used in three steps: (1) instrument the program by inserting TAU macros into the program (this can be done automatically for C++ programs); (2) run the program, a trace file is then automatically generated; and (3) view the trace file with a TAU visualizer (RACY) or a third-party visualizer (such as VAMPIR). TAU collects much more information than what is available through *prof* or *gprof*, the standard Unix utilities. Also available through TAU are: per-process, per-thread and per-host information (supports *pthreads*), inclusive and exclusive function times, profiling groups that allow you

to organize data collection, access to hardware counters on some systems, per-class and per-instance information, separate data for each template instantiation, start/stop timers for profiling arbitrary sections of code, and support for collection of statistics on user-defined events.

Our comments: The developers of TAU have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center. TAU can be used by anyone writing a C, C++, Fortran or Java application and that wants to understand where the performance bottlenecks are. The hurdles of inserting profiling functions is minimized by two factors: only a single statement is needed to profile a function, and profiling can be disabled completely so that profiling insertions don't need to be taken out. Users of TAU have been enthusiastic about its usefulness.

- **Tulip** (developed at LANL National Laboratories) is an implementation of the HPC++ [77] interface specification for high performance machines. It provides two kinds of capabilities. First, inter-machine communication is supported via one sided (put and get) communication operations on typed global pointers and remote method function invocation. Second, intra-machine lightweight parallel threads are available with both Java-like and POSIX-like (Pthread) APIs. Fast global synchronization of threads and collective operations tie the two capabilities together. Tulip uses Nexus in its communication layer and can support applications across heterogeneous architectures.

Our comments: Development of Tulip has ceased. Follow-on projects like InDEPS [50] will provide similar functionalities and more.

A.4. Tools That Support Library Development

- **ATLAS** (Automatically Tuned Linear Algebra Software, developed at the University of Tennessee, Knoxville) and **PHiPAC** (Portable High Performance ANSIC, developed at the University of California, Berkeley, and the International Computer Science Institute) are tools for the automatic generation of optimized numerical software for modern computer architectures and compilers. These tools have both initially focused on level-three BLAS operations (matrix-matrix multiplications) and also a few routines from LAPACK, which have high potential for optimization. Traditionally, the optimization of these routines has been a tedious, architecture-dependent, hand coding process. Codes automatically generated by ATLAS and PHiPAC have been able to meet and even exceed the performance of the vendor-supplied, hand-optimized BLAS on a range of platforms. ATLAS and PHiPAC achieve performance through loop unrolling, explicit removal of unnecessary dependencies in code blocks, and use of machine-sympathetic C constructs. Code generators are parameterized and scripts are used to find the optimal choice of parameters for a given architecture and compiler. Research continues on these tools to extend them to new BLAS 3 operations and other basic linear algebra calculations, including sparse operations.

Our comments: ATLAS and PHiPAC are of primary interest to computer architects and administrators and are not tools that the application programmer would use, although the technology that they are enabling may well eventually spill over into tools for application programmers. We do not have plans to evaluate ATLAS and PHiPAC for their usefulness to application programmers because these tools are not targeted for this community.

- **Nexus** [51] is a portable library that provides the multithreaded communication facilities required to implement advanced languages, libraries, and applications in heterogeneous parallel and distributed computing environments. Its implementation supports multiple communication protocols and resource characterization mechanisms that allow automatic selection of optimal protocols. Nexus is intended for use by compiler writers and library developers, rather than application programmers.

Our comments: Nexus has been fully incorporated in the Globus Toolkit but its earlier versions are no longer supported by the developers (such as Nexus 4.1.x and earlier). Because of the target audience, Nexus has not directly participated in the educational and outreach venues organized by us. We have, however, promoted this tool as part of the bag of services under Globus and provided documentation and pointers to the developers' website.

- **PADRE** (Parallel Asynchronous Data and Routing Engine, developed at LANL) is a C++ layer for interfacing with libraries that distribute data on parallel computers. PADRE attempts to provide a uniform interface to several parallel decomposition libraries (KeLP, Multiblock PARTI, etc.) and the communication mechanisms used by them.

Our comments: The developers of PADRE have participated in many of the activities organized by the ACTS Toolkit Project and have provided us with feedback and material for the ACTS Information Center. Given the fact that PADRE targets library developers and offers a specific functionality, we have only presented and promoted it as part of OVERTURE. PADRE was not suitable for installation (by itself) in the NERSC production systems.

- **PETE** (Portable Expression Template Engine, developed at LANL) is an extensible implementation of the expression template technique (C++ technique for passing expressions as function arguments [78]). This technique uses C++ recursively-defined templates for transforming certain kinds of C++ statements into other statements with the same effect but higher performance. Evaluating array expressions in a single loop is of fundamental importance for performance, for example. Using simple C++ overloading it is difficult (practically impossible) to obtain this kind of translation. Furthermore, PETE has the advantage that the entire transformation is done by the C++ compiler, without requiring separate tools.

Our comments: In general, PETE is useful to writers of C++ scientific libraries who want to implement high-level mathematical abstractions and simultaneously obtain high performance (for example to provide array statements in the style of Fortran 90 or HPF). However, because of the extensive use of templates, it often requires large amounts of time and memory to perform the necessary template expansions. PETE is used in the POOMA framework. There have been plans to use it in Overture as well, but currently Overture uses a separate preprocessor called ROSE.

Appendix B

Accomplishments (2000 and 2001)

This Appendix lists the activities related to the ACTS Toolkit Project and carried out by Osni Marques and Tony Drummond.

B.1. Software Installation and/or Updating on NERSC Platforms:

- IBM SP (*seaborg*):
 - Aztec
 - PETSc
 - PVODE
 - ScaLAPACK
 - SuperLU
 - TAU
- Cray T3E (*curie*):
 - Aztec
 - CUMULVS
 - Hypr
 - PETSc
 - PVODE
 - ScaLAPACK
 - SuperLU
 - TAO
 - TAU
- Linux Cluster (*alvarez*):
 - Aztec
 - PETSc

B.2. Presentations Related to the ACTS Toolkit:

- *The ACTS Toolkit, an Overview*, NPACI All-Hands Meeting, Programming Tools and Environment Tools Session, San Diego, CA. February 2000.
- *The ACTS Toolkit (utilization, status and plan)*, for the Executive Committee, and *The ACTS Toolkit (What can it do for you?)*, for NERSC users. NERSC User's Group Meeting, Oak Ridge, TN. June 2000.
- *Overview of the ACTS Project*, DOE EPSCoR Workshop, Berkeley, CA. August 2000.
- *Tutorial on the ACTS Toolkit*, Parallel Summer Institute, San Diego Scientific Computer Center, CA. August 2000.
- *Tools for the Development and Solution of Scientific Applications: The ACTS Toolkit*. University of São Paulo, São Carlos, Brazil. August 2000.
- *What Can ACTS Do for You?* NERSC Theater, SC 2000, Dallas, TX. November 2000.
- Birds-of-a-feather on the ACTS Toolkit, SC 2000, Dallas, TX. November 2000.
- *Tutorial on High Performance Numerical Linear Algebra: Fast and Robust Kernels for Scientific Computing*, NPACI All-Hands Meeting, San Diego, CA. February 2001.
- *Overview of the ACTS Project*, Vistas on Computational Physics, NSBP Conference, Stanford, CA. March 2001.
- *Frameworks for High Performance Computing: Directions of the Next Generation of Scientific Codes*. Universidad Politecnica de Valencia, Spain. June 2001.

- *Algorithms for the Solution of Large Scale Eigenvalue Problems*, Advanced School on Recent Developments in Large Scale Scientific Computing, Porto, Portugal. July 2001.
- *Advanced Computational Testing and Simulation (ACTS) Toolkit*, Welcome lectures series to LBNL 2001 summer students, Berkeley, CA. July 2001.
- *Trends in High Performance Computing and the ACTS Toolkit Project*, CENAPAD-SP, Campinas, Brazil. August 2001.
- *The ACTS Toolkit: What can it do for you?* NERSC Theater, SC 2001, Denver, CO. November 2001.

B.3. Workshops on the ACTS Toolkit (see Appendix C for further details):

- *How Can ACTS Work for You?* Berkeley, CA. September 2000.
- *Solving Problems in Science and Engineering*, Berkeley, CA. October 2001.
- *High Performance Libraries for Science and Engineering*, LACSI Symposium, Santa Fe, NM. October 2001.

B.4. Synergetic Activities:

- Mr. David Strait joined LBNL/NERSC to work in the ACTS Toolkit Project from June to August 2000. David was a summer intern with a fellowship from the DOE Office of Science, Community College Initiative in Biotechnology, Environmental Science, and Computing. David was given a briefing on the ACTS Toolkit, and his assignment was to select a tool from the ACTS Toolkit, and examine its functionality and easiness of use. Based on his background and interests, David selected TAU (Tuning and Analysis Utilities).
- Dr. David Blackston joined the ACTS Toolkit Project as a postdoctoral fellow for two months (August and September, 2000). During this period, David worked on PVODE (a solver for large systems of ordinary differential equations on parallel machines) and also on a portable parallel adaptive N-body solver (Pbody) for its potential inclusion in the ACTS Toolkit.
- Dr. Yves-Henri Sanejouand from the *Centre de Recherche Paul-Pascal*, Pessac, France, visited LBNL from February 23 to March 11, 2001. His visit was the first phase of a joint work towards the optimization of a computer program allowing for the fast calculation of the low-frequency normal model of biological macromolecules. One of the ACTS tools that can be used in this application is ScaLAPACK. The efficient integration of all the steps required in the analysis may also be a good study case for the CCA standards.
- Dr. Jon Saenz from the Department of Applied Physics, University of the Basque Country, Spain visited LBNL from February 26 to March 8, 2001. Dr. Saenz has an interest in the functionality of some of the numerical tools available in the ACTS Toolkit and uses these tools inside his Python-based Toolkit for the analysis of climate data called PyClimate [46].
- Dr. Jose Roman from the *Universidad Politecnica de Valencia*, Spain, visited LBNL in October 2001. During his visit Dr. Roman worked on SLEPc [76], which is a software library that extends the functionality of PETSc for the solution of large-scale sparse eigenvalue problems. The library is installed on top of PETSc by means of a makefile system that allows for portability to a wide range of platforms.
- Dr. Xiaoyong Zhan from the University of Kansas submitted and was granted a NERSC Allocation Request in December 2000 to start the development of a three dimensional finite element simulator for land subsidence problems and to test the simulator with real subsidence data from the field. The computational approach adopted consists of a Galerkin finite element method and some of the tools available in the ACTS Toolkit for dealing with data distribution and systems of equations on parallel computers. Dr. Zhan was an attendee of the workshop *How can ACTS work for you?* in September 2000.

- Prof. Eric Carlson from the Department of Chemical Engineering, the University of Alabama. Support for DOE EPSCoR [40] proposal on the development of anisotropic diffusion equation solver using tools under DOE's ACTS Toolkit. The proposal has been accepted and is pending allocation of funding.

B.5. Activities Spurred at Other Centers:

- After attending the ACTS Toolkit Workshop in September 2000, Mark McKeown, from the University of Virginia (UVa), installed several of the ACTS tools on the UVa IBM SP2. Mark reported that several researchers use the tools, and that researchers in the UVa Department of Economics have also shown a special interest in the tools. He would like to have a visit from ACTS support to help users with the porting of their legacy codes to high performance computing using ACTS tools.
- Paulo Goldfeld, who attended the ACTS Toolkit Workshop in September 2000, works at the Courant Institute of Mathematical Sciences on the development of domain decomposition based preconditioners for mixed finite elements, including Neumann-Neumann type preconditioners. Paulo reported that he had found ACTS tools like PETSc very valuable, which allowed him to implement parallel programs without having to deal with message passing details and other related issues. He also stated that he would not have been able to develop his codes in a reasonable amount of time without ACTS tools, in particular PETSc.
- Joshua Breslau, from the Princeton Plasma Physics Laboratory, attended the ACTS Toolkit Workshop in September 2000, and reported that with the expertise he acquired at the workshop he was able to upgrade some of the routines in the M3D code, which is a multi-level 3D plasma physics code that now makes extensive use of PETSc for the parallelization and solution of an unstructured mesh problem. Joshua also stated that the parallelization of M3D would have been very difficult without PETSc and at the very least would have required several scientists spending a significant amount of their time reinventing existing numerical algorithms instead of doing physics.
- Mary Ann Leung is a DOE Computational Science Graduate Fellow who attended the ACTS Toolkit Workshop in October 2001. Ms. Leung has expressed interest in spending her summer practicum at LBNL doing research on the use of ACTS tools in the algorithms and codes that she has been developing for material sciences calculations. During her practicum, Ms. Leung will interact with us and with Dr. Andrew Canning, from NERSC's Scientific Computing Group. Dr. Canning has acquired extensive experience on the use of ScaLAPACK for material sciences calculations (see Section 1).

Appendix C

Data from ACTS Toolkit Workshops

The following subsections contain a summary of three workshops organized by Osni Marques and Tony Drummond on the ACTS Toolkit. Each workshop has a website which contains slides of the presentations, full lists of participants and their institutions, as well as details on the background of the participants.

C.1. ACTS Toolkit Workshop I:

How Can ACTS Work for You?

September 28-30, 2000, Berkeley, CA

<http://acts.nersc.gov/presentations/Workshop2000>

The goal of this workshop was to promote the use of the ACTS tools and make scientists aware of how these tools can be use in their research area. The particulars are:

- Target audience: mostly postdoctoral fellows and students.
- Total number of participants: 45 (excluding speakers), see distribution in Figure 3
 - Total number of students: 23
 - Total number of postdoctoral fellows: 13
 - Total number of researchers: 4
 - Others (staff SC centers, NERSC): 5
 - LBNL and UC Berkeley participation: 8 (2 from UC Berkeley)
- Laboratories and universities represented by speakers:
 - LBNL, 7 speakers
 - ANL, 5 speakers
 - LLNL, 2 speakers
 - LANL, 1 speaker
 - ORNL, 1 speaker
 - PNNL, 1 speaker
 - SNL, 1 speaker
 - University of Oregon, 1 speaker
 - University of Southern California, 1 speaker

The following is a list of research and educational institutions represented in this workshop. Institutions are classified in three groups: universities and non-DOE research centers, industry (high performance research component), and DOE laboratories. A pie chart of this distribution is also provided in Figure 4.

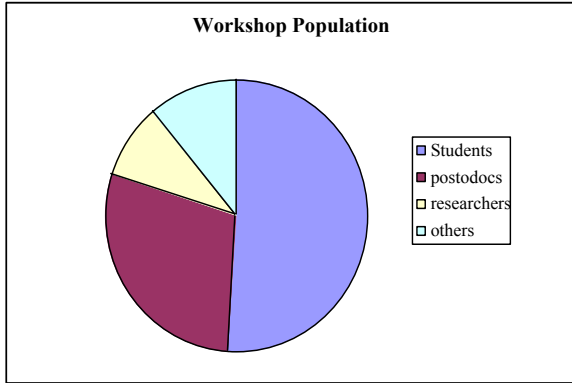


Figure 3. A breakdown of the participants by academic status




Figure 4. Workshop participation by institutional classification

Participating Universities and research centers:

- Auburn University
- Baylor College of Med
- Computer Sciences
- Courant Inst of Math
- Georgia Tech
- Hampton University
- Harvard University
- Iowa State University
- MIT
- NOAA/NOS/CS/CSDL/MMAP
- Old Dominion University
- Princeton University
- Purdue University
- San Diego Supercomputing
- Stanford University
- SUNY at Stony Brook
- Univ. MO-Kansas City
- University of California, Berkeley
- University of Houston
- University of Illinois
- University of Kansas
- University of Kentucky
- University of Michigan
- University of Southern California
- University of Tennessee
- University of Virginia
- Utah State University
- Vanderbilt University

Participating Industries:

- IBM Research
- Electrical Geodesics, Inc.

Participation DOE Laboratories:

- ANL
- Brookhaven National Laboratory
- LANL
- LBNL
- LLNL

Tool utilization:

The following tables summarize the list of tools that are being used (✓) by the participants and tools that they plan to use in the near future (X).

Application/ Field	ScalAPACK	SuperLU	Aztec	Hypre	PVODE	PETSc	TAO	Global Arrays	CUMULVS	Globus	Overture	PAWS	SILON	TAU
Computational Fluid Dynamics			✓			X								
Particle physics	✓								X					
Biophysics	✓ X						X	X	X	X				X
Earth Sciences (hydrology)	X	X	X			✓ X								
Mechanical Engineering	X	X				X X	X		X		X			
Astrophysics	✓	X	✓			X								X
Physics (solution of non-linear problems)	X						X							
Computational Physics						X								
Level Set Methods											X	X	X	

Application/ Field	TAU	SILON	PAWS	Overture	Globus	CUMULVS	Global Arrays	TAO	PETS _c	PVODE	HyPre	Aztec	SuperLU	ScalLAPACK
Plasma Physics	X					X			X					
Material Sciences				X					X					
Climate Modeling						X		X	X				X	X
Medical Physics						X		X	X				X	X
Computational Neurology						X		X	X				X	X
Atmospheric Sciences						X		X	X				X	X
Image Processing								X						X
Condensed matter theory					X									X
Atmospheric Physics					X									
Quantum Monte Carlo Techniques					X			X						
Ocean Modeling								X						
Meteorology				X					X					
Heavy Ion Fusion								X	X					X
Accelerator and Plasma simulator									X		X			
Solid States Physics									X					X
Chemistry									X					

C.2. *ACTS Toolkit Workshop II:
Solving Problems in Science and Engineering
October 10-13, 2001, Berkeley, CA*
<http://acts.nersc.gov/presentations/Workshop2001>

The goal of this workshop was to promote the use of the ACTS tools and make scientists aware of how these tools can be use in their research area. We extended the call for participation to researchers, in particular those funded by the SciDAC initiative. The particulars are:

- Target audience: postdoctoral fellows, students and SciDAC researchers.
- Total number of participants: 33 (excluding speakers), see distribution in Figure 5. Originally planned to host 38 participants but 5 people cancelled after the events of September 11, 2001.
 - Total number of students: 19
 - Total number of postdoctoral fellows: 6
 - Total number of researchers: 8
 - LBNL and UC Berkeley participation: 6 (2 from UC Berkeley)
 - Total number of SciDAC researchers: 5
 - Casey Meakin – University of Arizona (Adam Burrows supervisor)
 - Aimé Fournier – University of Maryland
 - Michel Kuhlen – University of California, Santa Cruz (Stan Woosley supervisor)
 - Water Elmer – University of Memphis
 - Bernhard Hientzsch – Courant Institute of Mathematical Sciences
- Laboratories and universities represented by speakers:
 - LBNL, 7 speakers
 - ANL, 5 speakers
 - LLNL, 5 speakers
 - LANL, 2 speakers
 - SNL, 2 speakers
 - ORNL, 1 speaker
 - PNNL, 1 speaker
 - University of Oregon, 1 speaker
 - Valencian University of Technology, 1 speaker

The following is a list of research and educational institutions represented in this workshop. Institutions are classified in three groups: universities and non-DOE research centers, industry (high performance research component) and DOE laboratories. A pie chart of this distribution is also provided in Figure 6.

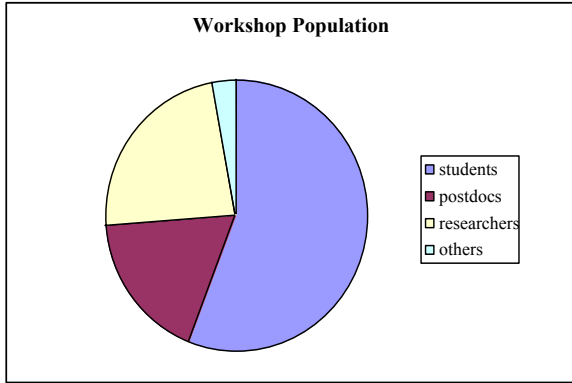


Figure 5. A breakdown of the participants by academic status.

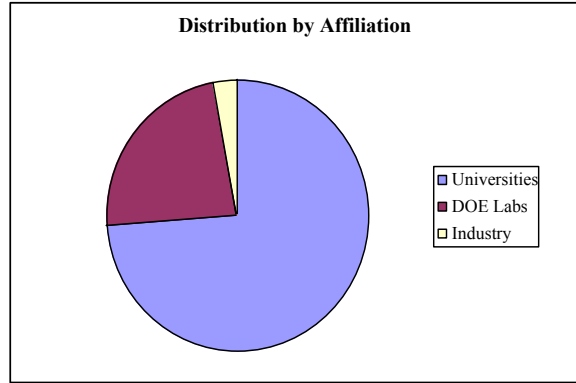


Figure 6. Workshop participation by institutional classification.

Participating institutions:

- Ames Laboratory
- Colorado State University
- Computational Science Graduate Fellowship Program (CSGP)
- Courant Institute of Technology
- Iowa State University
- MIT
- Old Dominion University
- San Jose State University
- Stanford University
- University of Arizona
- University of California, Berkeley
- University of California, Los Angeles
- University of California, Santa Cruz
- University of Colorado, Boulder
- University of Hawaii
- University of Kentucky
- University of Maryland
- University of Minnesota
- University of New Mexico
- University of New York
- University of Texas, Austin
- Wake Forest University

Participating industries:

- Science Applications International Corporation

Participating DOE laboratories:

- LANL
- LBNL
- LLNL
- SNL

Tool utilization:

The following tables summarize the list of tools that are being used (✓) by the participants and tools that they plan to use in the near future (X):

Application/ Field	ScalPACk	SuperLU	Trilinos	Hypre	PVODE	PETSc	TAO	Global Arrays	CUMULVS	Globus	Overture	PAWS	SILoon	TAU	Interoperability
Organic Chemistry	X					X	X							X	
Biological Oceanography	X				X	X	X								X ¹
Fusion Energy					X	X ²	X	X						X	
Astrophysics: Radiation Hydrodynamics	X	X				X									
Accelerator and Free electron laser Physics						X							X	X	
Astrophysics	X	X			X			X	X						
physics	✓	X													

¹ Making tools CCA compliant will make the developer's work much easier

² Definitely!

Application/ Field	ScalLAPACK	SuperLU	Trilinos	Aztec ³	Hypre	PVODE	PETSc	TAO	Global Arrays	CUMULVS	Globus	Overture	PAWS	SILON	TAU	Interoperability	ATLAS & PHiPAC ⁴
Scientific Computing	✓	X	X	✓			✓ X ⁵	X							X ⁷		
Electromagnetism	X	X			X ⁶		✓					X ⁶			X ⁷		✓
Finite Element									X						X		
Computational Chemistry	X							X							X		
CFD	X						X		X						X		
Chemistry						X	X		X						X		
Computational Biology	✓						X					X			X		
Computational Biophysics	X							X					X				
Combustion - Hydrodynamics					X		X					✓		X	X		
Applied mathematics	X	X			X		✓			X			X		X		
Computational Quantum Chemistry								X ⁸	X								

³ A tutorial on Trilinos was presented at the workshop and it included a brief introduction to AZTECOO which is currently intended to supersede AZTEC

⁴ ATLAS & PHiPAC were briefly introduced in Background Material and Next Steps talks

⁵ Will consider to using PETSc if it doesn't imply a major rewrite of my codes

⁶ May use only some parts from Hypre and Overture

⁷ May use some information from TAU

⁸ TAO can be used for the optimization procedure

Comments and suggestions on current *acts-support*:

- Would like some support prototyping their application: 12 participants
- Would possibly request some support: 4 participants
- Would like to visit LBNL/NERSC to get some help with ACTS tools: 2 participants
- Have been working with tool developers directly: 1 participant
- Have visited the ACTS Information Center: 15 participants
 - Has found this information useful: 15 participants

What is still missing from the ACTS Toolkit?

- More support for Fortran 90
- Multi-platform under PAWS
- Visualization tools
- More support for spectral methods
- A PDE toolbox at the level of discretization (approximating operators)
- Would like to use SLEPc (5 participants)
- Shared memory of BLAS and LAPACK to use within SMP nodes of large machines
- Tools to build a second order derivative for a Hessian or Jacobian routine

C.3. *ACTS Toolkit Workshop III:*

High Performance Numerical Libraries for Science and Engineering

LACSI Symposium, Santa Fe, New Mexico, 2001, October 18.

<http://acts.nersc.gov/lacsi>

This workshop presented talks on the development and application of state-of-the-art numerical tools that are currently used worldwide in science and engineering. Such mathematical tools for addressing more complex physical and societal phenomena, along with the growth of computing resources, are driving the continuous growth of the computational sciences community. This community includes scientists, designers and developers of high-end technology who require computerized modeling solutions.

The main objectives of this workshop can be summarized as:

- Promote the use of the ACTS tools,
- Present a few numerical frameworks like the Equation Solver Interface (ESI) and the Finite Element Interface (FEI), in which the numerical ACTS tools are used as numerical kernels
- Address issues of tool acceptance by the scientific and engineering communities.
- Address issues of interoperability
- Address issues of long tool development and support

The program of the workshop and the talks and abstracts are available at <http://acts.nersc.gov/lacsi>.

The morning talks were short tutorials on ScaLAPACK (Osni Marques), SuperLU (Sherry Li), Trillinos (Richard Lehoucq), PETSc and TAO (Tony Drummond) and an invited talk by Jose Roman from the Valencian University of Technology on SLEPc [76]. SLEPc is a tool for large-scale sparse eigenvalue problems built on top of PETSc. A 30-minute panel followed the presentations on *Promoting Reusability and Performance* (see next subsection for more details).

The afternoon session included mostly talks on numerical frameworks and interoperability. The first session presented talks on interoperability issues and Hypr (Andrew Cleary), the Equation Solver Interface, ESI (Robert Clay), followed by two talks on interoperability: SILOON (Craig Rasmussen) and the Common

Component Architecture (Matt Sottile). These presentations were also followed by a 30-minute panel discussion on *Promoting Software Interoperability*.

Summary from Panel I:

Tony Drummond was the moderator for this panel, and the panelists were Osni Marques, Sherry Li, Jose Roman and Rich Lehoucq. The following are some comments and questions posed to the panelists:

Q1. How does one balance the needs of performance against generality and portability?

- There is always an initial performance price to be paid when using a general-purpose tool, but in the long term, as users become more aware of their computational needs and use of the tools, the applications are better optimized.

Q2. Why are numerical libraries still difficult to install and port?

- Some of the ACTS tools use auto-configuration scripts that help users overcome these burdens.
- The more feedback developers receive from the user community, the higher the likelihood for developers to solve installation and porting problems across platforms.

Q3. Is there an item in your tool development agenda that addresses tool reusability?

- SLEPc uses functionality already available on other high performance numerical libraries like PETSc.
- Trilinos has plans to interoperate with other numerical libraries.
- While ScaLAPACK is a widely used tool, it is facing some development problems as their funding has ceased.

Q4. What are the driving parameters for tools to gain acceptance in the scientific and engineering communities?

- The kind of problems and environments where the tools are used.
- Real applications beyond the research group where they are initially designed.

Q5. Is it possible to extract modules/routines out of your toolkit without compromising the performance of the module?

- For most of tools represented at this panel, the answer is no.
- Is there a real need for this type of functionality?

Q6. How relevant are potential tool users in your development?

- Top applications and advancements in science influence the life of a tool.

Some additional comments from audience and panelists:

- What is still missing in all these efforts is the ability to tell the scientific world, particularly developers of long legacy code, what they have been doing wrong and why these new tools and programming paradigms are better. And even if it has been said, it needs to be promoted, perhaps through projects like ACTS.

Summary from Panel II:

Tony Drummond was the moderator for this panel, and the panelists were Robert Clay, Andrew Cleary, Matt Sottile, and Craig Rasmussen. The following are some comments and questions posed to the panelists.

Q1. Where is the starting point for developers of scientific and engineering applications?

- Think about code reusability, find out what is available out there, and learn how to use this technology.

- Users need to first define their computational needs well and then they should be able to select the tools they will use.
- The selection of the tool or tools to use should be independent of the language the users know. If there are any gaps, tools and mechanisms that support interoperability should provide the necessary bridges.

Q2. A message for other tool developers: What has worked? And what hasn't (in terms of interoperability)?

- Has worked:
 - Composability of the ESI is a good example.
 - Let users use their language of choice.
 - Modularity is preferred.
 - Object-Oriented technology has proven to be more efficient for developing libraries that provide a wide range of services to a wide range of communities.
 - Open source software is more likely to be used by other software developers.
 - Self-described calls.
- Hasn't worked:
 - Software designed to be optimal for a particular architecture, language or compiler.
 - Write giant monolithic software.
- Leverage standards that represent the consensus of a group of developers.
- No version 1.0 will be the right one for a long time.
- Writing code in a community environment is very hard because it requires several people with different backgrounds and interests to come to an agreement.

Q3. How does interoperability affects performance?

- Interoperability is still a hard problem between languages, operating systems and software tools. Making software interoperable and aware of the need to interoperate will only guarantee a future performance improvement and longevity of the software. For now the technology that promotes “interoperability” is still fighting several performance hurdles.
- The great payoff to interoperability is the code generations that will continue to benefit from interoperable software.

Appendix D

Early Interactions with DOE Projects and Others

The following is a summary of preliminary interactions with researchers participating in new high performance computing initiatives (such as SciDAC) and members of high performance computing centers. The goal of these interactions has been to establish initial grounds for future collaborations with the ACTS Project, promote the tools currently available in ACTS, and in the cases that involve software development, learn about plans for deployment of software technology and offer assistance and deployment of services that will be made available to a wide spectrum of users under this proposed and expanded ACTS Toolkit Project.

D.1. DOE Projects:

- ***Common Component Architecture Forum*** (CCA [41]) and ***Center for Component Technology for Terascale Simulation Software*** (CCTSS [52], SciDAC Project). Points of contact: Robert Armstrong (SNL, CCTSS PI), Ben Allen (SNL), David Bernhold (ORNL), James Kohl (ORNL), and Scott Kohn (LLNL).

This research focuses on the development of software component technology for high performance parallel scientific computing. The challenges being addressed by these research teams are of great relevance to the central software interoperability plans of the ACTS Toolkit Project as they propose a set of protocols for writing, maintaining and working with software component technology. Likewise, these research teams are equally interested in exploiting the infrastructure of the ACTS Project to provide long-term support to experienced users as well as means for centralizing and coordinating many of their activities. As one of our first steps in this collaboration, we have hosted under LBNL's Computing Infrastructure Support Department, the CCA repositories and testbed machines (initially two Linux-based computers) that will be instrumental for managing and coordinating the software development among the wide spectrum of institutions and scientists participating in the CCA Forum and the CCTSS.

In addition, we have been participating in several CCA related activities and e-mail lists, and have participated in a CCA tutorial that was presented at the workshop on the ACTS Toolkit at LBNL in October 2001. We will also participate in future CCA activities that involve technology transfer, education and dissemination.

- ***An Algorithmic and Software Framework for Applied Partial Differential Equations*** (APDEC [53], SciDAC Project). Point of contact: Phil Colella (LBNL, APDEC PI).

This project focuses on the development of a high performance algorithmic and software framework for solving partial differential equations arising from three important mission areas in the DOE Office of Science: magnetic fusion, accelerator design, and combustion. The Project PI has expressed interest in working with us and using the ACTS infrastructure for the deployment of the technology to be developed by APDEC, and in exploring and exploiting some tools currently available in the ACTS Toolkit. He also sees this infrastructure as a venue for interacting with other SciDAC-funded research groups and suggested that he would like active participation from a member of our staff in the APDEC activities and with his research group to better collect the necessary expertise in the algorithms and software tools being deployed here.

- ***Terascale Optimal PDE Simulations*** (TOPS [54], SciDAC Project). Points of contact: Prof. David Keyes (Old Dominion University, TOPS PI), Prof. Jack Dongarra (University of Tennessee, Knoxville), Prof. James Demmel (University of California, Berkeley) and Sherry Li (LBNL).

This research focuses on developing and implementing optimal or near-optimal schemes for PDE simulations, including optimization of PDE constrained systems, eigenanalysis, and adaptive time integrations. It will pay close attention to optimal complexity solvers for the nonlinear PDE that arise in many DOE Office of Science application areas like fusion, accelerator design, global climate change and reactive chemistry. This project is a compelling example of what has been promoted in the ACTS Toolkit Project, mainly software reusability, interoperability and scalability. Tools currently in the ACTS Toolkit like Hypr and PETSc (see Appendix A) will be used and extended in TOPS, and these and other components of TOPS have been actively participating in software component technologies to address interoperability issues. This gives both TOPS and ACTS a solid starting point for collaboration and exploitation of such technologies.

The early interactions and participation of most members of this project in the ACTS Toolkit have also enabled us to continue working and collaborating under the ACTS infrastructure. The Project PI expressed a great interest in the support provided by the ACTS infrastructure for the deployment and dissemination of the software tools and expertise to be developed and acquired during this project.

- ***Terascale Simulation Tools and Technology Center*** (TSTT [55], SciDAC Project). Points of contact: James Glimm (BNL, TSTT PI) and Lori Freitag (ANL, Co-PI).

The primary objective of TSTT is to develop technologies that enable application scientists to easily use multiple mesh and discretization strategies within a single simulation on terascale computers. The researchers involved in TSTT have expressed some interest in learning more about the infrastructure of the ACTS Project and also expressed their appreciation for these services to be available to them at no extra cost. Plans for future collaboration and integration with this team will be opportunely discussed as the TSTT's researchers identify their needs for deployment and outreach.

- ***Performance Evaluation Research Center*** (PERC [56], SciDAC Project). Points of contact: David Bailey (LBNL, PERC PI) and Erich Strohmaier (LBNL).

The goal of this SciDAC Center is to develop a science for understanding performance of scientific applications on high-end computer systems, and develop engineering strategies for improving performance on these systems. Initially, this project will interact with scientific simulations and experiments arising in the areas of climate modeling, high energy physics, astrophysics, genomics and proteomics. This project also builds on some of the tools available under the ACTS Toolkit like ATLAS (see Appendix A).

The PERC researchers have expressed interest in interacting with ACTS for long-term support and deployment of the tools and products to be developed by this center. There is also interest in using the mechanisms of the ACTS Information Center to educate, support and disseminate the results. The solid base of tools and the gallery of applications under ACTS could potentially enhance the research being done at PERC, and in turn PERC could further enhance the levels of optimal resource utilization at the application and tool levels.

ACTS tools and their users will provide more accurate and consistent parameters and metrics for the performance analysis, modeling and characterization of scientific applications. This is because the reusability of quality and near-optimal software removes behavioral effects caused by less experienced and often less engineered coding practices. For instance, there may be instructions in legacy and quickly

written codes that are not needed or are even the source of bugs. The codes using ACTS tools inherit the robustness and bug-free features of the ACTS tools, eliminating the aforementioned hurdles to performance modeling. Furthermore, ACTS makes available a suitable infrastructure for PERC to interact with other DOE projects as well as a large pool of users from the computational sciences community.

- **Scientific Data Management Enabling Technology Center** (SciDAC Project). Point of contact: Ari Shoshani (LBNL, Project PI).

This project focuses on the development of a framework for the unification, deployment and reuse of scientific data management software. The PI expressed that the project is still at an early stage to have a defined idea of what the software deliverables will be and their means for deployment. He is familiar with the activities under ACTS and is aware of our current plans for expansion of the project, and we agreed that it would be necessary to recontact his research group in the near future.

- **BABEL** [44]. Point of contact: Scott Kohn (LLNL).

This project addresses language interoperability and re-use for high performance parallel scientific software. Its purpose is to enable the creation, description, and distribution of language independent software libraries. As mentioned earlier in this proposal, we have plans for incorporating language interoperability and component technologies to the ACTS Toolkit Project to be able to reach out and educate a wider pool of computational scientists. As an initial step, we have participated in the BABEL e-mail list and Kohn presented a tutorial on BABEL at the ACTS Toolkit Workshop at LBNL in October 2001. In the immediate future, we will look for ways to make this technology available to existing ACTS Toolkit users and investigate the technology currently used in Alexandria for the exploitation and distribution of software component technology (see Subsection 5.4).

- **DOE Science Grid** [42]. Points of contact: Keith Jackson (LBNL) and William Johnson (LBNL).

This project's main goal is to provide the advanced distributed computing infrastructure based on Grid middleware and tools to enable the degree of scalability in scientific computing necessary for DOE to accomplish its missions in science. Currently, Globus is part of the ACTS Toolkit. However, it has not been previously fully deployed on DOE scientific computer centers like NERSC because of security and policy related issues that are now addressed by the DOE Science Grid. In addition to bringing the Grid technology and services to ACTS users, we will be able to naturally expand our scope of services to other DOE facilities, thus to a wider community of users. We will incorporate, when appropriate, tools from the DOE Science Grid to help ACTS users deploy Grid technology [13]. We also foresee using tools like PyGlobus [43] to educate users on the technology. Furthermore, our collaborations with the DOE Science Grid, CCTSS and BABEL research teams will deliver a fully interoperable infrastructure, both in software and hardware, for the development and testing of scientific applications.

- **Advanced Computational Research in Fusion** (SciDAC Project, PI Mitch Pindzola). Point of contact: Dario Mitnik (Dept. of Physics, Rollins College).

Dr. Dario Mitnik attended the workshop on the ACTS Toolkit in September 2000. Since then he has been actively using some of the ACTS tools, in particular ScaLAPACK [11], for which he has provided insightful feedback. Dario is currently working on the development, testing and support of new scientific simulation codes related to the study of atomic dynamics using time-dependent close coupling lattice and time-independent methods. He reports that this work could not be carried out in sequential machines and that ScaLAPACK is fundamental for the parallelization of these codes. Additionally, ScaLAPACK

allowed him to easily port his codes to a Sun Beowulf cluster and an IBM SP2. There are several tools currently present in the ACTS Toolkit for which Dario has already identified some applicability inside his codes, including Global Arrays and some of the Grid services provided by Globus. Dario is also interested in working closely with the ACTS Project in order to find support for his new developments and training in his endeavors with Grid technology.

D.2. DOE Institutions:

- **National Energy Research Scientific Computing Center** (NERSC [21]). Points of contact: Horst Simon (NERSC Director), Francesca Verdier (NERSC User Services Group) and William Miner (DOE Program Manager for NERSC).

NERSC provides high performance computing tools and expertise that enable computational science of scale, in which large, interdisciplinary teams of scientists attack fundamental problems in science and engineering that require massive calculations and have broad scientific and economic impacts. NERSC has plans to support some of the SciDAC teams. ACTS will interact with NERSC in the deployment of these tools and particularly will work closely with the User Services Group at NERSC to provide staff training and future support to users for the tools resulting from SciDAC. We will continue to coordinate *acts-support*, now to be extended to other DOE scientific computing centers like the Center for Computational Sciences at ORNL (See Subsection 5.5)

- **ORNL Center for Computational Sciences** (CCS/ORNL [37]). Points of contact: Arthur Bland (ORNL) and Kwai Wong (ORNL and University of Tennessee, Knoxville).

Mr. Bland is the Computer Resources Manager for the Center for Computational Sciences at ORNL. He appointed Kwai Wong as the point of contact for the ACTS Toolkit Project. Together, we have started collaboration for the deployment and support of ACTS tools at their center. We will first identify the current users of the ACTS Toolkit and incorporate relevant information from ORNL expertise into the ACTS Information Center. We will later coordinate weekly conference calls to work on the expanded ACTS infrastructure that will support projects like the ones outlined in this Appendix and future ones to be initiated upon funding of this proposal.

Similar contacts will also be established with the PNNL Environmental Molecular Sciences Laboratory (EMSL [38]) and ANL's Mathematics and Computer Science Division [39]. See Subsection 5.5.

D.3. Other Projects:

- **The Center for Information Technology Research in the Interest of Society** (CITRIS [57]). Point of contact: Prof. James Demmel (University of California, Berkeley, CITRIS Chief Scientist and Associate Director).

The goal of this project is to create information technology to tackle society's most critical needs. CITRIS research will focus initially on the areas of energy efficiency, transportation, seismic safety, education, health care, and environmental monitoring. CITRIS was conceived as a response to California Governor Gray Davis's call for visionary new University of California Centers for Science and Innovation. CITRIS will deliver the technical foundations to meet California's urgent public-sector needs in years, not decades, and lead the way in creating the information technology infrastructure needed for a vibrant 21st century society. Several of the numerical simulation, data driven problems to be addressed by this project can greatly benefit from tools already available in the ACTS Toolkit and, likewise, the ACTS Project can benefit from interactions with a larger pool of scientific and industrial users, thus pushing their promotion and acceptance outside the DOE community.

Prof. Demmel has participated in the ACTS Toolkit Project and is eager to continue this collaboration. He has long promoted the creation of an infrastructure like ACTS and the reusability of robust tools (see for example [8] and [9]). Prof. Demmel is also actively participating in several DOE-funded projects and initiatives.

- ***Equation Solver Interface*** (ESI [58]). Points of contact: Dr. Robert Clay (TeraScale, LLC) and Ben Allan (SNL).

TeraScale, LLC [59], is a software development company founded in New Mexico with offices in Cedar Crest, New Mexico and in Tracy, California. The company develops software for high performance, distributed memory parallel computing. This company has actively participated in the development of software specification for the Parallel Mesh Object (PMO) project, Finite Element Interface (FEI) to linear solvers, and the ESI for linear algebra software for engineering applications. The ESI has provided a well-founded example of building software technologies on top of tools available in the ACTS Toolkit (like ScaLAPACK, SuperLU, PETSc and Aztec). Both researchers have participated in ACTS related activities (see LACSI Workshop in Appendix C) and have expressed an interest in continuing to work with infrastructures like the ACTS Toolkit to promote their interfaces and educate users on the use and development of component software technology. ACTS benefits not only from their expertise but also from their interactions with a large pool of applications and users in the various DOE offices and the industry.

- ***Earth and Space Sciences (ESS), National Aeronautics and Space Administration*** [60]. Point of contact: Dr. James Fischer (ESS Program Manager).

This effort employs advanced computers to further the understanding of and ability to predict the dynamically interacting physical, chemical, and biological processes that drive these natural systems. NASA's Goddard Space Flight Center (GSFC) serves as the lead center for the ESS Project and collaborates with the Jet Propulsion Laboratory (JPL). Dr. Fischer has expressed an interest in learning more about the ACTS Project, and has participated in the ACTS Birds of a Feather at SC2000.

- ***Education Outreach and Training Partnership for Advanced Computational Infrastructure*** (EOT-PACI [82]). Point of contact: Roscoe Giles (Boston University).

The mission of the EOT-PACI is to develop human resources through the innovative use of emerging information technologies to understand and solve problems. We have plans to organize an introductory course on High Performance Computing and at least one Summer Institute where this course will be presented.

D.4. Other Institutions:

- ***Center for Computational Sciences, University of Kentucky*** (CCS/UK [61]). Point of contact: Prof. Craig Douglas (CCS/UK Associate Director).

Prof. Douglas has been a supporter of the ACTS Toolkit Project and his research team has participated in activities organized under ACTS. Prof. Douglas is the creator and maintainer of MGNet [27]. He edits the corresponding monthly electronic newsletter. MGNet is a repository for information related to multigrid, multilevel, multiscale, aggregation, defect correction, and domain decomposition methods. These methods are used primarily by scientists and engineers to solve partial differential equations on serial or parallel computers. Prof. Douglas is actively looking for ways to collaborate with the ACTS

Project and will host the Institute for Mathematics and its Applications' (IMA) summer graduate program, July 1-26, 2002 [62]. The institute will give emphasis to scientific computing and applications, including a week on biomathematical computing and a two-day workshop on the ACTS Toolkit. Lectures from IMA will be broadcasted via the Access Grid. We have plans for future collaborations that will bring the ACTS technology to CCS/UKY and integrate the MGNNet expertise into the ACTS Information Center.

- ***San Diego Supercomputing Center*** (SDSC [63]). Point of contact: Dr. Francine Berman (SDSC Director) and Mr. Donald Frederick (Leader, SDSC Scientific Computing Services)

SDSC provides computational resources for advancing knowledge through the development and application of advanced computing technologies. It also leads the National Partnership for Advanced Computing Infrastructure (NPACI [64]), which is an effort to deploy a national metacomputing infrastructure. Our contacts at this center have expressed some interest in the ACTS technology. As a consequence of our past interactions and participation in their educational events, SDSC has now installed on its machines a number of the tools currently available in ACTS. We plan to continue working with SCSC to reach out the scientific computing community by providing close support and information on the ACTS tools and participating in their educational programs, such as the SDSC's Computational Summer Institute, NPACI All-Hands Meetings, and Ptools workshops. We foresee our collaborations with this center as very important for reaching researchers outside the DOE community.

- ***Pittsburgh Supercomputing Center*** (PSC [65]). Point of contact: Dr. Ralph Roskies (PSC Scientific Director, Executive Management).

PSC joined a partnership with the National Energy Technology Laboratory (NETL [66]), Carnegie Mellon University, West Virginia University, and the West Virginia Governor's Office of Technology, creating the Super Computing Science Consortium. PSC also partners with the Commonwealth of Pennsylvania to offer free grants of supercomputing time to researchers in Pennsylvania. The PSC Industrial Affiliates Program has been providing world-class computing power and balanced infrastructure for industry for over a decade. The center provides an integrated array of services necessary for a company to approach problems previously viewed as beyond reach. The director and staff are quite eager to collaborate with us and promote the use of ACTS tools. This collaboration will benefit PSC with *acts-support* related services and anticipated better resource utilization by their users' applications. In turn, the ACTS Project will benefit by acquiring more expertise and recognition from a pool of scientific applications and users. We have plans to actively participate in PSC's educational and outreach programs.

- ***National Science Foundation*** (NSF [67]). Points of contact: Dr. Frederica Darema (NSF Senior Science and Technology Advisor, Program Manager for ITR-DDDA), Cliff Jackobs (NSF Program Manager for the National Center for Atmospheric Research)

NSF supports innovation by individual investigators and national centers alike, which educate and employ young computer specialists who will form the next generation of leadership for information technology. NSF began its Information Technology Research (ITR) initiative in fiscal year 2000. In 1999 the President's Information Technology Advisory Committee reported that without new federal investments like ITR, the U.S. would be in danger of losing its international leadership in computer science. The aforementioned program managers are interested in the activities and the technology being developed under the ACTS umbrella and have invited their participating research teams to attend workshops on ACTS. We have participated in several review panels for the NSF ITR initiative and other NSF programs.

- **European Center for Research and Advanced Training in Scientific Computing** (CERFACS [68])
Points of contact: Pierre-Henry Cros (CERFACS Associate Director), Prof. Iain Duff (CERFACS Parallel Algorithm Project, Leader).

CERFACS is located in Toulouse, France, and has around 65 international researchers working on efficient algorithms for solving large-scale scientific problems. This involves the evaluation of existing, and the development of new tools that exploit high performance serial, vector, and parallel computers. In order to optimize the exploration of high performance computing, CERFACS works in an interdisciplinary manner, involving practitioners, applied mathematicians, numerical analysts and computer scientists. Mr. Cros has showed interest in employing ACTS tools in industrial applications CERFACS has been working on. Prof. Duff's research interests are sparse matrices, parallel computing, scientific computation and mathematical software. He sees the ACTS Project as a good opportunity for technology transfer.

- **Parallel Computing Group, Department of Information Systems and Computation, Valencia University of Technology** [70]. Points of contact: Prof. Vicente Hernandez (Group Leader), Prof. Vicente Vidal and Prof. Jose Roman.

This research group has several high performance computing applications that serve both industry and academia in Spain and in European partnerships. The group is already familiar with several of the tools currently available in the ACTS Toolkit, has worked closely with some of the ACTS tool developers (such as PETSc) and would like to collaborate with the ACTS Project in the deployment of their software tools and dissemination of their research (see Subsections A.1, B4, C.2 and C.3).

- **Department of Computer Science and Artificial Intelligence, University of Alicante** [69]. Points of contact: Prof. Violeta Migallon, and Prof. Jose Penades.

The professors and their students have several projects in numerical analysis and artificial intelligence that could benefit from the software technology available in ACTS. In particular, there are two graduate students interested in deploying the ACTS technology in their research, image processing and data reanalysis.

D.5. Other Supporters of the Proposed Work:

- Prof. Jack Dongarra, Department of Computer Science, University of Tennessee, Knoxville; Innovative Computing Laboratory (ICL [71]), Director; Center for Information Technology Research (CITR [72]), Director; and Adjunct R&D Participant at ORNL, Computer Science and Mathematics Division.

Prof. Dongarra is a lead scientist in the field of scientific computing and an active member of several high performance computing organizations, institutions and committees worldwide. He has participated in some of the activities organized by the ACTS Toolkit Project and also in the development of tools like ScaLAPACK, ATLAS and PhiPAC, which are currently part of the ACTS Toolkit. He is also a co-PI in several of the SciDAC Projects. Prof. Dongarra has always expressed an interested in collaborating and supporting an infrastructure like the one proposed in this document.

- Dr. William Gropp, Associate Division Director and Senior Computer Scientist from the Mathematics and Computer Science Division, ANL.

Dr. Gropp has actively participated in several activities organized by the ACTS Toolkit Project. He has vast experience designing software standards and is a member of the PETSc development team and the Center for Programming Models for Parallel Computing [31]. Dr. Gropp has supported the ACTS Toolkit Project and has expressed an interest in continuing to participate and collaborate with us in making this project a continued success.

D.6. Interactions with Computer Vendors:

- **IBM.** Point of contact: Luis DeRose (IBM T.J. Watson Research Center).

Dr. DeRose works at IBM's Advanced Computing Technology Center (ACTC [73]) on application performance tools, in particular on the *Simulation Guided Memory Analyzer* and on the *Hardware Performance Monitor Toolkit*. The former is a data-centric tool under development for performance prediction and identification of bottlenecks, problems, and inefficiencies in a code due to data layout in the memory hierarchy, and for proposing performance improvement solutions in current and future architectures. The latter is a collection of libraries and utilities to support performance data capture, analysis, and presentation of hardware performance metrics from applications written in Fortran, C, and C++. These tools are important for complementing features that are currently available in ACTS (namely TAU, see Appendix A) as well as to evaluate the performance of ACTS tools. Dr. DeRose attended the workshop on the ACTS Toolkit in September 2000.

- **Compaq.** Point of contact: William Blake (Vice President of World Wide High Performance Technical Computing Group).

Contacted personnel in the High Performance Tool research department and we will continue initial interactions through the workshops and seminars organized by the Pittsburgh Supercomputing Center. Additional points of contact will be identified at a later date.

- **Intel.** Point of contact: Richard Wirt (Director of Microcomputer Software Labs, Santa Clara, CA).

Contacted at the Parallel Development Tools for High Performance Computing Workshop at SC 2001. He expressed an interest in participating in a forum for high performance tools and becoming more acquainted with the ACTS Project.

- **Fujitsu.** Point of contact: Clifford Chen (San Jose, CA).

After our initial contact, Dr. Chen has shared information on the ACTS Toolkit with his research team in Japan and has expressed, in the name of his colleagues, a strong interest in participating in forums and discussions on HPC tools.

References

- [1] T. Rescigno, M. Baertschy, W. Isaacs and W. McCurdy, 1999. *Collisional breakup in a quantum system of three charged particles*, Science, 286:2474-2479.
- [2] P. de Bernardis, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, B. P. Crill, G. De Gasperis, P. C. Farese, P. G. Ferreira, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, A. E. Lange, L. Martinis, S. Masi, P. V. Mason, P. D. Mauskopf, A. Melchiorri, L. Miglio, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, D. Pogosyan, S. Prunet, S. Rao, G. Romeo, J. E. Ruhl, F. Scaramuzzi, D. Sforza and N. Vittorio., 2000. *A flat Universe from high-resolution maps of the cosmic microwave background radiation*, Nature, 404: 955–959.
- [3] Canning, W. Mannstadt and A. J. Freeman, 2000. *Parallelization of the FLAPW Method*, Computer Physics Communications, 130:233-243, Elsevier, The Netherlands.
- [4] M. Heath, 2002. Scientific Computing: An Introductory Survey, 2nd Edition, McGraw-Hill, New York; <http://www.cse.uiuc.edu/heath/scicomp>.
- [5] M. Henderson, C. Anderson and S. Lyons, Edits, 1999. *Object Oriented Methods for Interoperable Scientific and Engineering Computing*, Proc. of the 1998 SIAM Workshop, SIAM, Philadelphia, PA.
- [6] W. Press, S. Teukolsky, W. Vetterling and B. Flannery, 1992. *Numerical Recipes in Fortran 77*, Cambridge University Press.
- [7] H. Zha, O. A. Marques and H. D. Simon, 1998. *Large-Scale SVD and Subspace-based Methods for Information Retrieval*, Solving Irregularly Structured Problems in Parallel (Proceedings of the Irregular'98 Conference), Lecture Notes in Computer Science, 1457:29-42, Springer Verlag.
- [8] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine and H. Van der Vorst, 1994. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, 2nd Edition, SIAM, Philadelphia, PA.
- [9] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, Editors, 2000. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, Philadelphia.
- [10] M.W. Berry and M. Browne, 1999. *Understanding Search Engines: Mathematical Modeling and Text Retrieval*, SIAM Book Series: Software, Environments, and Tools.
- [11] D. M. Mitnik, D. C. Griffin and N. R. Badnel, 2001. *Electron-impact excitation of Ne^{5+}* , J. Phys. B: At. Mol. Opt. Phys. 34:4455-4473 (28 November).
- [12] I. Foster and C. Kesselman, Eds, 1999. *The Grid, Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, Inc, San Francisco, CA.
- [13] Ian Foster, 2002. *The Grid: A New Infrastructure for 21st Century Science*, Physics Today, February.
- [14] MADCAP, <http://www.nersc.gov/~borrill/cmb/madcap.html>
- [15] SciDAC/DOE, <http://www.er.doe.gov/scidac>
- [16] LACSI, <http://lacs.lanl.gov/symposium>
- [17] ACTS Toolkit Review, <http://acts.nersc.gov/presentations/PI2000>
- [18] ACTS, Next Steps, http://acts.nersc.gov/presentations/Workshop2001/slides/drummond_2.pdf
- [19] CAN/NASA, <http://sdcd.gsfc.nasa.gov/ESS/CAN.html>
- [20] ITR/NSF, <http://www.itr.nsf.gov>

- [21] NERSC, <http://www.nersc.gov>
- [22] Netlib, <http://www.netlib.org>
- [23] HPC-Netlib, <http://www.nhse.org/hpc-netlib>
- [24] NHSE, <http://www.nhse.org>
- [25] CSIR, <http://www.csir.org>
- [26] GAMS, <http://gams.nist.gov>
- [27] MGNet, <http://www.mgnet.org>
- [28] NEOS, <http://www-fp.mcs.anl.gov/otc/Guide>
- [29] Ptools, <http://www.ptools.org>
- [30] PTLib, <http://www.nhse.org/ptlib>
- [31] PMODELS, <http://www.pmodels.org>
- [32] CSEP, <http://csep1.phy.ornl.gov>
- [33] CS267, http://www.cs.berkeley.edu/~demmel/cs267_Spr99 and <http://www.cs.berkeley.edu/~dbindel/cs267ta>
- [34] SMA5505, <http://www.mit.edu/~cly/18.337>
- [35] CS454, <http://www.cse.uiuc.edu/cs454/notes>
- [36] Trilinos, <http://www.cs.sandia.gov/Trilinos>
- [37] CCS/ORNL, <http://www2.ccs.ornl.gov>
- [38] EMSL/PNL, <http://www.emsl.pnl.gov>
- [39] MCS/ANL, <http://www.mcs.anl.gov/computing>
- [40] EPSCoR, <http://www.science.doe.gov/EPSCoR>
- [41] CCA-Forum, <http://www.cca-forum.org>
- [42] DOE Science Grid, <http://www-itg.lbl.gov/Grid>
- [43] pyGlobus, <http://www-itg.lbl.gov/Grid/projects/pyGlobus>
- [44] Babel, <http://www.llnl.gov/CASC/components/babel.html>
- [45] Alexandria, <http://www.llnl.gov/casc/components/alexandria.html>
- [46] PyClimate, <http://www.pyclimate.org>
- [47] ODE, <http://www.netlib.org/ode>
- [48] LAPACK, <http://www.netlib.org/lapack>
- [49] NWChem, <http://www.emsl.pnl.gov:2080/docs/nwchem>
- [50] InDEPS, <http://z.ca.sandia.gov/~indeps>
- [51] Globus, <http://www.globus.org>
- [52] CCTTSS, <http://www.cca-forum.org/ccttss>

- [53] APDEC, <http://davis.lbl.gov/APDEC>
- [54] TOPS, <http://www.math.odu.edu/~keyes/scidac>
- [55] TSTT, <http://www.tstt-scidac.org>
- [56] PERC, <http://perc.nersc.gov>
- [57] CITRIS, <http://www.citris.berkeley.edu>
- [58] ESI, <http://z.ca.sandia.gov/esi>
- [59] TeraScale, <http://www.terascale.com>
- [60] ESS/NASA, <http://ess.gsfc.nasa.gov>
- [61] CCS/UKY, <http://www.ccs.uky.edu>
- [62] IMA 2002, <http://www.ima.umn.edu/PI/2002summergrad.html>
- [63] SDSC, <http://www.sdsc.edu>
- [64] NPACI, <http://www.npaci.edu>
- [65] PSC, <http://www.psc.edu>
- [66] FETC/DOE, <http://www.fetc.doe.gov>
- [67] NSF, <http://www.nsf.gov>
- [68] CERFACS, <http://www.cerfacs.fr>
- [69] DCCIA, <http://www.dccia.ua.es/dccia/frontal.ing>
- [70] DSIC/UPV, <http://www-copa.dsic.upv.es>
- [71] CS/UTK, <http://icl.cs.utk.edu>
- [72] CITR/UTK, <http://icl.cs.utk.edu/citr>
- [73] ACTC/IBM, <http://www.research.ibm.com/actc>
- [74] PITAC, <http://www.ccic.gov/ac/report>
- [75] NUGEX, <http://hpcf.nersc.gov/about/NUG>
- [76] SLEPc, <http://acts.nersc.gov/workshop/slides/roman.pdf>
- [77] HPC++, <http://www.extreme.indiana.edu/hpc++>
- [78] Expression Templates, <http://osl.iu.edu/~tveldhui/papers/Expression-Templates/exptrtmpl.html>
- [79] OO Numerics, <http://oonumerics.org/oon>
- [80] NITRD, <http://www.itrd.gov/pubs/blue02/index.html>
- [81] DADS, <http://www.nist.gov/dads>
- [82] EOT-PACI, <http://www.eot.org>