

A Workshop on the ACTS Collection

Robust and High Performance Tools for Scientific Computing

Osni Marques

Lawrence Berkeley National Laboratory (LBNL)

oamarques@lbl.gov

How can ACTS work for you?

- Computational problems at hand.
- A sample of scientific applications.
- Why do we need all these tools?
- Agenda for the Workshop.

What needs to be computed?

ScaLAPACK

Aztec/Trilinos

SuperLU

$$Ax = b$$

$$Az = \lambda z$$

$$A = U\Sigma V^T$$

$$\min \left\{ \frac{1}{2} \|r(x)\|^2 : x_l \leq x \leq x_u \right\}$$

PDEs

ODEs

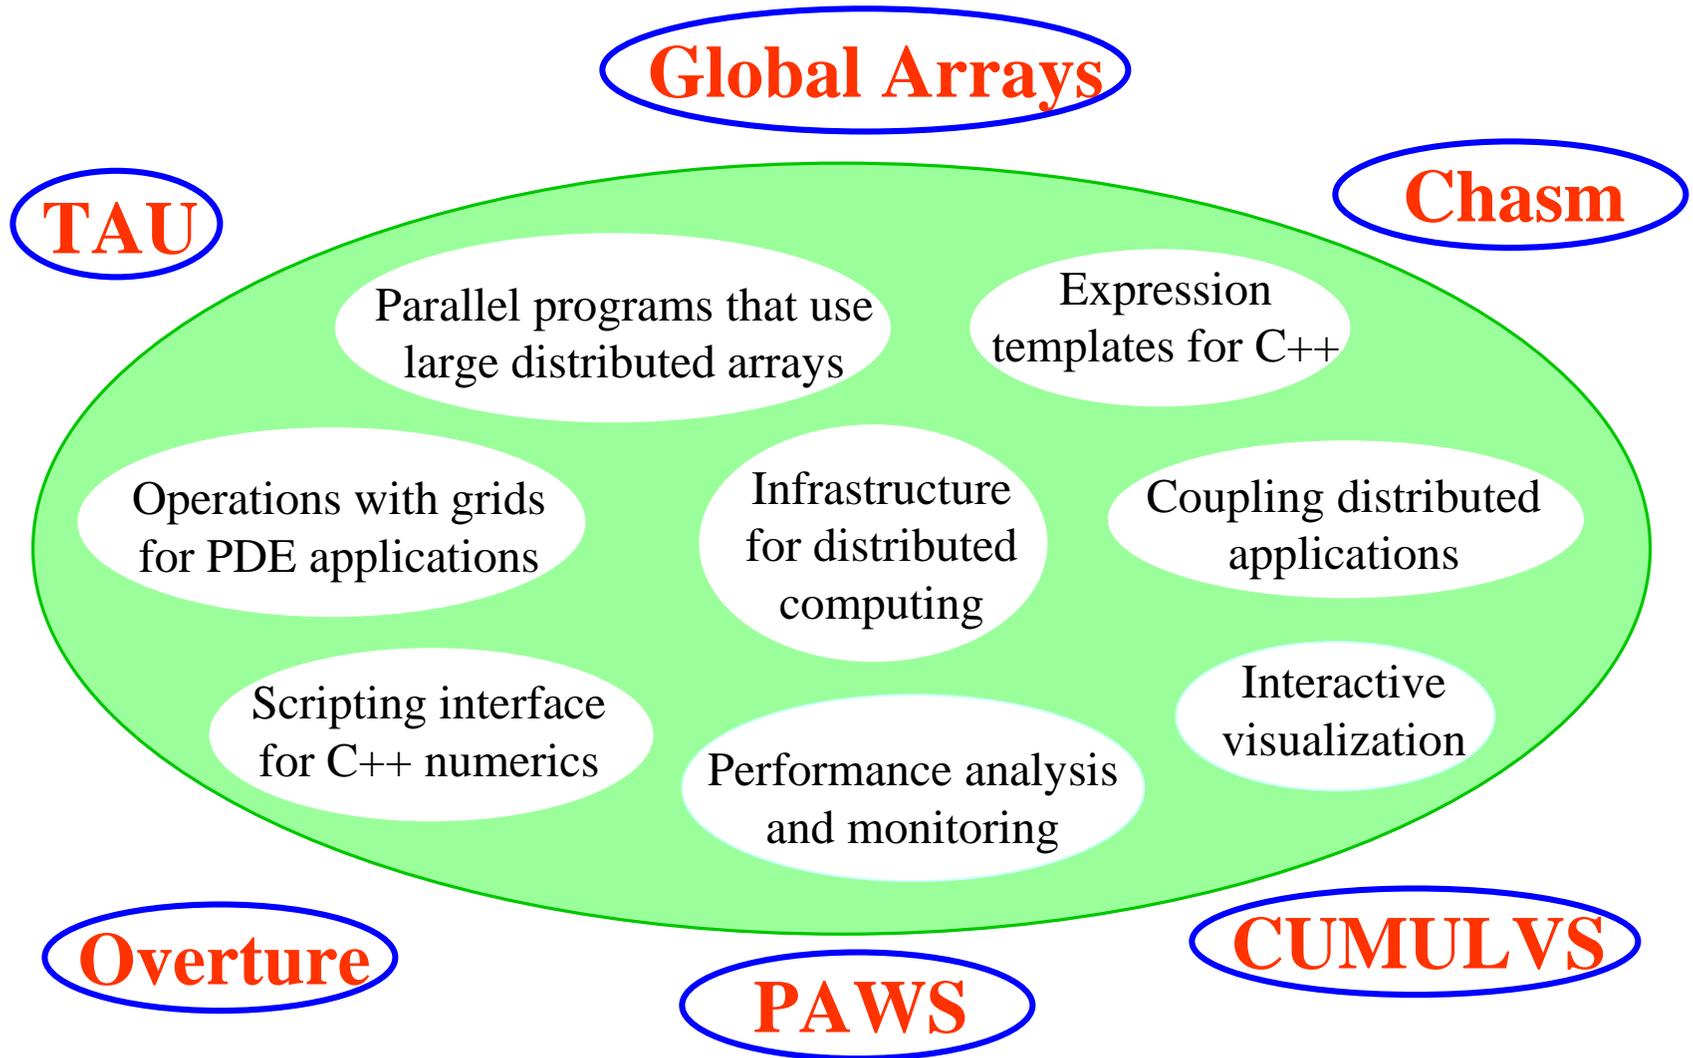
PETSc

Hypre

TAO

SUNDIALS

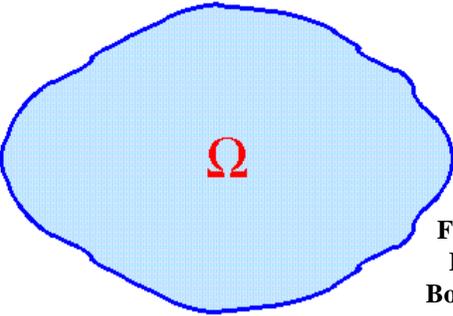
What codes are being developed?



Use of ACTS Tools

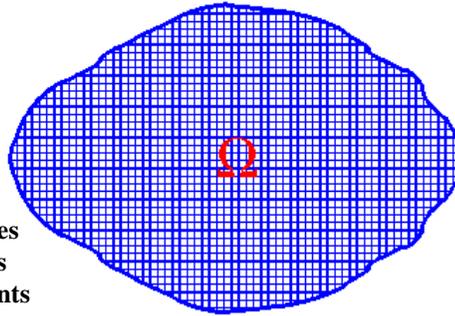
continuous problem

discrete problem



grid

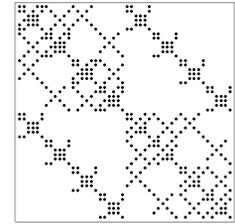
Finite Differences
Finite Elements
Boundary Elements
Fourier
⋮



$$m \frac{d^2 x}{dt^2} + c \frac{dx}{dt} + kx = f(t)$$

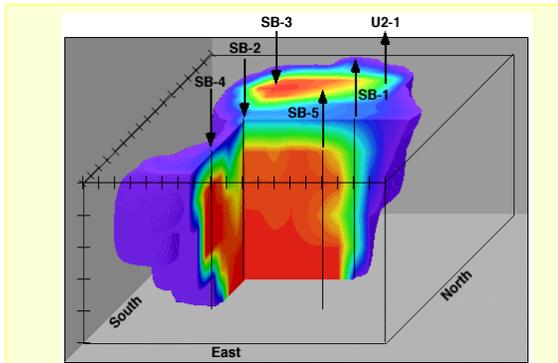
$$\frac{\partial v}{\partial t} = -v \frac{\partial v}{\partial x} + \Gamma \frac{\partial^2 v}{\partial x^2}$$

$$\vdots$$

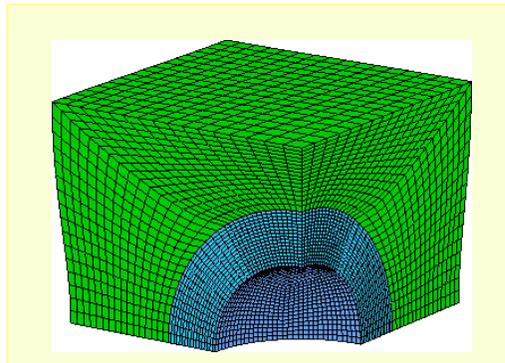


= A

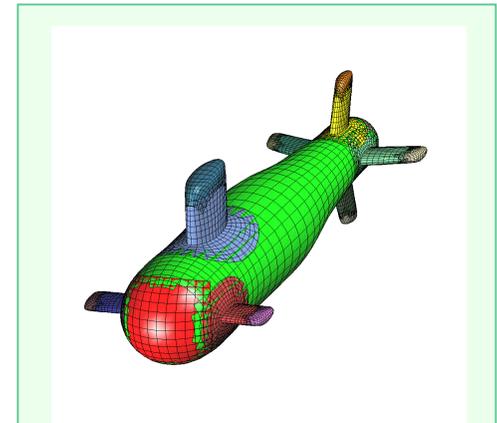
(sparse matrix)



Multiphase flow using *PETSc*, 4 million cell blocks, 32 million DOF, over 10.6 Gflops on an IBM SP (128 nodes), entire simulation runs in less than 30 minutes (Pope, Gropp, Morgan, Seperhrnoori, Smith and Wheeler).



Model of a "hard" sphere included in a "soft" material, 26 million d.o.f. (Adams and Demmel, Prometheus and *PETSc*, unstructured meshes in solid mechanics).

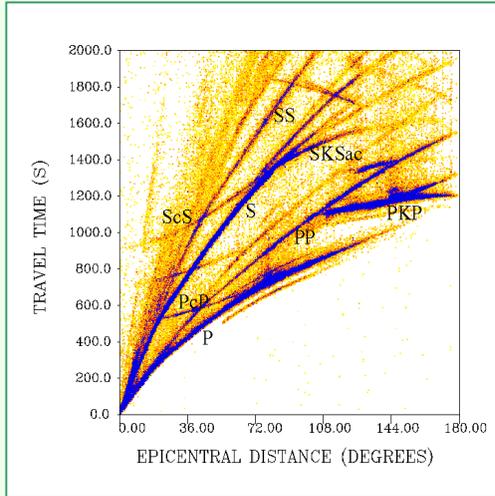


3D overlapping grid for a submarine produced with *Overture's* module *ogen*.

Use of ACTS Tools

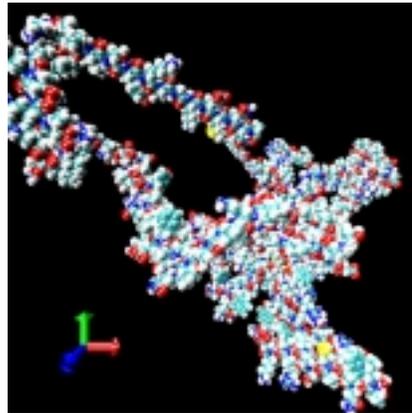
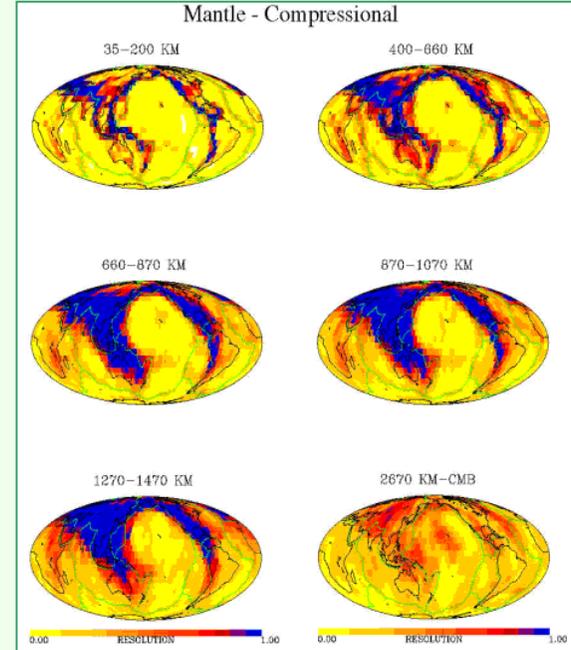
Data: travel times of sound waves generated by earthquakes used to infer structure in the entire Earth (crust, mantle and core).

Goal: model for the internal structure of the Earth



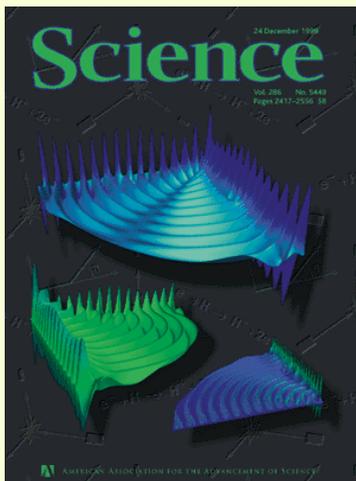
more than 1×10^6 data points and 2×10^5 parameters

ScaLAPACK used for the solution of eigenvalue problems.

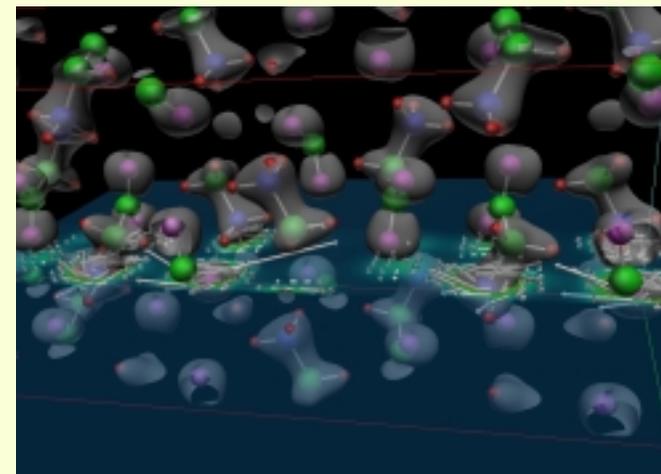


OPT++ is used in protein energy minimization problems (shown here is protein T162 from CASP5)

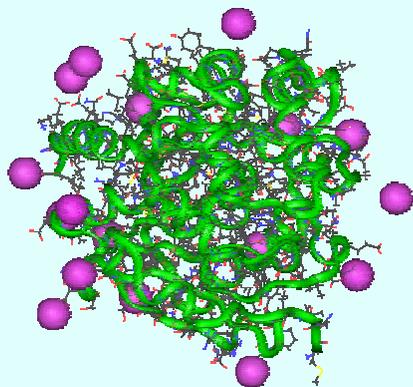
Use of ACTS Tools



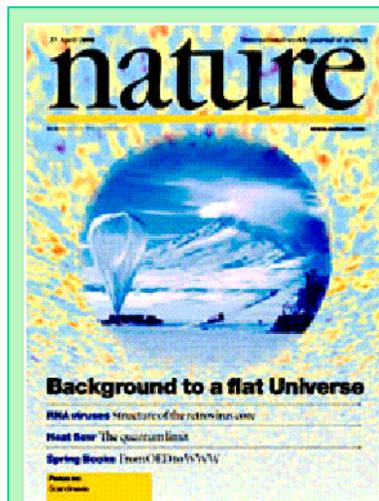
Collaborators at the LBNL, LLNL and the UC Davis obtained a complete solution of the ionization of a hydrogen atom by collision with an electron, the simplest nontrivial example of the long-standing unsolved problem of scattering in a quantum system of three charge particles. The images show representative wave functions of the problem (Rescigno, Baertschy, Isaacs and McCurdy, Dec. 24, 1999). *SuperLU* was used to build preconditioners for complex, unsymmetric and ill-conditioned sparse linear systems of order up to 2 million.



Induced current (white arrows) and charge density (colored plane and gray surface) in crystallized glycine due to an external field (Louie, Yoon, Pfrommer and Canning), eigenvalue problems solved with *ScaLAPACK*.



Molecular dynamics simulation of a haloalkane dehalogenase enzyme performed by a molecular dynamics code developed by T. Straatsma at PNNL, and using *Global Arrays*.

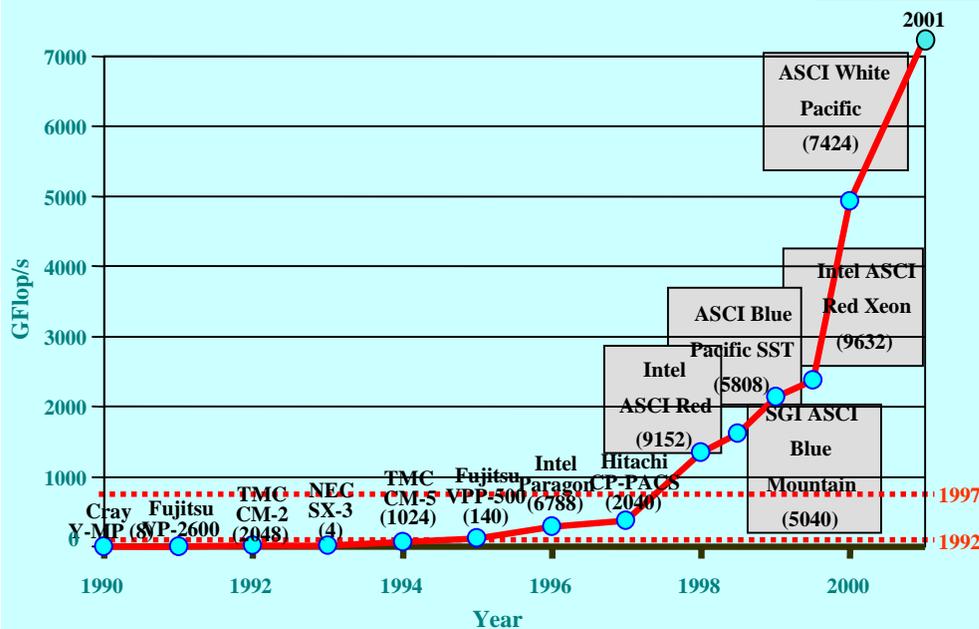


The international BOOMERanG collaboration announced results of the most detailed measurement of the cosmic microwave background radiation (CMB), which strongly indicated that the universe is flat (Apr. 27, 2000). Likelihood methods implemented in the MADCAP software package, using routines from *ScaLAPACK*, were used to examine the large dataset generated by BOOMERanG.

Why do we need these tools?

Top 5 Machines:
<http://www.top500.org>
 (June 2003)

Rank	Manufacturer	Computer	R_{max} [TF/s]	Installation Site	Country	Year	Area of Installation	# Proc
1	NEC	Earth-Simulator	35.86	Earth Simulator Center	Japan	2002	Research	5120
2	HP	ASCI Q, AlphaServer SC	13.88	Los Alamos National Laboratory	USA	2002	Research	8192
3	Linux NetworX	MCR Linux Cluster Xeon	7.63	Lawrence Livermore National Laboratory	USA	2002	Research	2304
4	IBM	ASCI White SP Power3	7.30	Lawrence Livermore National Laboratory	USA	2000	Research	8192
5	IBM	SP Power3	7.30	NERSC/LBNL	USA	2002	Research	6656



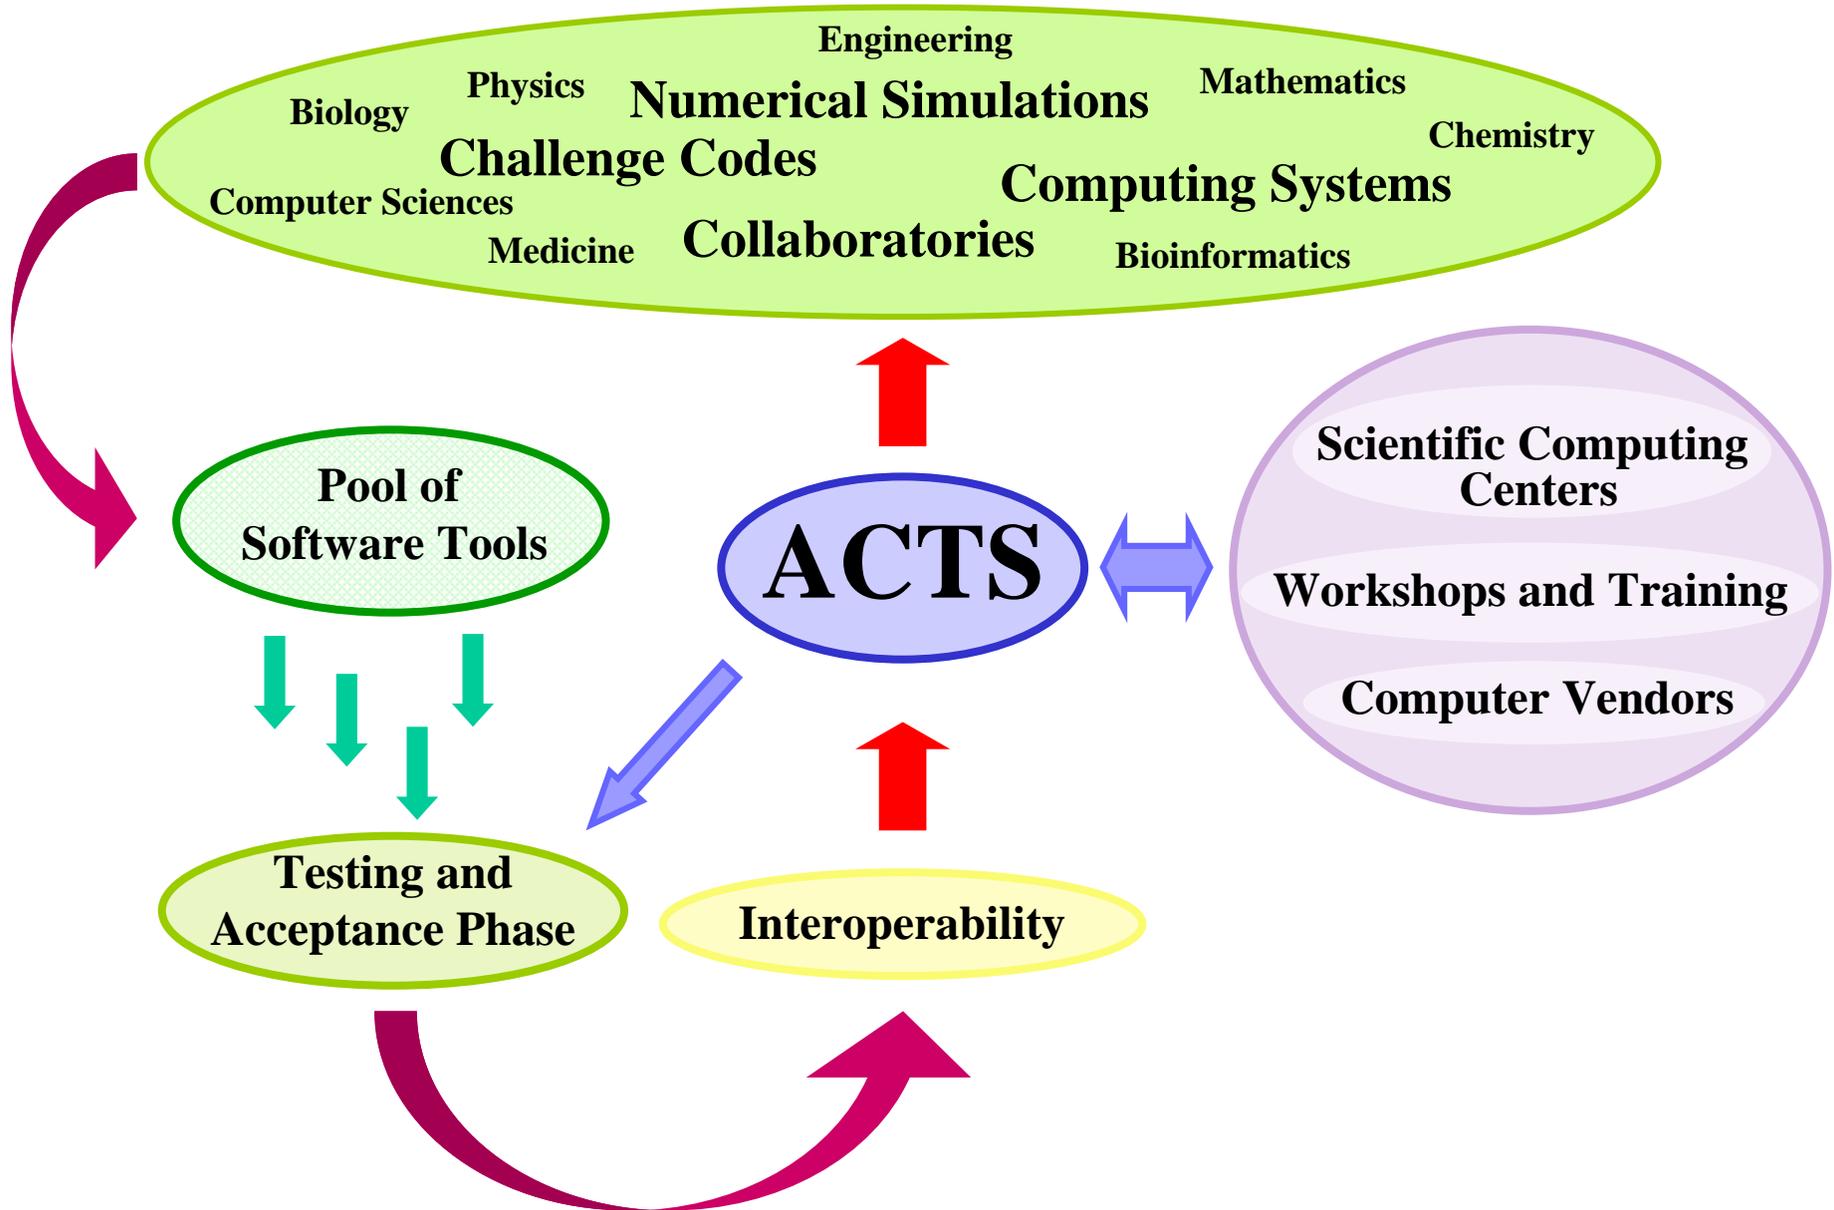
A computation that took 1 full year to complete in 1980 could be done in ~ 10 hours in 1992, in ~ 16 minutes in 1997 and in ~ 27 seconds in 2001!

- High Performance Tools
 - portable
 - library calls
 - robust algorithms
 - help code optimization
- More code development in less time!
- More simulation in less computer time!

Lessons Learned

- *There is still a gap between tool developers and application developers which leads to duplication of efforts.*
- *The tools currently included in the ACTS Collection should be seen as dynamical configurable toolkits and should be grouped into toolkits upon user/application demand.*
- *Users demand long-term support of the tools.*
- *Applications and users play an important role in making the tools mature.*
- *Tools evolve or are superseded by other tools.*
- *There is a demand for tool interoperability and more uniformity in the documentation and user interfaces.*
- *There is a need for an intelligent and dynamic catalog/repository of high performance tools.*

User Community



The DOE SciDAC Program

<http://www.science.doe.gov/scidac>

- Scientific Discovery through Advanced Computing:
 - Creation of a new generation of *Scientific Challenge Codes* for terascale computers to tackle critical scientific problems.
 - Creation of *Mathematical and Computing Systems Software* to enable the Scientific Challenge Codes to take full advantage of the computing power advances.
 - Creation of a *Collaboratory Software Infrastructure* to enhance remote collaboration and access to facilities and data.
- Program Areas:
 - Advanced Scientific Computing
 - Basic Energy Sciences
 - Biology and Environment
 - Fusion Energy Sciences
 - High-Energy and Nuclear Physics
- SciDAC is one important source of new tools and new users for the ACTS Collection.

Outreach (some past activities)

- ACTS Toolkit Workshop: *Solving Problems in Science and Engineering*. LBNL, October 9-13, 2001.
- ACTS Workshop: High Performance Libraries for Science and Engineering. LACSI Symposium, Santa Fe, October 18, 2001.
- Tutorial on High-Performance Math Libraries. NPACI All-Hands Meeting, San Diego, March 6, 2002.
- Workshop on High Performance Numerical Libraries for Science and Engineering, Seventh Copper Mountain Conference on Iterative Methods, March 25, 2002.
- Scaling to New Heights Workshop, Pittsburgh Supercomputing Center, May 20-21, 2002.
- ACTS Workshop at the IMA 2002 Summer Program, University of Kentucky Center for Computational Sciences, July 18-19, 2002.
- ACTS Workshop: Robust and High Performance Tools for Scientific Computing. LBNL, September 4-7, 2002.

Some Spin-offs and Collaborations

- Workshop 2000 Graduates:
 - Mark McKeown, from the University of Virginia (UVa), installed several of the ACTS tools on the UVa IBM SP2 and reported that several researchers at UVa use the tools.
 - Paulo Goldfeld, from the Courant Institute of Mathematical Sciences, started an active collaboration with the PETSc group, for the development of domain decomposition based preconditioners for mixed finite elements.
 - Joshua Breslau, from the Princeton Plasma Physics Laboratory, reported that with the expertise he acquired at the workshop he was able to easily upgrade some of the routines in the M3D code (multi-level 3D plasma physics) using PETSc.
- Workshop 2001 Graduates:
 - Mary Ann Leung, a DOE Computational Science Graduate Fellow, spent 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.
 - Ron Choy and Tong Wen, from MIT, are using ACTS tools in MATLAB*P (a parallel version of Matlab).
- Collaborations:
 - Dr. Jose Roman from the *Universidad Politecnica de Valencia*, Spain, visited LBNL in October 2001, and worked on SLEPc, which is a software library that extends the functionality of PETSc for the solution of large-scale sparse eigenvalue problems.
 - Prof. Eric Carlson from the Department of Chemical Engineering, University of Alabama, was granted a DOE EPSCoR proposal for the development of anisotropic diffusion equation solver using ACTS Tools.

In the works...

- Special issue of ACM Transactions on Mathematical Software (TOMS) containing a paper on the aims, scope, and successes of the ACTS Project, together with a collection of papers, each of which focusing on a specific tool.
- Special issue of International Journal on High Performance Computers for the “non numerical” tools.
- SC2003 (Phoenix):
 - Presentation at the LBNL booth
 - Poster
- We would like to hear from you and your computational problems!

Workshop Agenda

- Numerical Tools:
 - Hypre
 - OPT++
 - PETSc
 - ScaLAPACK
 - SuperLU
 - TAO
- Tools for code development:
 - Global Arrays
 - Overture
- Tools for code execution:
 - Chasm
 - CUMULVS
 - PAWS
 - TAU
- Grid:
 - GLOBUS
- SciDAC:
 - *Center for Terascale Optimal PDE Simulations (TOPS)*
 - *Terascale Simulation Tools and Technology Center (TSTT)*
 - *An Algorithmic and Software Framework for Applied Partial Differential Equations (APDEC)*
 - *Performance Evaluation Research Center (PERC)*
 - *Center for Component Technology for Terascale Simulation Software (CCTTSS)*
- Invited Talks:
 - *SCiDAC at NERSC*
 - *HPC Middleware*

Have a nice Workshop!