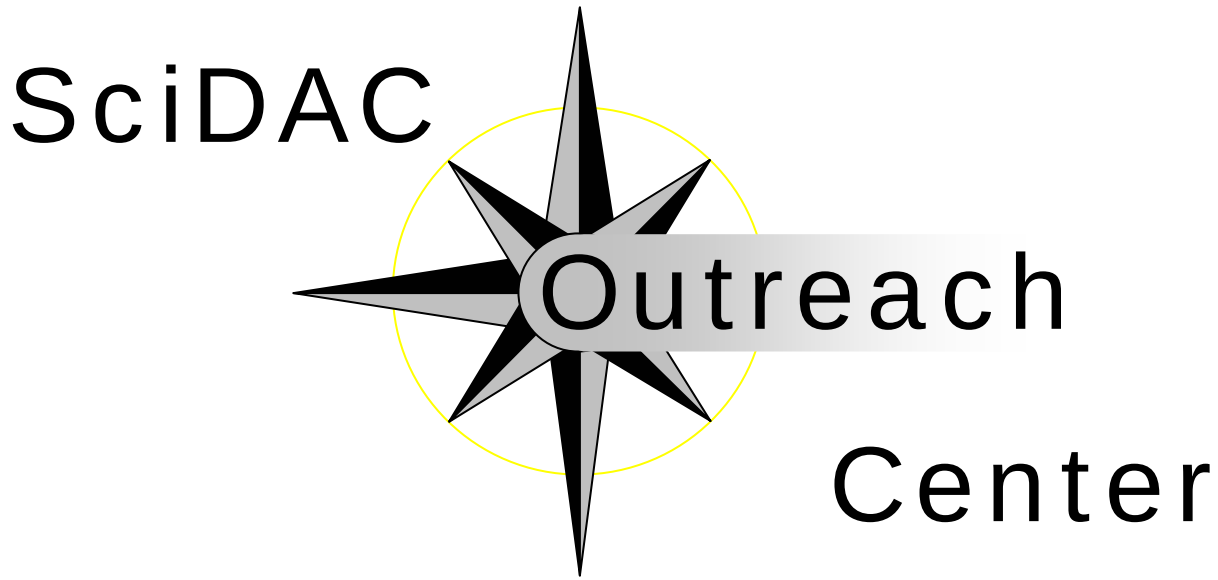

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY

Two topics for CSGF

- Trends in computing for the young researcher to consider
- Resources in SciDAC that can help

Computational Science :
The computer as microscope

"An important development in sciences is occurring at the intersection of computer science and the sciences that has the potential to have a profound impact on science. It is a leap from the application of computing … to the *integration of computer science concepts, tools, and theorems* into the very fabric of science." -*Science* 2020 Report, March 2006

Nature, March 23, 2006

2020
SCIENCE

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY

Computational Science Engineering

Three Examples

- **simulation replacing experiment that is too difficult**
- **simulation replacing experiment that is too dangerous**
- **analyzing massive amounts of data with new tools**

For this to happen we need (besides a brilliant research idea)

- A computer and access to it
- Someone who can operate the computer
- An algorithm implemented in software

Computational Science and Engineering (CSE)

- **CSE is a widely accepted label for an evolving field concerned with the science of and the engineering of systems and methodologies to solve computational problems arising throughout science and engineering**
- **CSE is characterized by**
  - **Multi - disciplinary**
  - **Multi - institutional**
  - **Requiring high-end resources**
  - **Large teams**
  - **Focus on community software**
- **CSE is not "just programming" (and not CS)**
- **Teraflop/s computing is necessary but not sufficient**

Reference: Petzold, L., *et al*., Graduate Education in CSE, *SIAM Rev*., 43(2001), 163-177

CSE is expanding the scope of science

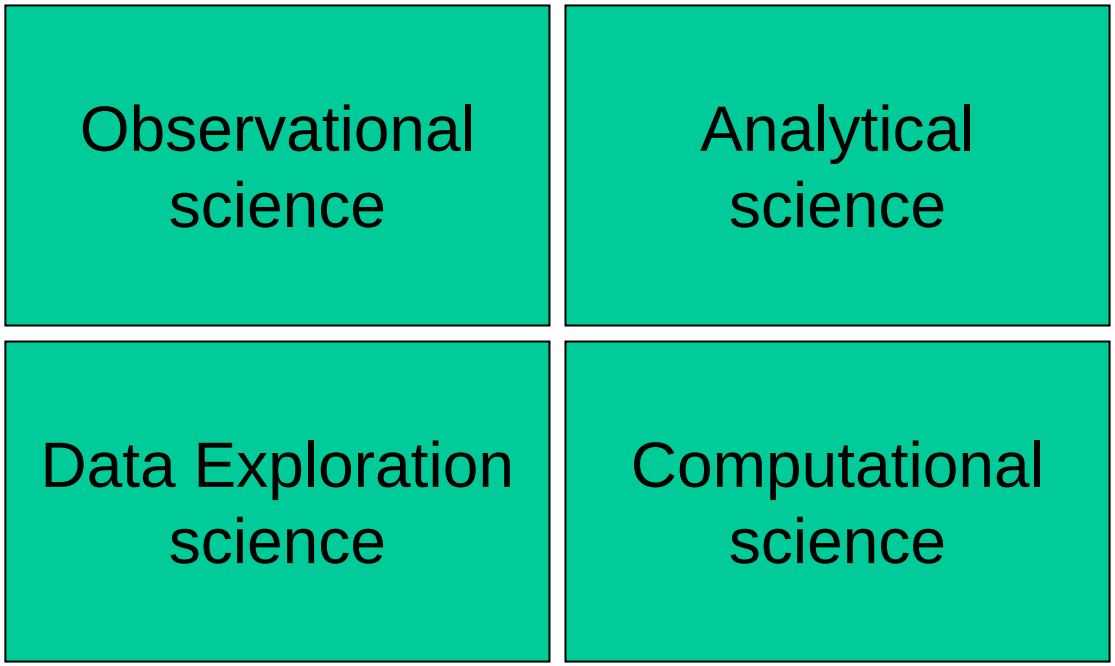

Computing means both Flops and Data

Technology Trends: Microprocessor Capability

Moore's Law

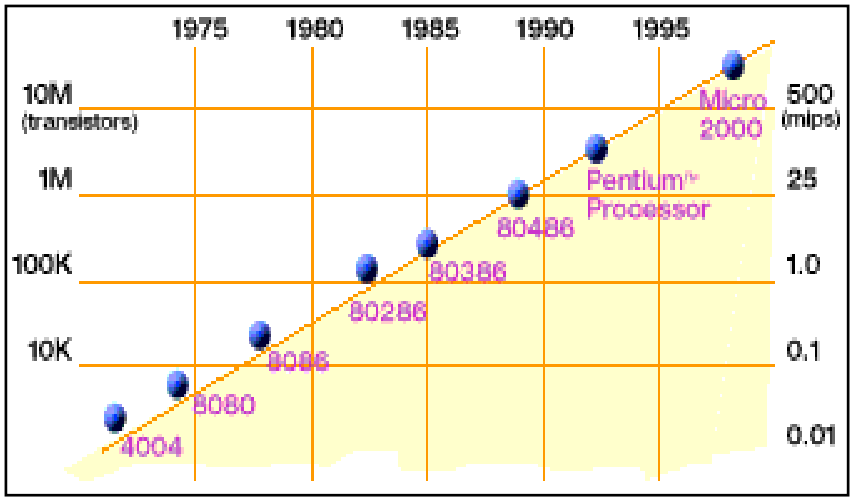

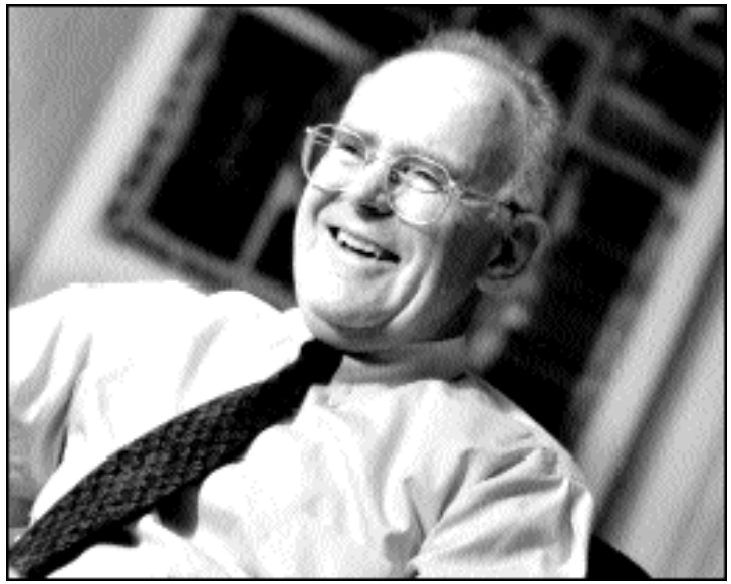

2X transistors/chip every 1.5 years
Called "Moore's Law"

Microprocessors have become smaller, denser, and more powerful.

Gordon Moore (co-founder of Intel) predicted in 1965 that the transistor density of semiconductor chips would double roughly every 18 months.

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY

Traditional Sources of Performance Improvement are Flat-Lining

- New Constraints
  - **15 years of *exponential* clock rate growth has ended**
- But Moore's Law continues!
  - **How do we use all of those transistors to keep performance increasing at historical rates?**
  - **Industry Response: #cores per chip doubles every 18 months *instead* of clock frequency!**
- No more automatic performance boost over time
- How will this constrain your research?

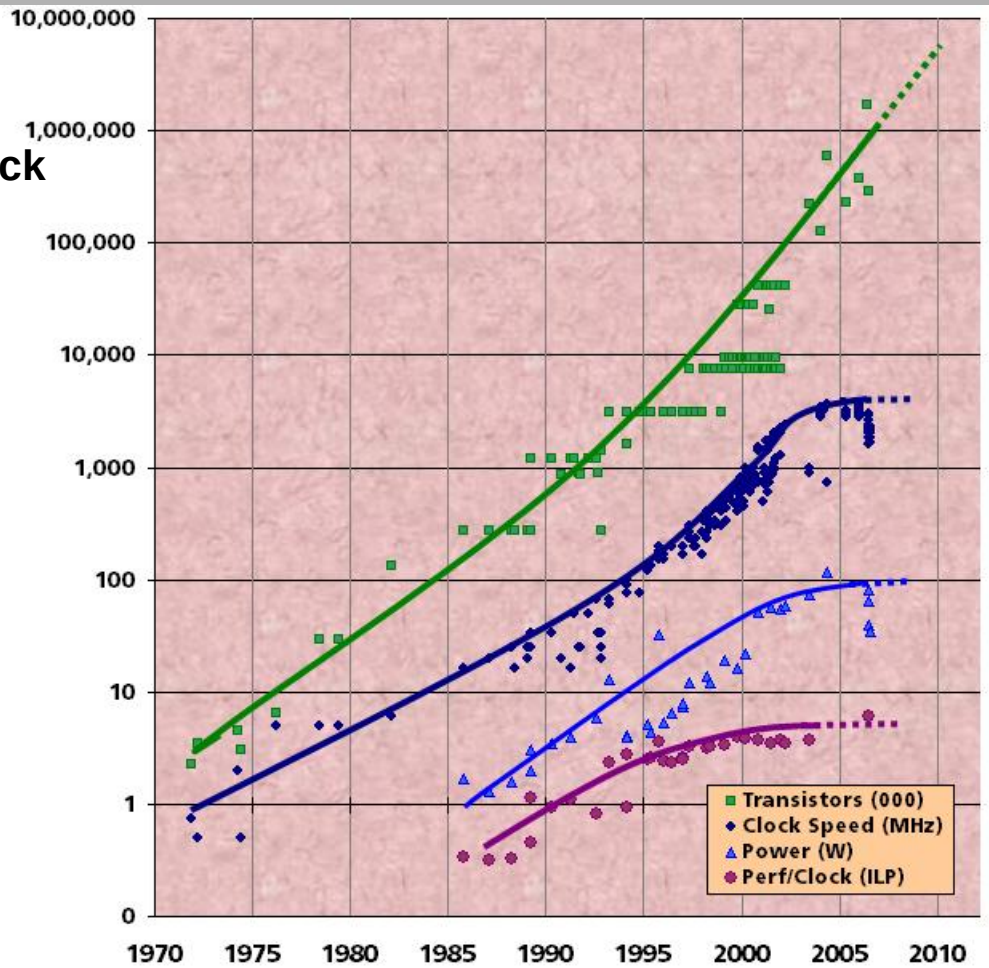

Hockeystick Parallelism

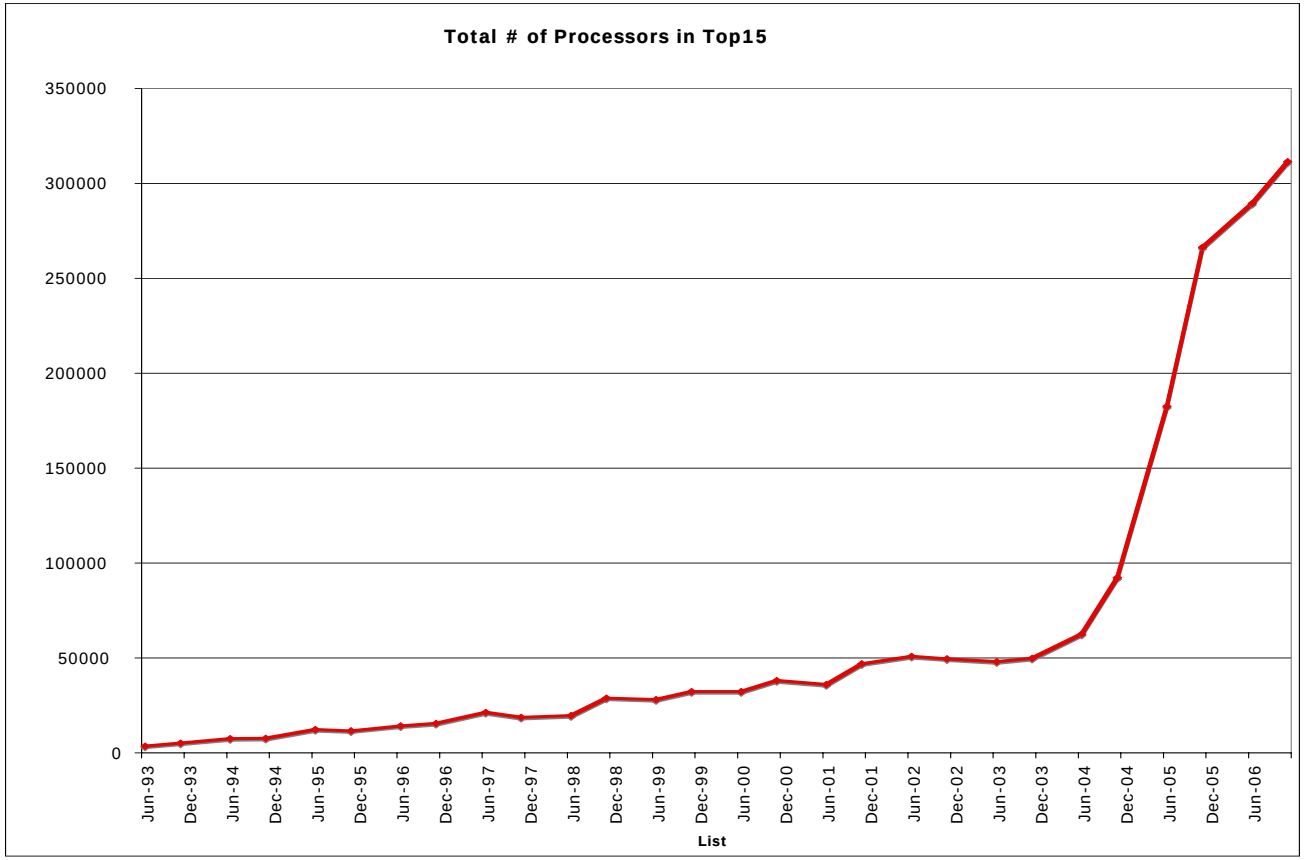

- Future performance gains will come from many-core parallelism
- Who will implement scientific software and algorithms that exploit that parallelism?

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY

Supercomputing Today & Tomorrow

- Microprocessors have made desktop computing in 2007 what supercomputing was in 1995.
- Massive Parallelism has changed the “high-end” completely.
- Most of today's standard supercomputing architecture are “hybrids”, clusters built out of commodity microprocessors and custom interconnects.
- The microprocessor revolution will continue with little attenuation for at least another 10 years
- Like it or not, the future will be massively parallel, based on multicore

Data Trends

- Simulation sizes are increasing
  - Time spent managing data can bottleneck research
  - Increased attention on the number of copies of data, how data is moved, and archived
- Experimental sources of data are expanding
  - # sources, data rates

Software Trends

- More choices in languages, libraries, and compilers
- Software is becoming more layered
- All of these imply a need for organization, attention to program design, and software engineering

These trends imply changing research skill sets

- Old
  - How to compile codes w/ scalar optimization
  - Keeping workstation backed up
- New
  - Scaling to 10K-100K cores
  - Data movement, streaming analysis, archiving
  - Managing code sources, SW engineering
- Not adapting means your research possibilities are confined to a (relatively) smaller and smaller box.

Workstations and HPC

- Use the right tool for the job
- Proficiency with a range of tools expands your research options. Learn to be nimble in moving from serial to parallel.
- HPC allows you to
  - Think big, outside the (workstation) box
  - Be lazy, run a 112GB perl script
  - Focus on research, not on being a system/cluster admin

What is SciDAC?

- Scientific Discovery through Advanced Computing
- Where applied mathematicians and computer scientists find jointly funded collaboration
- Home to many innovative HPC software and algorithm resources
  - AMR, performance tools, data management tools
  - scalable solvers, visualization, advanced meshing
  - and lots more : www.scidac.gov

SciDAC
First Federal Program to Implement CSE

SciDAC (Scientific Discovery through Advanced Computing) program created in 2001

- **About $50M annual funding**
- **Berkeley (LBNL+UCB) largest recipient of SciDAC funding**

Global Climate

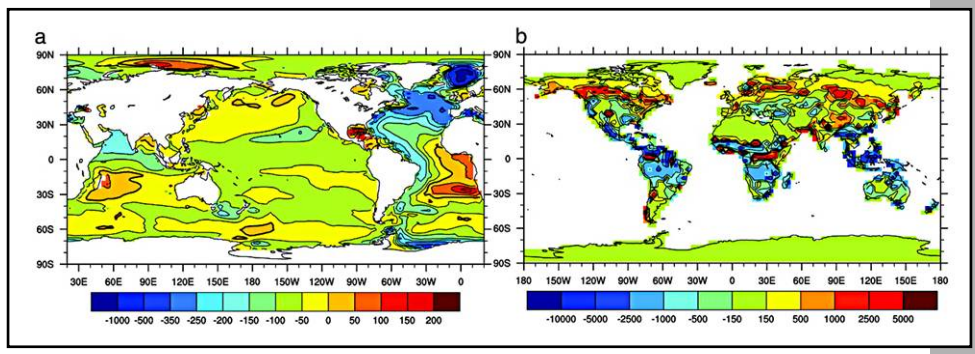

Biology

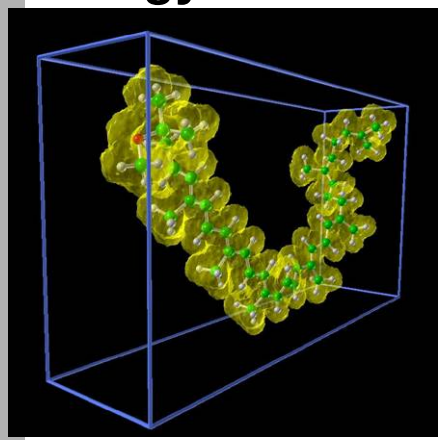

Nanoscience

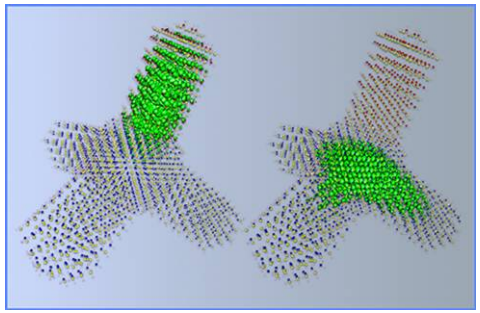

Combustion

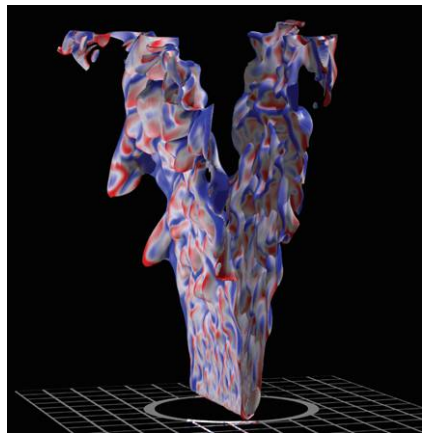

Astrophysics

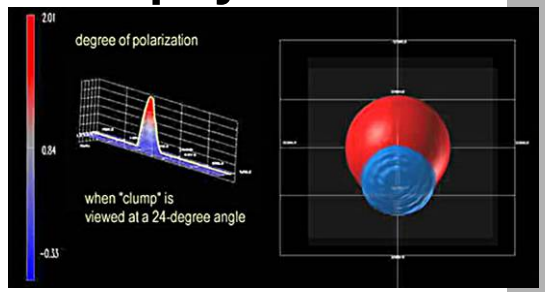

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY

SciDAC Outreach Center Mission

- Provide services that make SciDAC supported technologies more accessible both within and outside the SciDAC community
- Field inquiries about SciDAC that range from general information to technical specifics
- Assist in deployment and bridge gaps between SciDAC stakeholders
- Provide a central orientation for all things SciDAC Get interested parties to the right resources
- Foster awareness and education about HPC

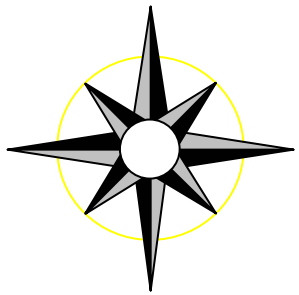

Outreach can mean many things

We are currently focused on two approaches

- Innovative web and software services
  - Tools which make SciDAC researchers more effective at delivering their technologies
  - Information services which provide an easy interface to SciDAC for all involved
- In person outreach
  - Workshops, trainings, and event coordination
  - Getting the right people to the right audiences

Collaborative e-Services

- Author Documents
- Inform Collaborators
- Inform the public
- Develop software
- Test software
- Package software
- Distribute software

How?
http://outreach.scidac.gov/

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY

Gettting an HPC allocation

- Not as hard as you might think
  - If you have an abstract of your research goals applying will take you 30 min or so
- A small allocation is a stepping stone toward a large allocation when you need it. It helps you build a computing relationship with DOE and project reviewers.
- NERSC
  - https://nim.nersc.gov/newpi.php
- ANL
  - https://accounts.alcf.anl.gov/accounts/projects/intrepid.htm
- ORNL
  - http://www.nccs.gov/user-support/access/project-request

Is HPC hard?

- As with most areas of science “it depends”
- Crafting your own parallel app from scratch may be a challenge, but maybe you don’t need to
- There are many existing applications and libraries that may solve the problem you’re tackling

Summary and Contact

We reinforce the SciDAC mission by bringing its information and technologies to new venues in DOE, academia, and industry through electronic and in person outreach

Email : help@outreach.scidac.gov

Phone : 1-866-470-5547

SciDAC
Scientific Discovery through Advanced Computing

Office of Science
U.S. DEPARTMENT OF ENERGY