

UT-BATTELLE



AMES LABORATORY



ARGONNE NATIONAL LABORATORY



Los Alamos NATIONAL LABORATORY



Lawrence Livermore National Laboratory



NCAR

Pacific Northwest National Laboratory



PPPL



Sandia National Laboratories



Leadership Computing Facility at Oak Ridge National Laboratory

Studham RS, Kuehn J, White JB, Fahey M, Carter S, Bland A, Nichols JA, Zacharia T


CUG
May 17th 2005

THE CENTER FOR
COMPUTATIONAL SCIENCES

OAK RIDGE NATIONAL LABORATORY
U. S. DEPARTMENT OF ENERGY

Outline

- What is “Leadership” computing
- What is the “Leadership Computing Facility at ORNL”
- Overview of the LCF computing architectures and long term plans
- Benchmarks of LCF computing systems
- Example science results from LCF systems.
- LCF proposal process and challenges.




CCS first Cray
1985 -1992
1 Processor
8 MB memory
180 MFLOPS

The user paradigm must change to do Big Science (AKA "Supercomputing")

DOE-SC's primary computing resource is a 10TF (peak) system that supports ~363 projects

The average a project gets:

$$\frac{10\text{TFyears}}{363 \text{ Projects}} = 27 \text{ GFyears/Project}$$



What is “Leadership Computing”?

Leadership computing systems will provide computational capability that is at least **100 times greater** than what is generally available for advanced scientific and engineering problems.


| | Generally Available TFy/Project | Leadership Computing TFy/Project | TF at NLCF | Number of Leadership Computing Projects |
|------|---------------------------------|----------------------------------|------------|---|
| 2005 | 0.027 | 2.7 | 6.8 | 2 |
| 2006 | 0.055 | 5.5 | 44 | 8 |
| 2007 | 0.110 | 11.0 | 118 | 11 |
| 2008 | 0.220 | 22.0 | 368 | 17 |
| 2009 | 0.441 | 44.1 | 1000 | 23 |

Leadership computing is a White House priority

“The goal of such systems [leadership systems] is to provide computational capability that is at least 100 times greater than what is currently available.”

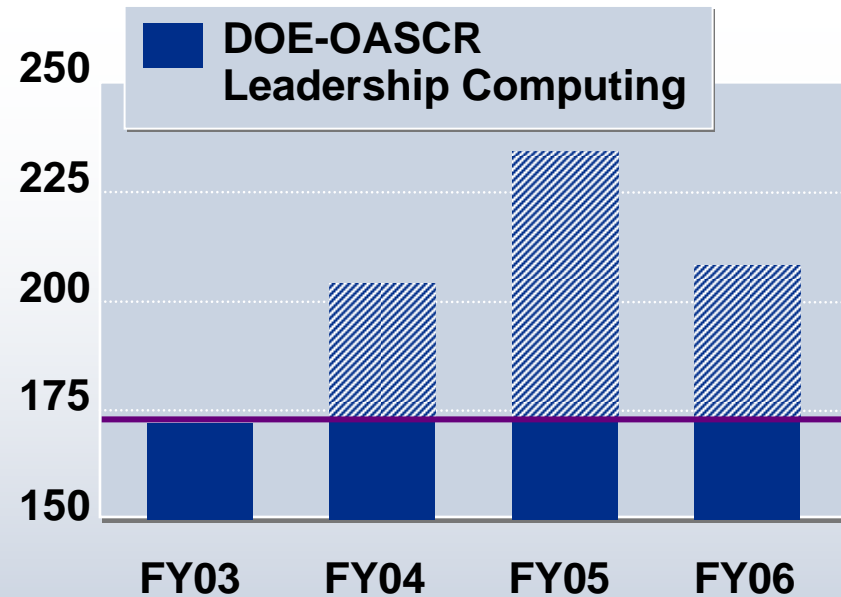
“...Leadership Systems are expensive, typically costing in excess of \$100 million per year....”

– Page 29, Federal Plan for High-end Computing



Leadership computing is a congressional priority

- Appropriated \$30M in FY04 for leadership computing
- Additional \$30M appropriated in FY05
- Public Law 108-423 Department of Energy High-End Computing Revitalization Act of 2004



Leadership computing is a State of Tennessee priority

- **\$9M State of Tennessee Investment in the Joint Institute for Computational Sciences**
- **\$10M for National Academy Level Joint Faculty**
- **\$12M for high speed networks for research and education**
- **\$1M/year for Computational Science Initiative for graduate student training and outreach**



“I have recommended funds ...to attract more nationally-recognized faculty members (jointly with ORNL).... There is an opportunity today.... to rapidly become world class in some areas like supercomputers, materials science, and nanotechnology.

....Our pioneer ancestors wouldn't have known what supercomputers were, but I believe they would have understood our aspirations perfectly.”



- Gov. Bredesen, State of the State Speech, January 31, 2005

Leadership computing is the highest domestic priority for Office of Science

- **Ray Orbach has articulated his philosophy for the SC laboratories**
 - Each lab will have world-class capabilities in one or more areas of importance to Office of Science
 - ORNL: SNS and NLCF will underpin world-class programs in materials, energy, and life sciences
- **20-year facilities plan being used to set priorities among projects**

“I am committed to the concept of a Leadership Class Computing facility at Oak Ridge National Laboratory. The facility will be used to meet the missions of the Department and those of other agencies. I can assure you that I understand the important role supercomputing plays in scientific discovery.”

Secretary Bodman



Office of Science notice to SC Laboratories

Leadership-Class Computing Capability for Science



“The focus of the proposed effort should be on capability computing in support of high-end science – rather than on enhanced computing capacity for general science users”

“The proposed effort must be a user facility providing leadership class computing capability to scientists and engineers nationwide independent of their institutional affiliation or source of funding.”

“Proposals must include specific information on architecture or architectures to be provided over the life of this project as well as a list of targeted applications.”

Our plan of action to deliver leadership computing for DOE

- **Rapidly field most powerful open capability computing resource for scientific community**
 - Providing clear upgrade path to at least 100 teraflop/s (TF) by 2006 and 250 TF by 2007/2008
- **Deliver outstanding access and service to research community**
 - Utilizing most powerful networking capability extant coupled with secure and highly cost-effective operation by proven team
- **Deliver much higher sustained performance for major scientific applications than currently achievable**
 - Developing next generation models and tools
 - Engaging computer vendors on hardware needs for scientific applications
- **Engage research communities in climate, fusion, biology, materials, chemistry, and other areas critical to DOE-SC and other federal agencies**
 - Enabling high likelihood of breakthroughs on key problems
- **Conduct in-depth exploration of most promising technologies for next-generation leadership-class computers**
 - Providing pathways to petaflop/s (PF) computing within decade

National Leadership Computing Facility

May 12, 2004



Oak Ridge gets super win

"Feds spark high-end computing resurgence"

By Scott Chamberlain

With great high-end computing power, Oak Ridge National Laboratory (ORNL) has won the bid to build a supercomputer that will be the most powerful in the world. The \$250 million project will be the largest ever for the U.S. government, and it will be the largest ever for the U.S. government.

U.S. Moves to Build Top Supercomputer

Wednesday, May 12, 2004

By Scott Chamberlain

The U.S. government is set to build the world's most powerful supercomputer, a project that will cost \$250 million and take three years to complete. The computer will be built at Oak Ridge National Laboratory in Tennessee.

media News

CRAY X1

U.S. Energy Dept. Awards \$25M Grant for Supercomputer Build

ComputerWeekly

U.S. Energy Dept. Awards \$25M Grant for Supercomputer Build

The U.S. Department of Energy has awarded a \$25 million grant to Oak Ridge National Laboratory (ORNL) to build the world's most powerful supercomputer.

COMPUTERWORLD

Energy Department Grant Spurs Work on Supercomputer

The \$25M grant went to the Oak Ridge National Laboratory.

BusinessWeek

Where No Computer Has Gone Before

the secrets of fusion energy, subatomic behavior, and being attacked by supercomputers capable of superintelligence

U.S. to Build Civilian Supercomputer

By Scott Chamberlain

The U.S. government is set to build the world's most powerful supercomputer, a project that will cost \$250 million and take three years to complete. The computer will be built at Oak Ridge National Laboratory in Tennessee.

The Atlanta Journal-Constitution

Computer to be more than a bit faster

Supercomputer-generated "power wall" at Oak Ridge, Tenn., National Laboratory illuminates climate change

Uncle Sam Wants World's Fastest Computer

Want To Surpass Japanese Supercomputer

The U.S. government is set to build the world's most powerful supercomputer, a project that will cost \$250 million and take three years to complete. The computer will be built at Oak Ridge National Laboratory in Tennessee.

Tennessee is a Rising Star in Supercomputing

ORNL Nabs Supercomputer Deal

LAB CHIEF: 'Oak Ridge will help lay the foundation for our country's next generation of scientific discovery.'

WASHINGTON TECHNOLOGY

DOE Orders World's Fastest Supercomputer

By Scott Chamberlain

The U.S. government is set to build the world's most powerful supercomputer, a project that will cost \$250 million and take three years to complete. The computer will be built at Oak Ridge National Laboratory in Tennessee.

San Francisco Chronicle

U.S. Government Moves to Build World's Fastest Civilian Computer

By Scott Chamberlain

The U.S. government is set to build the world's most powerful supercomputer, a project that will cost \$250 million and take three years to complete. The computer will be built at Oak Ridge National Laboratory in Tennessee.

矽谷科技

全球最快速民用科研電腦

能源部將在田納西州的美國橡樹嶺國家實驗室 (Oak Ridge National Laboratory) 興建全球最快速的民用科研電腦，該電腦將安裝在四萬平方呎的運算科學中心內。該系統每小時運算能力高達10^16次。這台電腦將於2006年開始運作。(美聯社)

GLOBEANDMAIL.COM

U.S. Aims to Regain Computer Speed Lead

WASHINGTON — The United States is looking to regain the lead in the competition for the world's fastest supercomputer.

DOE Leadership-Class Computing Capability for Science will be Developed at Oak Ridge National Laboratory

Office of Science

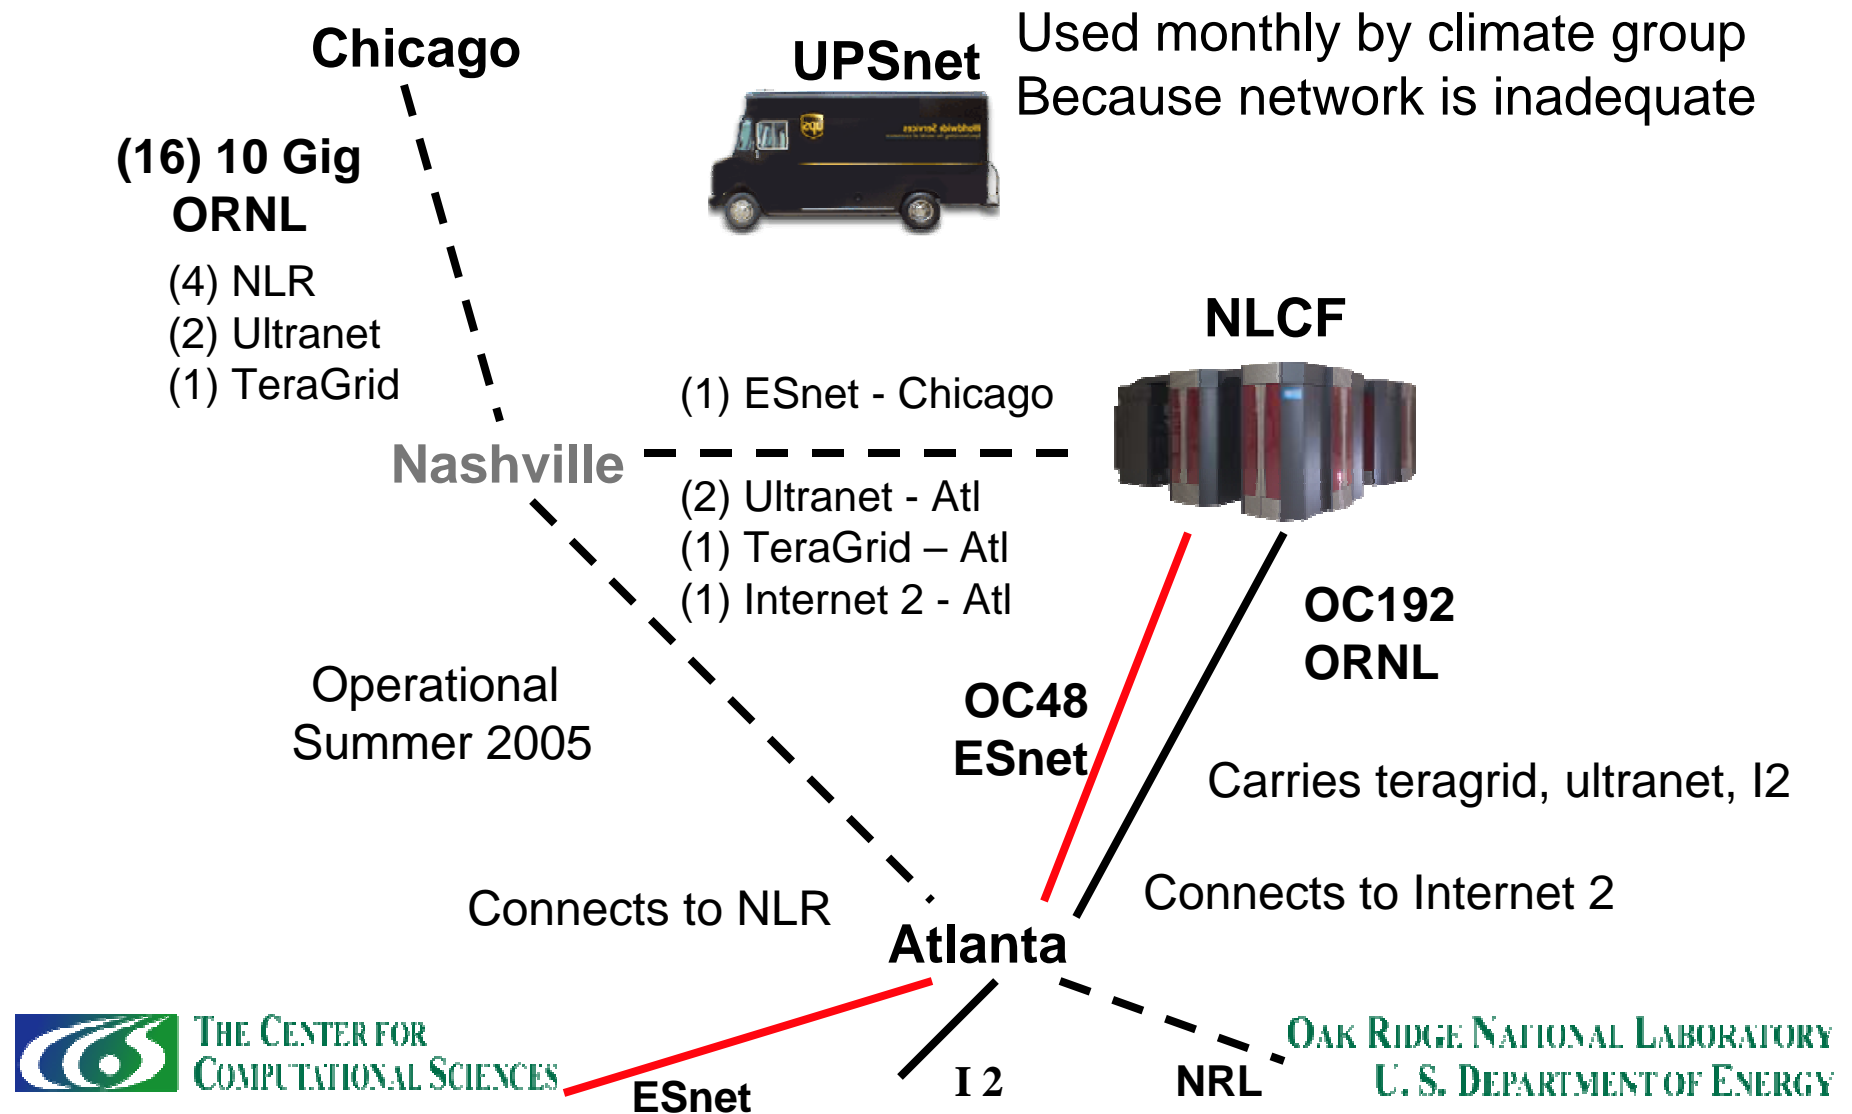
U.S. DEPARTMENT OF ENERGY

Leadership-class Computing Requires Leadership-class Facilities

- \$72M private sector investment in support of leadership computing
- Space and power:
 - **40,000 ft² computer center with 36-in. raised floor, 18 ft. deck-to-deck**
 - **8 MW of power (expandable) @ 5c/kWhr**
- High-ceiling area for visualization lab (Cave, Powerwall, Access Grid, etc.)
- Separate lab areas for computer science and network research



Leadership-class Computing Requires Leadership-class Connectivity




NLCF hardware roadmap

In 2005,

- Deploy 18.5TF Cray X1E and 25.1TF Cray XT3 systems
- Cray forms and supports “Supercomputing *Center of Excellence*”
- Develop/deploy complementary software environment
- Full operational support of NLCF as a capability computing center
- Deliver computationally intensive projects of large scale and high scientific impact through competitive peer review process

In 2006, deploy 100TF Cray XT3

In 2007-8, deploy 250+TF system



Phoenix – The CCS Cray X1/X1e

- Largest Cray X1 in the world. Upgrading to 18TF X1E this summer.

- Highest bandwidth communication with main memory



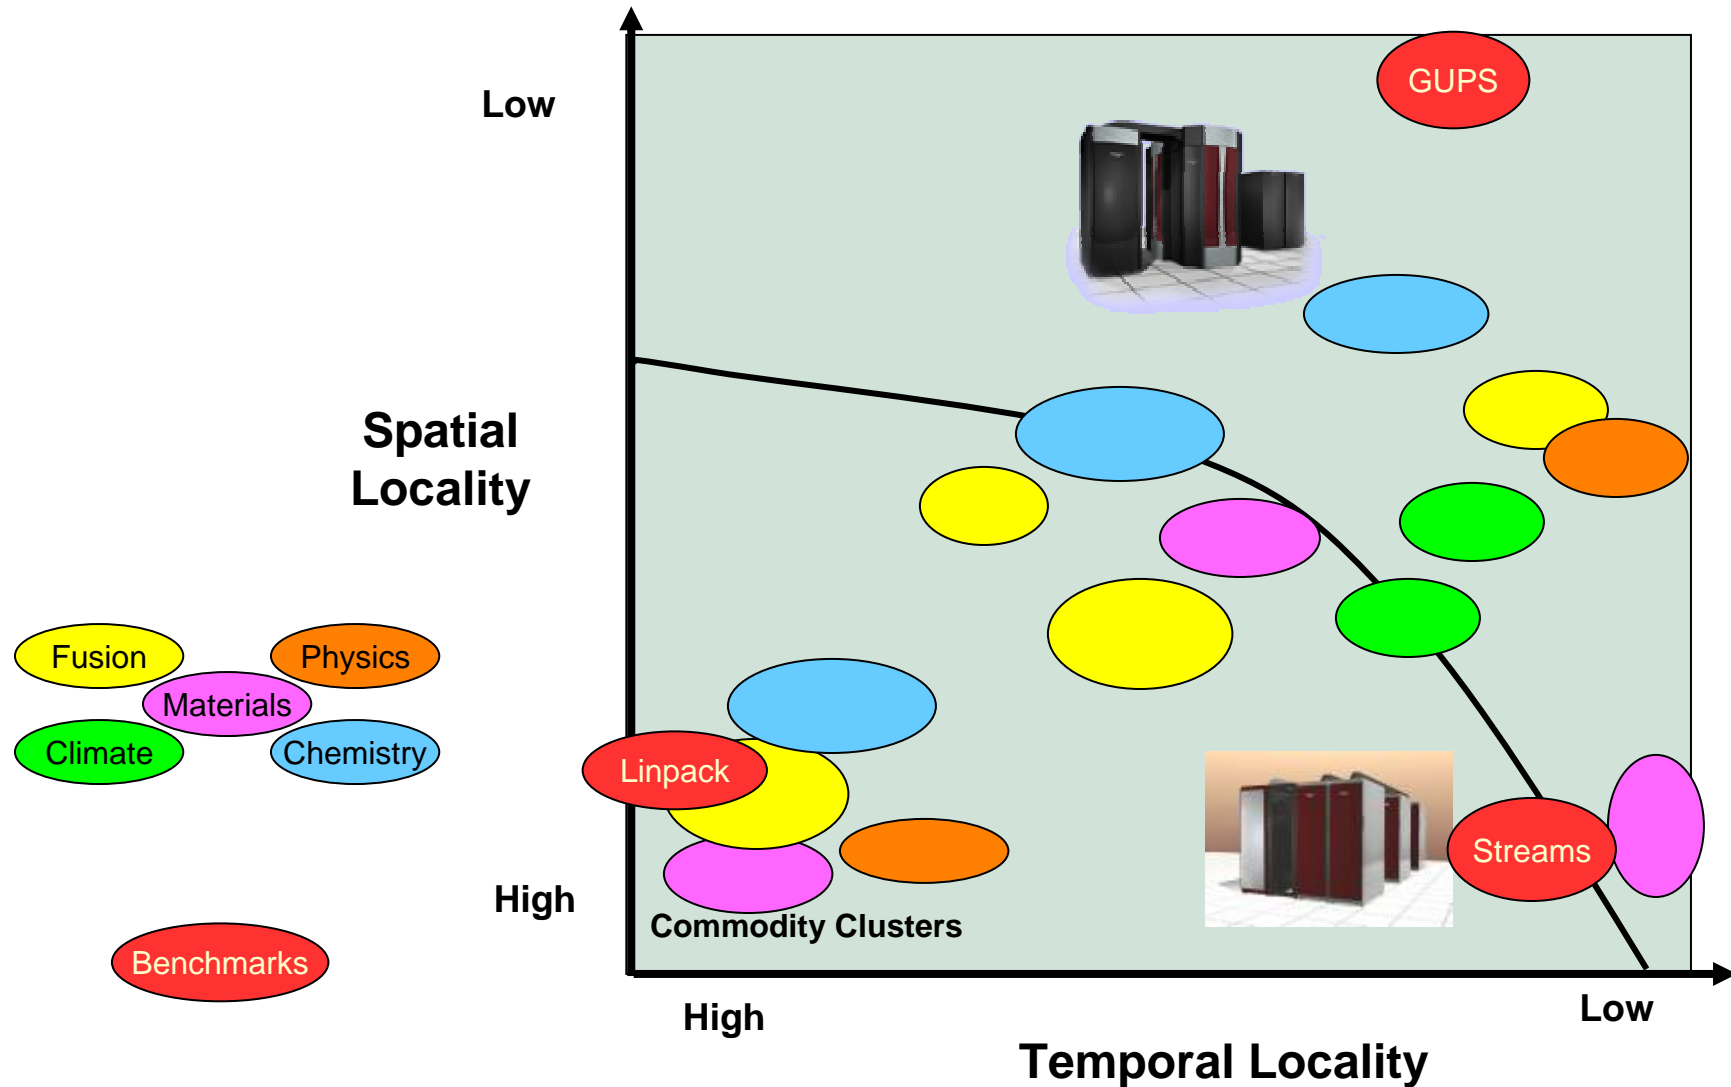
- Highly scalable hardware and software
- High sustained performance on real applications

Jaguar – The CCS XT3 system

| Status | Performance | Processors | Memory | I/O Bandwidth |
|------------|-------------|------------|---------|---------------|
| Summer 05 | 25 TF | 5,304 | 10.5 TB | 15 GB/s |
| Fall 05* | 50 TF | 11,374 | 23 TB | 30 GB/s |
| Spring 06* | 100 TF | 22,748 | 46 TB | 60 GB/s |



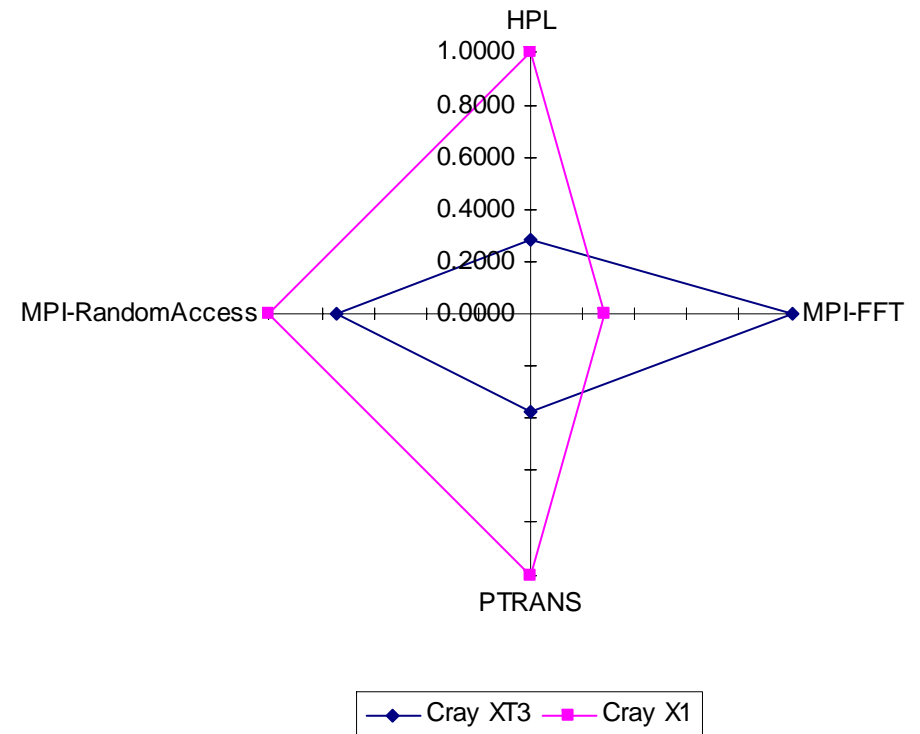
Different applications have different memory access patterns – “Once size does not fit all”



X1/XT3 Relative Global Benchmark Comparison

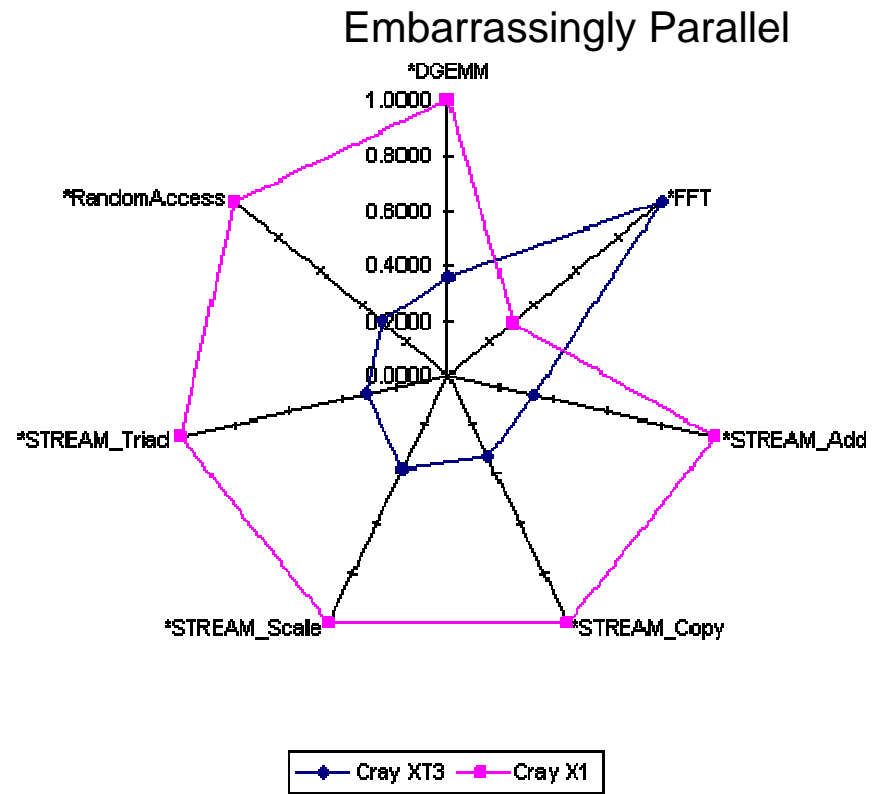
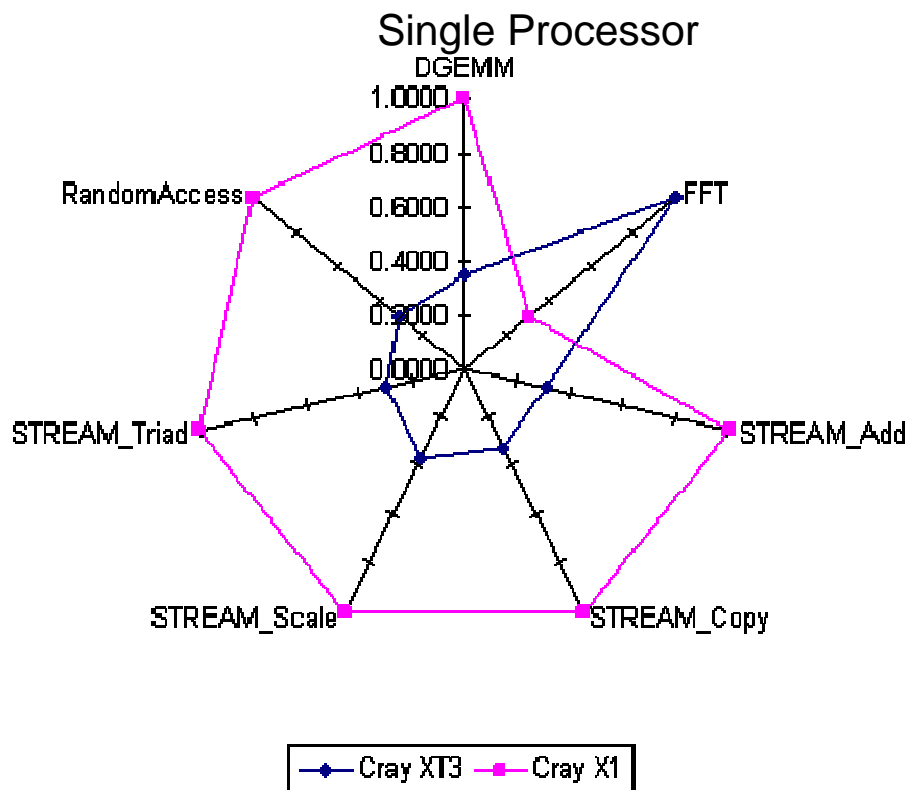
Relative performance on 64 CPUs

- Opteron scalar performance dominates MPI-FFT for XT3
- X1 bandwidth dominates MPI-RandomAccess and PTRANS
- X1 processor performance dominates HPL

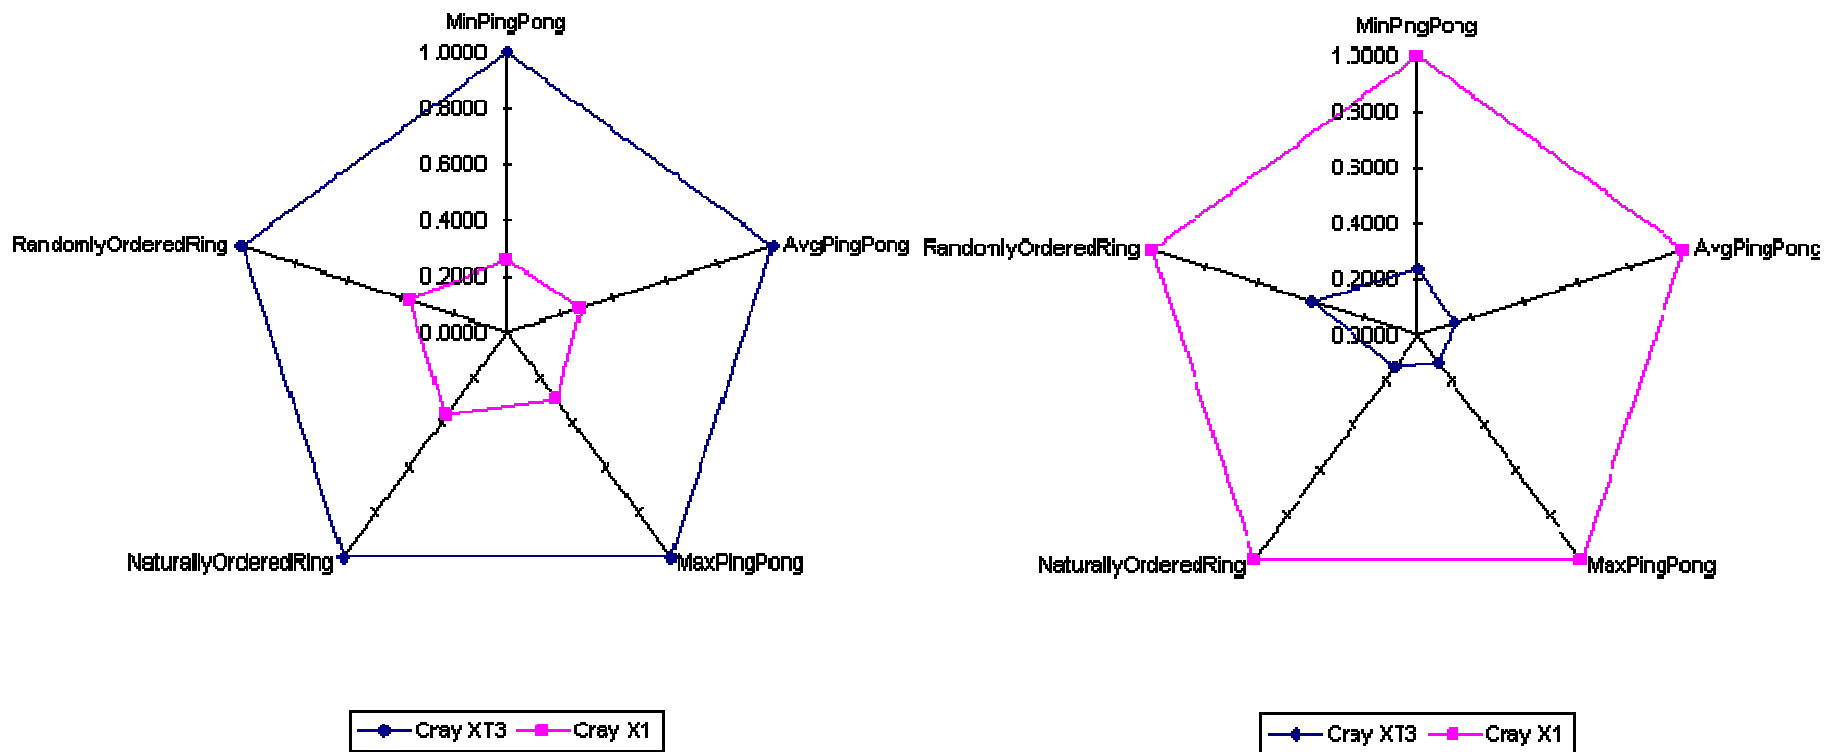


XT3 and X1 Relative Processor Performance


1D FFT can't be vectorized so the Opteron's 2.4GHz scalar performance is highlighted over the 400MHz X1 scalar performance



XT3/X1 Relative interconnect performance



GYRO: A 5-D gyrokinetic-Maxwell solver



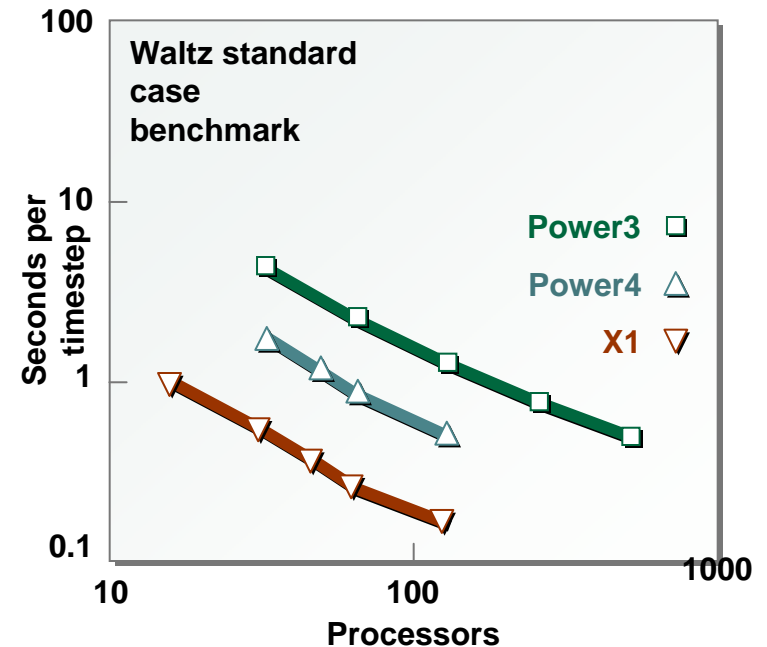

$$\frac{\partial f}{\partial t} = L_a f + L_b \Phi + \{f, \Phi\} \quad \text{where} \quad F\Phi = \iint dv_1 dv_2 f$$

GYRO employs Implicit-Explicit Runge-Kutta scheme discretized on Eulerian grid

For typical simulations, GYRO performance is at least 4x faster on Cray X1 than SGI Altix or IBM Power4;

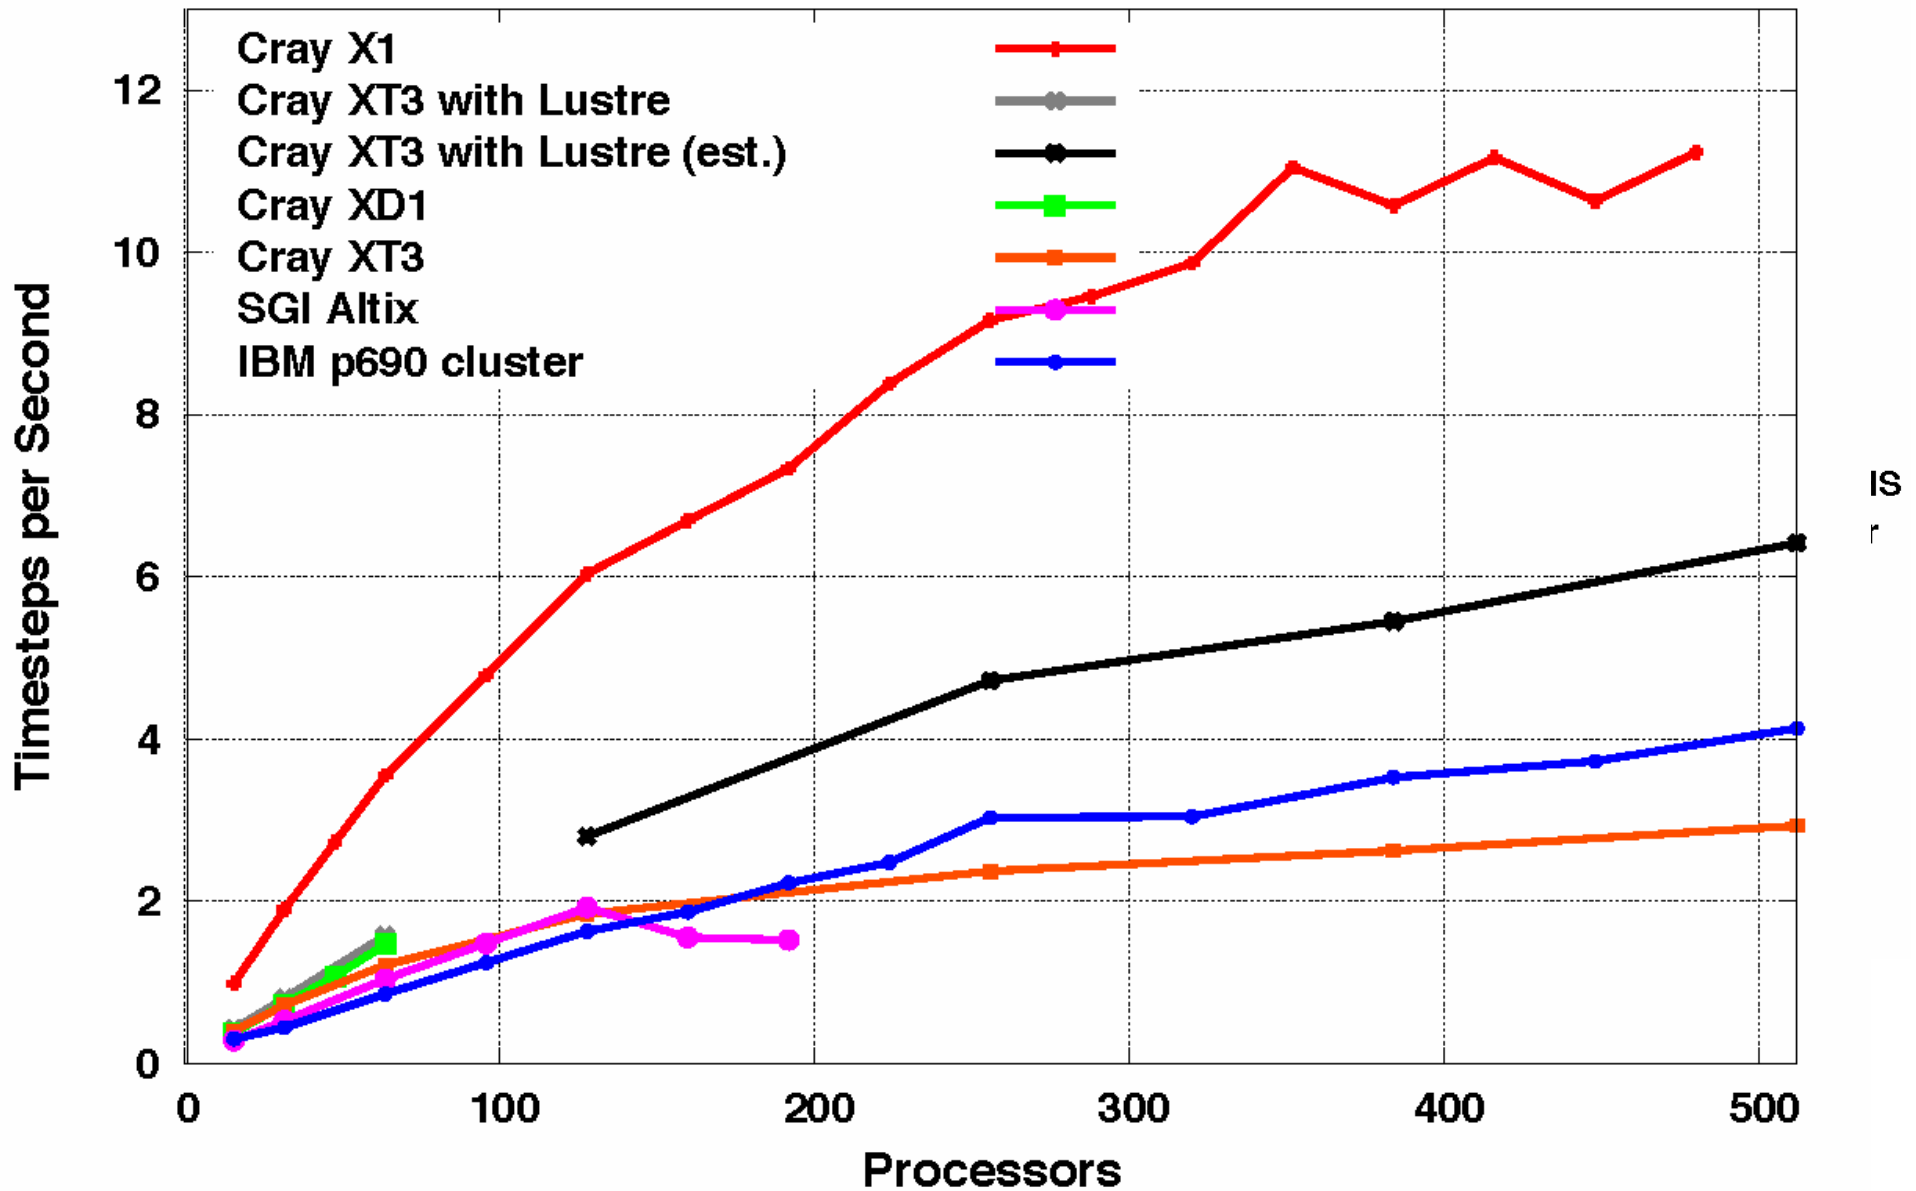
X1 is significantly faster for larger problems.

Recent accomplishments include: completed the most comprehensive tokamak turbulence simulation, showed transport is smooth across an $s=0$ (minimum- q surface) due to the appearance of gap modes, found that χ_i matched the "Cyclone value" at small ρ_* resolving an earlier conflicting report while simultaneously finding there is a long transient period for which χ_i exceeds the statistical average in GK simulations.



Impact of using NFS for IO on XT3

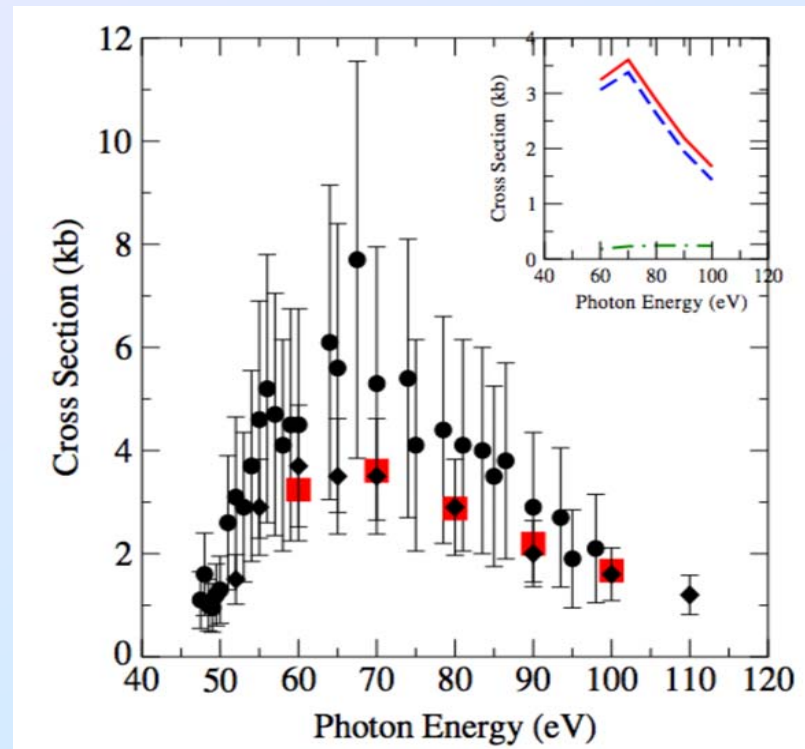
GYRO performance for B1-std



Double photoionization of H₂

James Colgan (LANL) and Mitch Pindzola (Auburn University)

- First ever fully *ab initio* calculation
- Important step in exploring the physics inherent in molecular break-up by a single photon
- 4D wave function propagated in time using Schroedinger equation
- Required hundreds of Cray X1 MSPs




“Time-dependent close-coupling calculations for the double photoionization of He and H₂.”
J Colgan, M S Pindzola, and F Robicieux. *J. Phys. B: At. Mol. Opt. Phys.* 37 (2004) L377–L384.

“Neutron star spin-up discovered with 3D simulations on Cray X1”

<http://astro.physics.ncsu.edu/TSI/>

- Terascale Supernova Initiative
- 3D hydrodynamics using 600 million zones
- Shading shows extent of supernova shock
- Streamlines show collapsing stellar material
 - Flow decelerates at shock
- Angular momentum of flow creates spinning neutron star



Some NLCF Science Highlights

- Fusion: Gyrokinetic simulation
 - **“GYRO Performance on a Variety of MPP Systems” (yesterday)**
- Climate
 - **“Porting and Performance of the CCSM3 on the Cray X1” (8:30 AM Thursday)**
- High-Temperature Superconductors
 - **Significant results submitted for publication**
- Chemistry
 - **Full-Configuration Interaction at 5.5 TF on 432 MSPs (60% of peak)**
 - **The new code already enables computations 40+x larger and 100+x faster than previous work**
- Combustion
 - **38% performance improvement using targeted Co-Array Fortran**

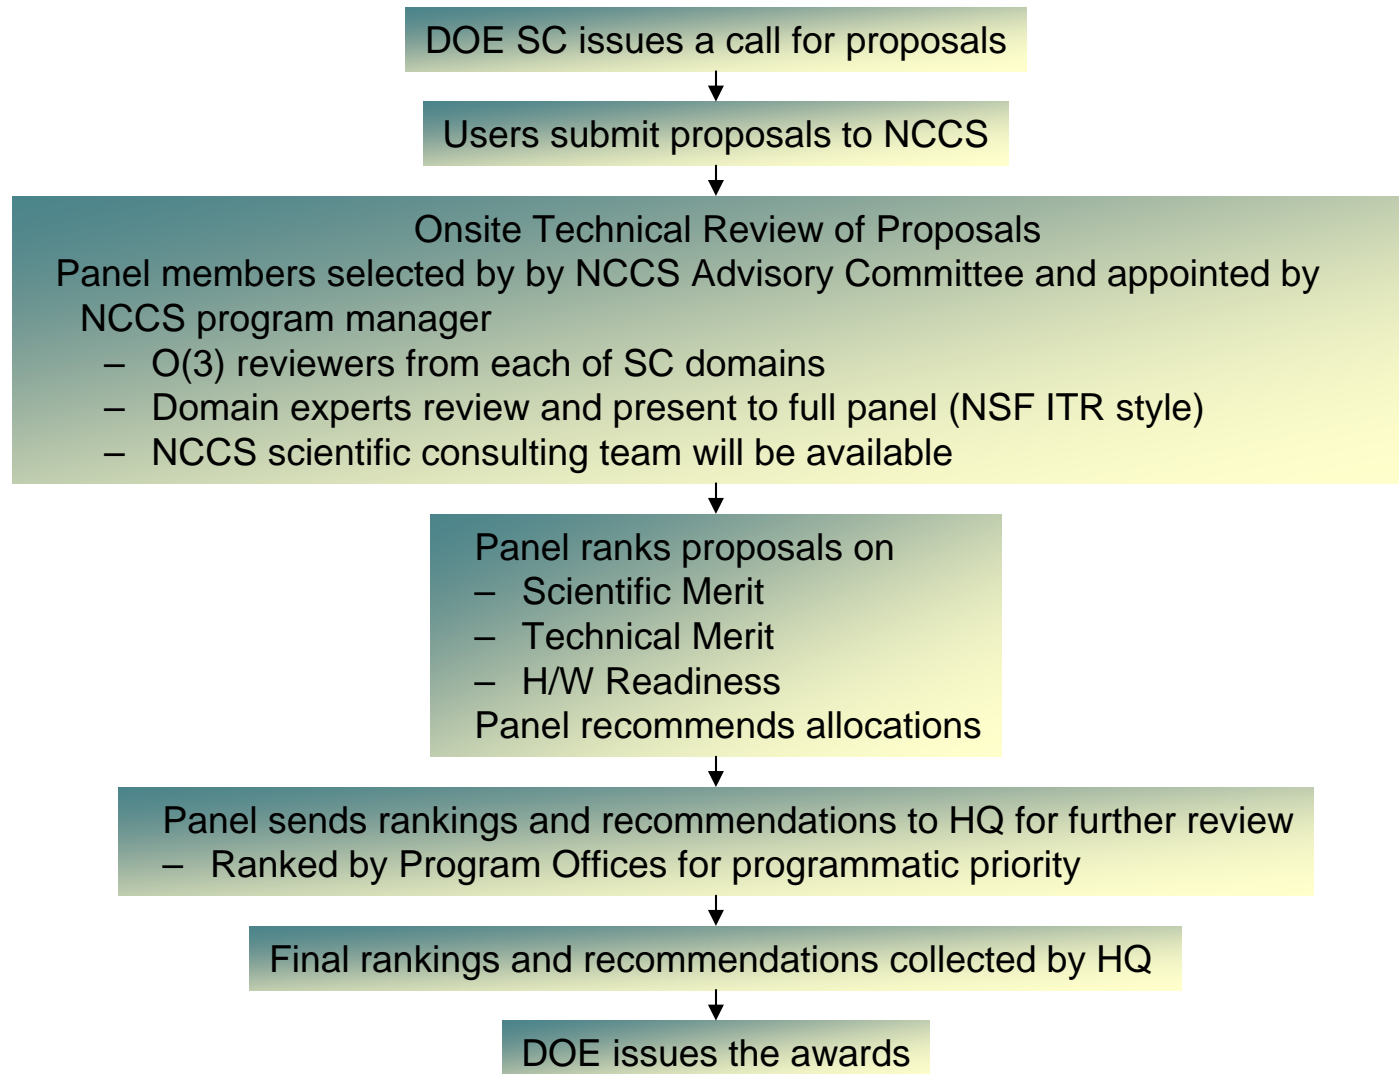
NLCF Applications at CUG 2005

- Early Evaluation of the Cray XD1 (2PM)
- Early Evaluation of the Cray XT3 at ORNL (3PM)
- Comparative Analysis of Interprocess Communication on the X1, XD1, and XT3 (2PM Wednesday)
- High-Speed Networking with Cray Supercomputers at ORNL (2:30PM Wednesday)
- Comparisons of the X1 to X1E on Major Scientific Applications (3PM Wednesday)
- Optimization of the PETSc Toolkit and Application Codes on the Cray X1 (11:30AM Thursday)

National Leadership Computing Facility 2006 Call for Proposals

- **Largest, most computationally intensive scientific endeavors to be conducted on high capability Leadership systems of DOE's National Leadership Computing Facility (NLCF)**
- **Seeks high priority, challenging, high-payoff, and heretofore intractable, computationally-intensive experiments, where capability of NLCF systems can enable new breakthroughs in science**
- **Expectation that Leadership systems will enable U.S. to be "first to market" with important scientific and technological capabilities, ideas, and software**
- **Selection of limited set of scientific applications (perhaps 10 per year) and given substantial access to leadership system**

NCCS Proposal Review Process



Summary

- The “Leadership” computing facility at ORNL is deploying systems to provide 100X more computational resources to critical problems.
- The Cray X1 and Cray XT3 complement each other by servicing different types of calculations.
- The NLCF at ORNL has the necessary support and is ready to deliver leadership computing to the nation.
- The largest challenge is proving to be getting scientists to think about problems that are 100x larger than what they are currently solving.

Thank You

To join our team please go to:
<http://computing.ornl.gov/Employment/postings.html>