The OU Supercomputing Center for Education & Research:
Building a Community of Computational Science & Engineering

Henry Neeman, Director

February 21, 2003

EDUCAUSE Southwest Regional Conference 2003

Outline

- Who, What, Where, When, Why, How
- OSCER efforts
  - Education
  - Research
  - Marketing
  - Resources
- OSCER’s future

What is CSE?

__Computational Science & Engineering__ (CSE) is the use of computers to simulate physical phenomena, or to optimize how physical systems are structured, or to discover new information hidden within them.

Most problems that are interesting to scientists and engineers are problems that are very very big, even though some of them are very very small.

For example, studying the relationships between the atoms in a few tens of thousands of molecules, or the movement of tornados across a state, or the formation of galaxies, can require TB of RAM, tens of TB of storage and weeks of CPU time.

What is Supercomputing?

- ***Supercomputing*** is the biggest, fastest computing **right this minute**.
- Likewise, a ***supercomputer*** is the biggest, fastest computer right this minute.
- So, the definition of supercomputing is constantly changing.
- ***Rule of Thumb***: a supercomputer is 100 to 10,000 times as powerful as a PC.
- ***Jargon***: supercomputing is also called ***High Performance Computing*** (HPC).

What is Supercomputing About?

Size **Speed**

What is Supercomputing About?

- **Size:** many problems that are interesting to scientists and engineers can't fit on a PC – usually because they need more than a few GB of RAM, or more than a few 100 GB of disk.
- **Speed:** many problems that are interesting to scientists and engineers would take a very very long time to run on a PC: months or even years. But a problem that would take a month on a PC might take only a few hours on a supercomputer.

What is HPC Used For?

- Simulation of physical phenomena, such as
  - Weather forecasting
  - Galaxy formation
  - Hydrocarbon reservoir management
- Data mining: finding needles of information in a haystack of data, such as
  - Gene sequencing
  - Signal processing
  - Detecting storms that could produce tornados
- Visualization: turning a vast sea of data into pictures that a scientist can understand

[1]

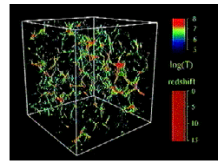

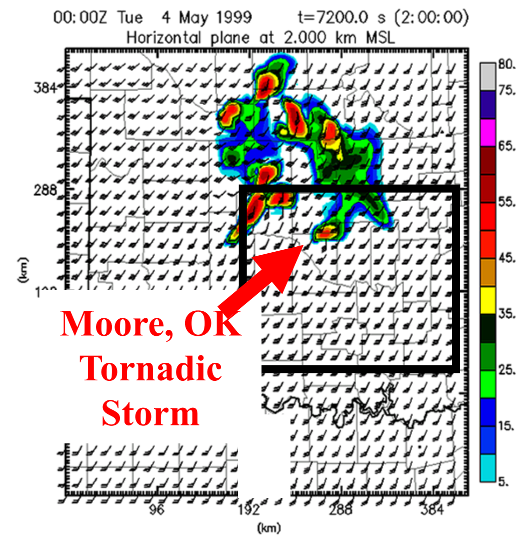

May 3 1999[2]

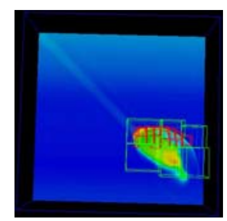

[3]

Linux Clusters

- Linux clusters are much cheaper than proprietary HPC architectures: factor of 5 to 10 in price/performance.
- They're largely useful for:
  - Distributed parallelism (message passing): hard to code!
  - Large numbers of single-processor applications

- MPI software design is not easy for inexperienced programmers because:
  - difficult programming model
  - lack of user-friendly documentation – emphasis on technical details rather than broad overview
  - hard to find good help
- BUT: at the national level, a few million dollars for MPI programmers is much much cheaper than tens or hundreds of millions for big iron – and the payoff lasts much longer.

What is OSCER?

- New, multidisciplinary center within OU's Department of Information Technology
- OSCER provides:
  - Supercomputing **education**
  - Supercomputing **expertise**
  - Supercomputing **resources:** hardware, storage, software
- OSCER is for:
  - Undergrad students
  - Grad students
  - Staff
  - Faculty

Who is OSCER? Departments

- Aerospace Engineering
- Astronomy
- Biochemistry
- Chemical Engineering
- Chemistry
- Civil Engineering
- Computer Science
- Electrical Engineering
- Industrial Engineering
- Geography
- Geophysics
- Management
- Mathematics
- Mechanical Engineering
- Meteorology
- Microbiology
- Molecular Biology
- OK Biological Survey
- Petroleum Engineering
- Physics
- Surgery
- Zoology

Colleges of Arts & Sciences, Business, Engineering, Geosciences and Medicine – with more to come!

Expected Biggest Consumers

- Center for Analysis & Prediction of Storms: daily real time weather forecasting
- Advanced Center for Genome Technology: on-demand genomics
- High Energy Physics: Monte Carlo simulation and data analysis

Who Are the Users?

161 users so far:

- 30 faculty
- 32 staff
- 93 students
- 6 off campus users

Comparison: National Center for Supercomputing Applications, with over $100M funding, has about 600 users.

OSCER Structure

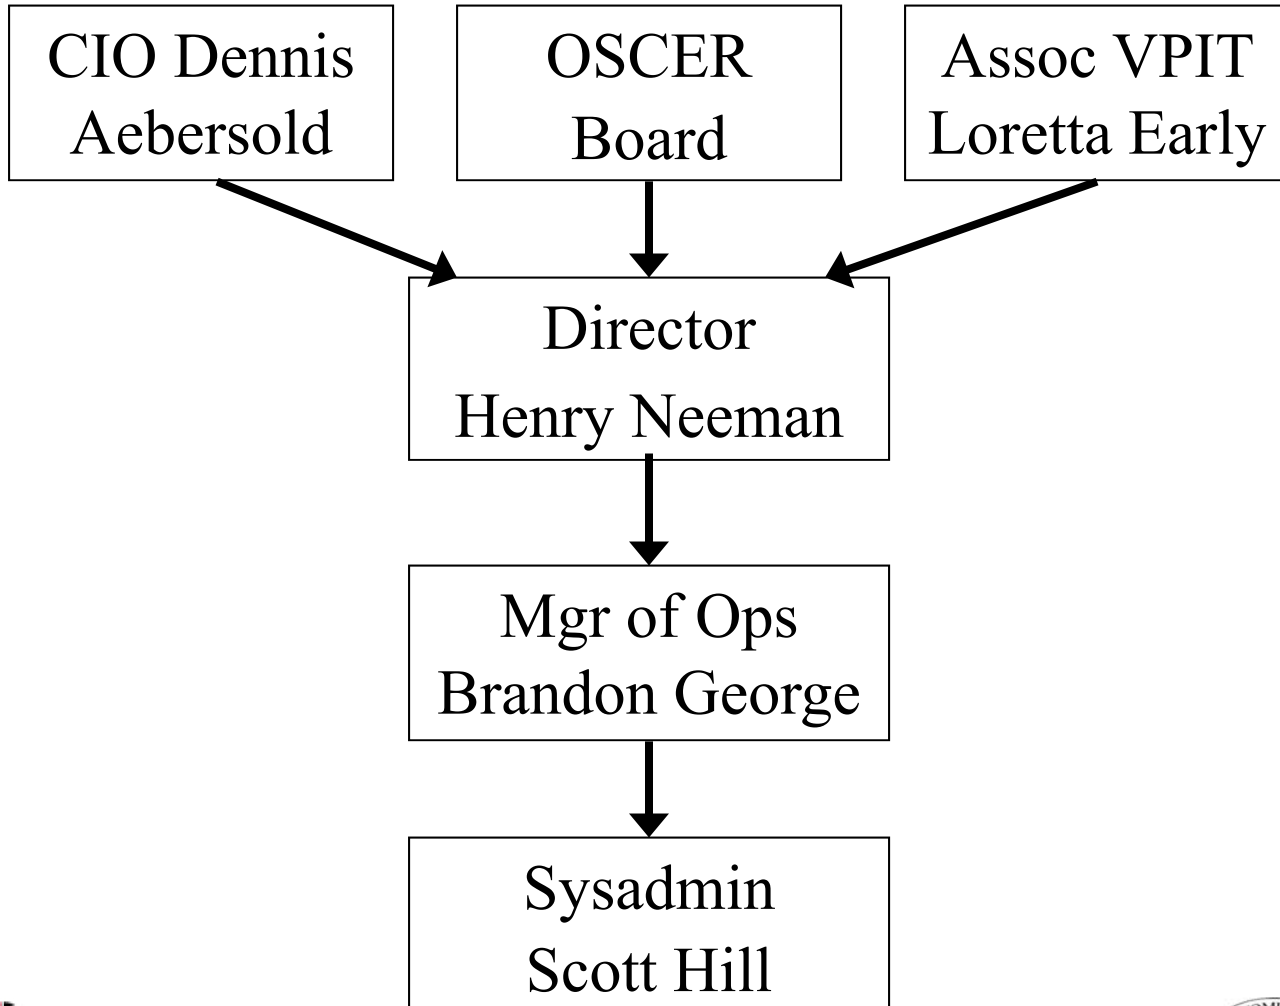

Who Works for OSCER?

- Director: Henry Neeman
- Manager of Operations: Brandon George
- System Administrator: Scott Hill (funded by CAPS)

Left to right:
Henry Neeman,
Brandon George,
Scott Hill

OSCER Board

- Arts & Sciences
  - Tyrrell Conway, Microbiology
  - Andy Feldt, Physics & Astro
  - Pat Skubic, Physics & Astro
- Engineering
  - S. Lakshmivarahan, Comp Sci
  - Dimitrios Papavassiliou, Chem Engr
  - Fred Striz, Aerospace & Mech Engr
- Geosciences
  - Kelvin Droegemeier, Meteorology/CAPS
  - Tim Kwiatkowski, CMRP
  - Dan Weber, CAPS

L to R: Papavassiliou, IBM VP for HPC Sales Peter Ungaro, Skubic, Striz, Neeman, Droegemeier, Weber

OSCER is Long Term

- OU recently broke ground on a new weather center complex, consisting of a Weather Center building and the Peggy and Charles Stephenson Research and Technology Center, which will house genomics, computer science (robotics), the US Geological Survey and OSCER.
- OSCER will be housed on the ground floor, in a glassed-in machine room and offices, directly across from the front door – a showcase!
- Scheduled opening: Spring 2004

Stephenson Center Floor Plan

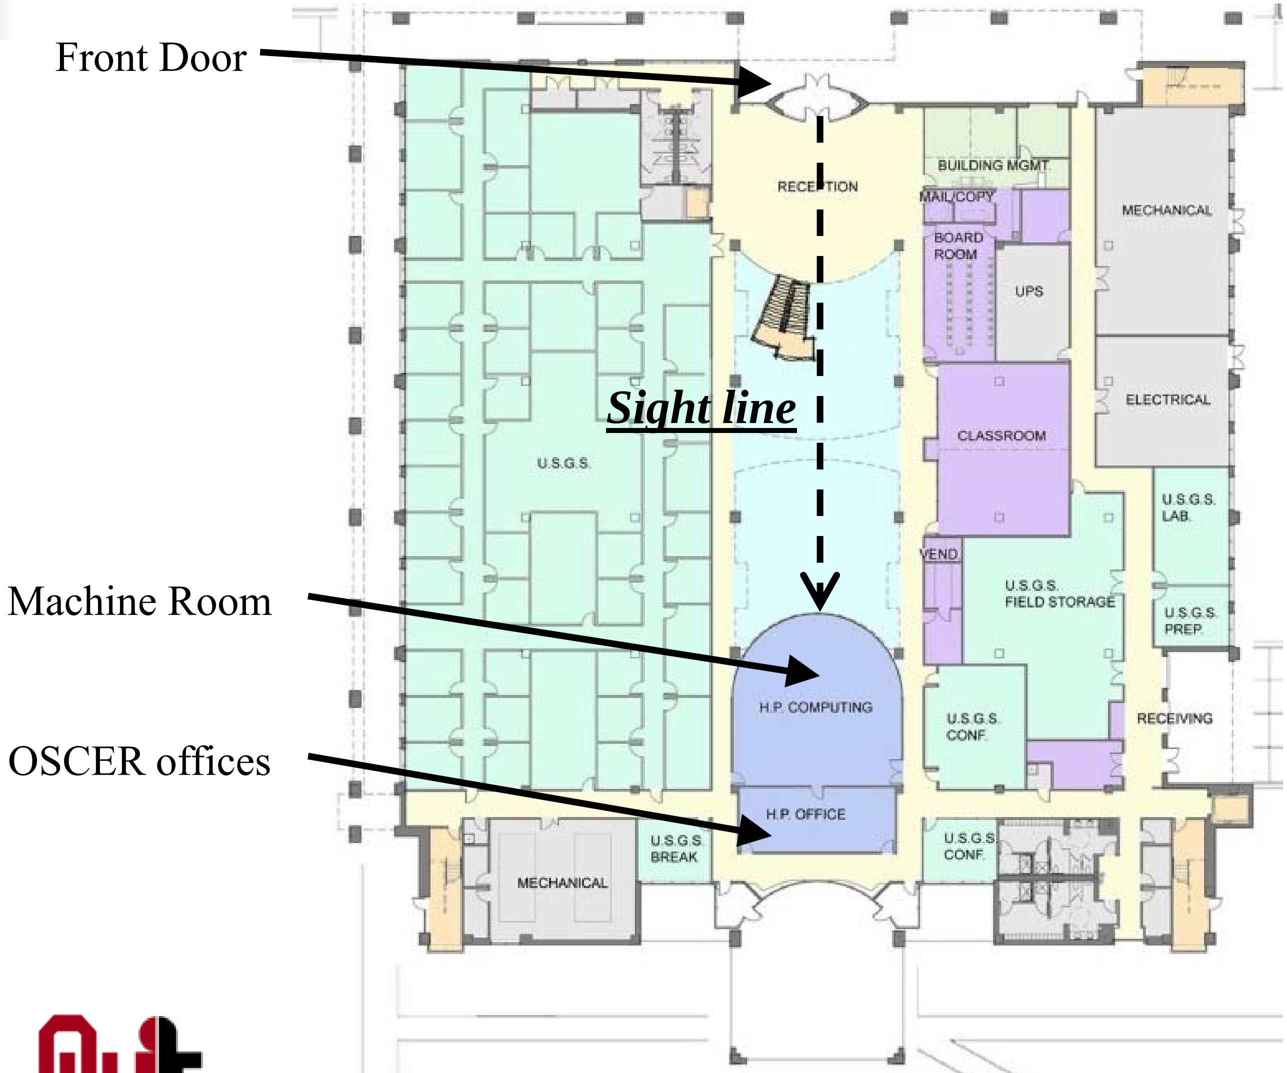

How Did OSCER Happen?

Cooperation between:

- OU High Performance Computing group: currently 119 faculty and staff in 19 departments within 5 Colleges
- OU CIO Dennis Aebersold
- OU VP for Research Lee Williams
- Williams Energy Marketing & Trading Co.
- OU Center for Analysis & Prediction of Storms
- OU School of Computer Science
- Encouragement from OU President David Boren, OU Provost Nancy Mergler, Oklahoma Congressman J.C. Watts Jr. (now retired), OU Assoc VPIT Loretta Early

Why OSCER?

- CSE has become sophisticated enough to take its place alongside observation and theory.
- Most students – and most faculty and staff – don't learn much CSE, because it's seen as needing too much computing background, and needs HPC, which is seen as very hard to learn.
- HPC can be hard to learn: few materials for novices; most documentation written for experts as reference guides.
- We need a new approach: HPC and CSE for computing novices – OSCER's mandate!

OSCER History

- Aug 2000: founding of OU High Performance Computing interest group
- Nov 2000: first meeting of OUHPC and OU Chief Information Officer Dennis Aebersold
- Jan 2001: Henry's "listening tour:" learning about what science & engineering researchers needed – education!!!
- Feb 2001: meeting between OUHPC, CIO and VPR; draft white paper about HPC at OU
- Apr 2001: Henry appointed OU IT's Director of HPC
- July 2001: draft OSCER charter released

OSCER History (continued)

- Aug 31 2001: OSCER founded; first supercomputing education workshop presented
- Sep 2001: OSCER Board elected
- Nov 2001: hardware bids solicited and received
- Dec 2001: OU Board of Regents approval
- March – May 2002: machine room retrofit
- Apr & May 2002: supercomputers delivered
- Sep 12-13 2002: 1st annual OU Supercomputing Symposium
- Oct 2002: first paper about OSCER's education strategy published

What Does OSCER Do?

- Teaching
- Research
- Marketing
- Resources

What Does OSCER Do? Teaching

Supercomputing in Plain English

An Introduction to High Performance Computing

Henry Neeman, Director
OU Supercomputing Center for Education & Research

Why is HPC Hard to Learn?

- HPC software technology changes very quickly:
  - Pthreads: 1988 (POSIX.1 FIPS 151-1) [4]
  - PVM: 1991 (version 2, first publicly released) [5]
  - MPI: 1994 (version 1) [6,7]
  - OpenMP: 1997 (version 1) [8,9]
  - Globus: 1998 (version 1.0.0) [10]
- Typically a 5 year lag (or more) between the standard and documentation readable by experienced computer scientists who aren't in HPC
  1. Description of the standard
  2. Reference guide, user guide for experienced HPC users
  3. Book for general computer science audience
- Documentation for novice programmers: very rare
- Tiny percentage of physical scientists & engineers ever learn this stuff

Why Bother Teaching Novices?

- Application scientists & engineers typically know their applications very well, much better than a collaborating computer scientist would ever be able to.
- Because of Linux clusters, CSE is now affordable.
- Commercial code development lags far behind the research community.
- Many potential CSE users don’t need full time CSE and HPC staff, just some help.
- Today’s novices are tomorrow’s top researchers, especially because today’s top researchers will eventually retire.

Educational Strategy

Workshops:

- Supercomputing in Plain English
  - Fall 2001: 87 registered, 40 – 60 attended each time
  - Fall 2002: 66 registered, c. 30 – 60 attended each time
  - Slides adopted by R. Wilhelmson of U. Illinois for Atmospheric Sciences' supercomputing course
  - Videos currently being used by OU School of Petroleum Engineering
- Performance evaluation workshop (fall 2002)
- Parallel software design workshop (fall 2002)

… and more to come.

Educational Strategy (cont'd)

Web-based materials:

- "Supercomputing in Plain English" (SiPE) slides
- Links to documentation about OSCER systems
- Locally written documentation about using local systems (coming soon)
- Introductory programming materials (developed for CS1313 Programming for Non-Majors)
- Introductions to Fortran 90, C, C++ (some written, some coming soon)
- Multimedia: SiPE workshops videotaped, soon available on DVD

Educational Strategy (cont'd)

Coursework

- Scientific Computing (S. Lakshmivarahan)
- Computer Networks & Distributed Processing (S. Lakshmivarahan)
- Nanotechnology & HPC (L. Lee, G.K. Newman, H. Neeman)
- Advanced Numerical Methods (R. Landes)
- Industrial & Environmental Transport Processes (D. Papavassiliou)
- Supercomputing presentations in other courses (e.g., undergrad numerical methods, U. Nollert)

Educational Strategy (cont'd)

Rounds: regular one-on-one (or one-on-few) interactions with several research groups

- Brainstorm ideas for applying supercomputing to the group's research
- Develop code
- Learn new computing environments
- Debug
- Papers and posters

Spring 2003: meeting with 20 research groups weekly, biweekly or monthly

Research

- OSCER's Approach
- Collaborations
- Rounds
- Funding Proposals
- Symposia

OSCER's Research Approach

- Typically, supercomputing centers provide resources and have in-house application groups, but most users are more or less on their own.
- OSCER's approach is unique: we partner directly with research teams, providing supercomputing expertise to help their research move forward faster. No one else in the world does this.
- This way, OSCER has a stake in each team's success, and each team has a stake in OSCER's success.

New Collaborations

- OU Data Mining group
- OU Computational Biology group – Norman campus and Health Sciences (OKC) campus working together
- Grid Computing group: OSCER, CAPS, Civil Engineering, Chemical Engineering, High Energy Physics, Aerospace Engineering
- … and more to come

Education & Research: Rounds

From left: Civil Engr undergrad from Cornell; CS grad student; OSCER Director; Civil Engr grad student; Civil Engr prof; Civil Engr undergrad

Rounds Participants: Fac & Staff

- John Antonio, Comp Sci
- Muhammed Atiquzzaman, Comp Sci
- Scott Boesch, Chemistry
- Dan Brackett, Surgery
- Bernd Chudoba, Aerospace Engr
- Yuriy Gusev, Surgery
- Randy Kolar, Civil Engr
- S. Lakshmivarahan, Comp Sci
- Lloyd Lee, Chem Engr
- Janet Martinez, Meteorology
- David Mechem, Cooperative Inst for Mesoscale Meteorological Studies
- Fekadu Moreda, Civil Engr
- Pia Mukherjee, Astronomy
- Jerry Newman, Chem Engr
- Dean Oliver, Petroleum Engr
- Dimitrios Papavassiliou, Chem Engr
- Tom Ray, Zoology
- Horst Severini, Physics
- Donna Shirley, Aerospace Engr
- Fred Striz, Aerospace Engr
- William Sutton, Mechanical Engr
- Baxter Vieux, Civil Engr
- Francie White, Mathematics
- Luther White, Mathematics
- Yun Wang, Astronomy
- Dan Weber, CAPS
- Ralph Wheeler, Chemistry
- Chenmei Xu, Zoology
- Mark Yeary, Electrical Engr

TOTAL TO DATE: 29 faculty & staff

Rounds Participants: Students

- Aerospace & Mechanical Engineering: 12
- Chemical Engineering & Materials Science: 6
- Chemistry & Biochemistry: 3
- Civil Engineering & Environmental Science: 5
- Computer Science: 3
- Electrical Engineering: 2
- Management: 1
- Meteorology: 2
- Petroleum Engineering: 3

TOTAL TO DATE: 31 students (undergrad, grad)

Research: Proposal Writing

- OSCER provides boilerplate text about not only resources but especially education and research efforts (workshops, rounds, etc).
- Faculty write in small amount of money for:
  - funding of small pieces of OSCER personnel;
  - storage (disk, tape);
  - special purpose software.
- In many cases, OSCER works with faculty in proposal development and preparation.

OSCER-Related Proposals #1

Funded:

- R. Kolar, J. Antonio, S. Dhall, S. Lakshmivarahan, "A Parallel, Baroclinic 3D Shallow Water Model," DoD - DEPSCoR (via ONR), $312K
- L. Lee, J. Mullen (Worcester Polytechnic), H. Neeman, G.K. Newman, "Integration of High Performance Computing in Nanotechnology," NSF, $400K
- J. Levit, D. Ebert (Purdue), C. Hansen (U Utah), "Advanced Weather Data Visualization," NSF, $300K
- D. Papavassiliou, "Turbulent Transport in Wall Turbulence," NSF, $165K
- M. Richman, A. White, V. Lakshmanan, V. De Brunner, P. Skubic, "A Real Time Mining of Integrated Weather Data," NSF, $950K
- D. Weber, H. Neeman, et al, "Continued Development of the Web-EH Interface and Integration with Emerging Cluster and Data Mining," Nat'l Ctr for Supercomputing Applications, $360K

TOTAL TO DATE: $2.4M to 15 OU faculty & staff

OSCER-Related Proposals #2

Submitted, decision pending

- D. Papavassiliou, H. Neeman, M. Zaman, "Multiple Scale Effects and Interactions for Darcy and Non-Darcy Flow," DOE, $436K
- D. Papavassiliou, H. Neeman, M. Zaman, "Integrated, Scalable Model Based Simulation for Flow Through Porous Media," NSF, $313K
- D. Papavassiliou, H. Neeman, "Development of a Lagrangian Methodology for Transport in Microscales using High Performance Computing," NSF, $500K
- H. Neeman, K. Droegemeier, K. Mish, D. Papavassiliou, P. Skubic, "Acquisition of an Itanium Cluster for Grid Computing," NSF, $465K
- D. Papavassiliou, R. Braatz (U. Illinois), J. McLaughlin (Clarkson), H. Neeman, T. Trafalis, "Development of a Grid-enabled Problem Solving Environment for Engineering Management Systems," NSF, $4M

TOTAL SUBMITTED: $5.7M

OSCER-Related Proposals #3

To be submitted:

- M. Atiquzzaman, H. Neeman, "Development of a Data Networks Course with On-site Mentoring by Network Professionals," NSF, $400K
- B. Chudoba, A. Striz, H. Neeman, "Development of a Parallel Design Environment for Preliminary Aerospace Design and Optimization," NSF, $300K
- D. Papavassiliou, M. Zaman, H. Neeman, "Integration of Computational Transport Processes and High Performance Computing Education," NSF, $400K
- H. Neeman et al, "Expansion of OSCER," NSF, $2M
- H. Neeman et al, "Incorporation of Computational Science & Engineering with High Performance Computing in Multidisciplinary Graduate Research," NSF, $2.95M

… and many more to come.

OSCER-Related Proposals #4

Rejected:

- "A Study of Moist Deep Convection: Generation of Multiple Updrafts in Association with Mesoscale Forcing," NSF
- "Use of High Performance Computing to Study Transport in Slow and Fast Moving Flows," NSF
- "Integrated, Scalable Model Based Simulation for Flow Through Reservoir Rocks," NSF
- "Hybrid Kilo-Robot Simulation Space Solar Power Station Assembly," NASA-NSF
- "Understanding and Interfering with Virus Capsid Assembly," NIH
- "Hydrologic Evaluation of Dual Polarization Quantitative Precipitation Estimates," NSF
- "A Grid-Based Problem Solving Environment for Multiscale Flow Through Porous Media in Hydrocarbon Reservoir Simulation," NSF

NOTE: Most of these will be resubmitted, or already have been in some form.

Supercomputing Symposium 2002

- Participating Universities: OU, Oklahoma State, Cameron, Langston, U Arkansas Little Rock
- Participating companies: Aspen Systems, IBM
- Other organizations: OK EPSCoR, COEITT
- 69 participants, including 22 students
- Roughly 20 posters
- Leveraging to build regional collaborations
- This was the first annual – we plan to do this every year.
- Symposium 2003 already planned and funded.

OSCER Marketing

- Media
- Other

OSCER Marketing: Media

- Newspapers
  - Norman Oklahoman, Dec 2001
  - OU Daily, May 2002
  - Norman Transcript, June 2002
- OU Football Program Articles
  - Fall 2001
  - Fall 2002 (OU-Texas)
- Television
  - "University Portrait" on OU's cable channel 22
- Press Releases

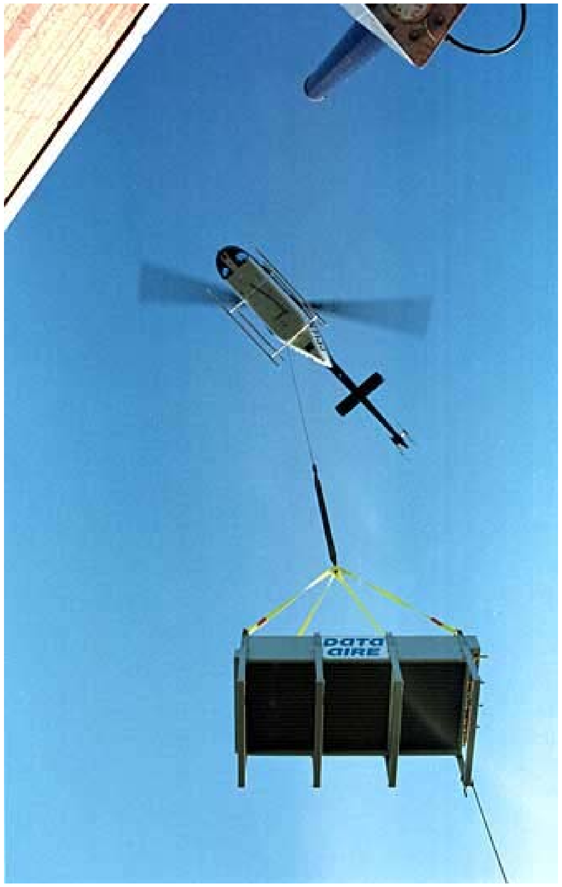

Norman Transcript 05/15/2002
Photo by Liz Mortensen

OSCER Marketing: Other

- OU Supercomputing Symposium
- OSCER webpage: **`www.oscer.ou.edu`**
- Participation at conferences
  - Supercomputing 2001
  - Alliance All Hands Meeting 2001
  - Scaling to New Heights 2002
  - Linux Clusters Institute HPC 2002
- Phone calls, phone calls, phone calls
- E-mails, e-mails, e-mails

OSCER Resources

- Purchase Process
- Hardware
- Software
- Machine Room Retrofit

Hardware Purchase Process

- Visits from and to several supercomputer manufacturers ("the usual suspects")
- Informal quotes
- Benchmarks (ARPS weather forecast code)
- Request for Proposals
- OSCER Board: 4 meetings in 2 weeks
- OU Board of Regents
- Negotiations with winners
- Purchase orders sent
- Delivery and installation

Machine Room Retrofit

- SEC 1030 was the best **existing** machine room for OSCER.
- But, it was nowhere near good enough when we started.
- Needed to:
  - Move a workstation lab out
  - Knock down a dividing wall
  - Install air conditioner piping
  - Install 2 large air conditioners (19 tons)
  - Install large Uninterruptible Power Supply (100 kVa)
  - Had it professionally cleaned – lots of sheetrock dust
  - Other miscellaneous stuff

OSCER Hardware

- IBM Regatta p690 Symmetric Multiprocessor
- Aspen Systems Pentium4 Linux Cluster
- IBM FAStT500 Disk Server
- Qualstar TLS-412300 Tape Library

OSCER Hardware: IBM Regatta

- 32 Power4 CPUs (1.1 GHz)
- 32 GB RAM
- 218 GB internal disk
- OS: AIX 5.1
- Peak speed: 140.8 GFLOP/s*
- Programming model:
  - shared memory
  - multithreading (OpenMP)
  - (also supports MPI)

*GFLOP/s: billion floating point operations per second

OSCER Hardware: Linux Cluster

264 Pentium4 XeonDP CPUs
264 GB RAM
8.7 TB disk (includes scratch)
OS: Red Hat Linux 7.3
Peak speed: > 1 TFLOP/s*
Programming model:
  distributed multiprocessing
  (MPI)

Aspen Systems Inc.

*TFLOP/s: trillion floating
  point operations per second

Linux Cluster Storage

Hard Disks

- EIDE 7200 RPM
  - Each Compute Node: 40 GB (operating system & local scratch)
  - Each Storage Node: 3 × 120 GB (global scratch)
  - Each Head Node: 2 × 120 GB (global home)
  - Management Node: 2 × 120 GB (logging, batch)
- SCSI 10,000 RPM
  - Each Non-Compute Node: 18 GB (operating sys)
  - RAID: 3 × 73 GB (realtime and on-demand systems)

IBM FAStT500 FC Disk Server

- 2200 GB hard disk: 30×73 GB FiberChannel
- IBM 2109 16 Port FiberChannel-1 Switch
- 2 Controller Drawers (1 for AIX, 1 for Linux)
- Room for 60 more drives: researchers buy drives, OSCER maintains them
- Expandable to 11 TB

Tape Library

- Qualstar TLS-412300
- Reseller: Western Scientific
- Initial configuration
  - 100 tape cartridges (10 TB)
  - 2 drives
  - 300 slots (can fit 600)
- Room for 500 more tapes, 10 more drives: researchers buy tapes, OSCER maintains them
- Software: Veritas NetBackup DataCenter, Storage Migrator
- Driving issue for purchasing decision: weight!

What Next?

- Waiting for to hear about submitted proposal for more hardware funding from NSF
- More users
- More rounds
- More workshops
- More collaborations (intra- and inter-university)
- **MORE PROPOSALS!**

A Bright Future

- OSCER's approach is unique, but it's the right way to go.
- People at the national level are starting to take notice.
- We'd like there to be more and more OSCERs around the country:
  - local centers can react better to local needs;
  - inexperienced users need one-on-one interaction to learn how to use supercomputing in their research.

References

[1] Image by Greg Bryan, MIT: http://zeus.ncsa.uiuc.edu:8080/chdm_script.html

[2] "Update on the Collaborative Radar Acquisition Field Test (CRAFT): Planning for the Next Steps." Presented to NWS Headquarters August 30 2001.

[3] See http://scarecrow.caps.ou.edu/~hneeman/hamr.html for details.

[4] S.J. Norton, M. D. Depasquale, *Thread Time: The MultiThreaded Programming Guide*, 1st ed, Prentice Hall, 1996, p. 38.

[5] A. Geist, A. Beguelin, J. Dongarra, W. Jiang, R. Manchek, V. Sunderam, PVM: Parallel Virtual Machine: A Users' Guide and Tutorial for Networked Parallel Computing. The MIT Press, 1994. http://www.netlib.org/pvm3/book/pvm-book.ps

[6] Message Passing Interface Forum, *MPI: A Message Passing Interface Standard.* 1994. http://www.openmp.org/specs/mp-documents/fspec10.pdf

[7] P.S. Pacheco, Parallel Programming with MPI. Morgan Kaufmann Publishers Inc., 1997.

[8] OpenMP Architecture Review Board, "OpenMP Fortran Application Program Interface." 1997. http://www.openmp.org/specs/mp-documents/fspec10.pdf

[9] R. Chandra, L. Dagum, D. Kohr, D. Maydan, J. McDonald, R. Menon, Parallel Programming in OpenMP. Morgan Kaufmann Publishers Inc., 2001.

[10] Globus News Archive. http://www.globus.org/about/news/