


OU Supercomputing Center for Education & Research

State of the Center Address 2002

Henry Neeman, Director
September 12, 2002





Outline

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





What is OSCER?

- New, multidisciplinary center within OU's Department of Information Technology
- OSCER provides:
 - Supercomputing **education**
 - Supercomputing **expertise**
 - Supercomputing **resources**: hardware and software
- OSCER is for:
 - OU undergrad students
 - OU grad students
 - OU staff
 - OU faculty
 - Their collaborators





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!





Who is OSCER? Centers

- Advanced Center for Genome Technology
- Center for Analysis & Prediction of Storms
- Center for Aircraft & Systems/Support Infrastructure
- Coastal Meteorology Research Program
- Center for Engineering Optimization
- Cooperative Institute for Mesoscale Meteorological Studies
- DNA Microarray Core Facility
- High Energy Physics
- Institute of Exploration & Development Geosciences
- National Severe Storms Laboratory
- Oklahoma EPSCoR



Expected Biggest Consumers

- Center for Analysis & Prediction of Storms: daily realtime weather forecasting



- Advanced Center for Genome Technology: on-demand bioinformatics
- High Energy Physics: Monte Carlo simulation and data analysis



Who Works for OSCER?

- Director: Henry Neeman
- Manager of Operations: Brandon George
- System Administrator: Scott Hill (funded by CAPS)
- Student Programmer: Lyal Grissom, Comp Sci



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. Lakshmiarahan, 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 Peter Ungaro, Skubic, Striz, Neeman, Droegemeier, Weber





Where is OSCER?

- Machine Room: Sarkeys Energy Center 1030 (shared with Geosciences Computing Network; Schools of Meteorology, Geography, Geology & Geophysics; Oklahoma Climatological Survey) – for now ...

Take the tour!

- Henry's office: SEC 1252
- Brandon & Scott's office: SEC 1014





Where Will OSCER Be?

OU is about to break ground on a new weather center complex, consisting of a weather center building and a “G+” building housing genomics, computer science (robotics) 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!





Why OSCER?

- Computational Science & Engineering (CSE) is 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!





How Did OSCER Happen?

Cooperation between:

- OU High Performance Computing group
- OU CIO Dennis Aebersold
- OU VP for Research Lee Williams
- OU President David Boren
- Williams Energy Marketing & Trading
- OU Center for Analysis & Prediction of Storms
- OU School of Computer Science





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"
- Feb 2001: meeting between OUHPC, CIO and VP for Research Lee Williams; draft white paper about HPC at OU released
- Apr 2001: Henry Neeman named Director of HPC for Department of Information Technology
- July 2001: draft OSCER charter released



OSCER History (continued)

- Aug 31 2001: OSCER founded; first supercomputing education workshop presented
- Nov 2001: hardware bids solicited and received
- Dec 2001: OU Board of Regents approves purchase of supercomputers
- 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?

- Education
- 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





Educational Strategy

Workshops:

- Supercomputing in Plain English
 - Fall 2001: 87 registered, 40 – 60 attended each time
 - Fall 2002: 64 registered, c. 60 attended Sep 6
 - Slides adopted by R. Wilhelmson of U. Illinois for Atmospheric Sciences' supercomputing course
 - All day IBM Regatta workshop (fall 2002)
 - Performance evaluation workshop (fall 2002)
 - Parallel software design workshop (fall 2002)
 - Introductory batch queue workshops (soon)
- ... and more to come.





Educational Strategy (cont'd)

Web-based materials:

- “Supercomputing in Plain English” slides
- SiPE workshops being videotaped for streaming
- 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)





Educational Strategy (cont'd)

Coursework

- Scientific Computing (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 optimization and parallelization
 - Develop code
 - Learn new computing environments
 - Debug
 - Papers and posters





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 partners directly with research teams, providing supercomputing expertise to help their research move forward faster.
- 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 and Health Sciences campuses working together
- Chemical Engineering and High Energy Physics: Grid computing
- ... 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, Computer Science
- Scott Boesch, Chemistry
- Randy Kolar, Civil Engineering
- S. Lakshmivaran, Comp Sci
- Lloyd Lee, Chemical Engineering
- Janet Martinez, Meteorology
- David Mechem, CIMMS
- Fekadu Moreda, Civil Engineering
- Dimitrios Papavassiliou, Chemical Engineering
- Tom Ray, Zoology
- Horst Severini, Physics
- Fred Striz, Aerospace & Mechanical Engineering
- William Sutton, Aerospace & Mechanical Engineering
- Baxter Vieux, Civil Engineering
- Francie White, Mathematics
- Luther White, Mathematics
- Yun Wang, Astronomy
- Dan Weber, CAPS
- Ralph Wheeler, Chemistry
- Chenmei Xu, Zoology

TOTAL TO DATE: 22 faculty & staff





Rounds Participants: Students

- Aerospace & Mechanical Engineering: 10
- Chemical Engineering & Materials Science: 5
- Chemistry & Biochemistry: 3
- Civil Engineering & Environmental Science: 5
- Computer Science: 3
- Electrical Engineering: 2
- Management: 1
- Meteorology: 2

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 some cases, OSCER works with faculty in proposal development and preparation.





OSCER-Related Proposals #1

Accepted:

- 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

TOTAL TO DATE: \$2.1M to 14 OU faculty & staff





OSCER-Related Proposals #2

Pending

- A. Zlotnick et al, “Understanding and Interfering with Virus Capsid Assembly,” NIH, \$1.25M
- B. Vieux et al, “Hydrologic Evaluation of Dual Polarization Quantitative Precipitation Estimates,” NSF, \$438K
- H. Neeman, D. Papavassiliou, M. Zaman, R. Alkire (UIUC), J. Alameda (UIUC), “A Grid-Based Problem Solving Environment for Multiscale Flow Through Porous Media in Hydrocarbon Reservoir Simulation,” NSF, \$592K
- D. Papavassiliou, H. Neeman, M. Zaman, “Multiple Scale Effects and Interactions for Darcy and Non-Darcy Flow,” DOE, \$436K

TOTAL PENDING: \$2.7M





OSCER-Related Proposals #3

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

NOTE: Some of these will be resubmitted.





Supercomputing Symposium 2002

- Participating Universities: OU, OSU, TU, UCO, Cameron, Langston, U Arkansas Little Rock, Wichita State
- Participating companies: Aspen Systems, IBM
- Other organizations: OK EPSCoR, COEITT
- 60 – 80 participants
- Roughly 20 posters
- Let's build some multi-institution collaborations!
- This is the first annual – we plan to do this every year.






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, 2002
 - 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





Purchase Process Heroes

- Brandon George
- OSCER Board
- Florian Giza & Steve Smith, OU Purchasing
- Other members of OUHPC
- Vendor sales teams
- OU CIO Dennis Aebersold





Machine Room Retrofit

- SEC 1030 is the best machine room for OSCER.
- But, it was nowhere near good enough when we started.
- Needed to:
 - Move the AMOCO workstation lab out
 - Knock down the dividing wall
 - Install 2 large air conditioners
 - Install a large Uninterruptible Power Supply
 - Have it professionally cleaned – lots of sheetrock dust
 - Other miscellaneous stuff





Retrofit Heroes

- Brandon George, OSCER
- OU Physical Plant
 - Gary Ward
 - Dan Kissinger
 - Brett Everett
 - OU Electrical
- Warden Construction: Dan Sauer
- Natkin Piping
- SealCo Cleaning (Dallas)





OSCER Hardware

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



OSCER Hardware: IBM Regatta

32 Power4 CPUs

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





IBM Regatta p690

- 32 Power4 1.1 GHz CPUs (4.4 GFLOP/s each)
- 1 MB L1 Data Cache (32 KB per CPU)
- 22.5 MB L2 Cache (1440 KB per 2 CPUs)
- 512 MB L3 Cache (32 MB per 2 CPUs)
- 32 GB ChipKill RAM
- 218 GB local hard disk (global home, operating system)
- Operating System: AIX 5.1
- Peak Computing Speed: 140.8 GFLOP/s
- Peak Memory Bandwidth: 102.4 GB/sec



OSCER Hardware: Linux Cluster

264 Pentium4 Xeon CPUs

264 GB RAM

2.5 TB global disk

OS: Red Hat Linux 7.3

Peak speed: > 1 TFLOP/s

Programming model:
distributed multiprocessing
(MPI)



*TFLOP/s: trillion floating
point operations per second





Linux Cluster

- 264 Pentium4 XeonDP CPUs (4 GFLOP/s each)
- 2 MB L1 Data Cache (8 KB per CPU)
- 132 MB L2 Cache (512 KB per CPU)
- 264 GB RAM (1 GB per CPU)
- 2500 GB hard disk available for users
- Myrinet-2000 Interconnect (250 MB/sec)
- Operating System: Red Hat Linux 7.3

- Peak Computing Speed: 1,056 GFLOP/s
- Peak Memory Bandwidth: 844 GB/sec
- Peak Interconnect Bandwidth: 32 GB/sec



Linux Cluster Nodes

- Breakdown of Nodes
 - 132 Compute Nodes (computing jobs)
 - 8 Storage Nodes (Parallel Virtual File System)
 - 2 Head Nodes (login, compile, debug, test)
 - 1 Management Node (PVFS control, batch queue)
- Each Node
 - 2 Pentium4 XeonDP CPUs (2 GHz, 512 KB L2 Cache)
 - 2 GB RDRAM (400 MHz, 3.2 GB/sec)
 - Myrinet-2000 adapter





Linux Cluster Storage

Hard Disks

- EIDE 7200 RPM
 - Each Compute Node: 40 GB (operating system & local scratch)
 - Each Storage Node: 2×120 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: 3×36 GB (realtime and on-demand systems)



IBM FASTT500 FC Disk Server

- 2,190 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



Tape Library

- Qualstar TLS-412300
- Reseller: Western Scientific
- Initial configuration: very small
 - 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
- Driving issue for purchasing decision: weight!





Software: IBM Regatta p690

- **Campus Base Product Group:** AIX 5.1, XLC compiler, CSet++ compiler, XLFortran compiler & Runtime Environment, Performance Toolbox/Aide, Engineering & Scientific Subroutine Library (ESSL)
- **Campus Scalable POWERparallel (SP) Group:** Parallel System Support Program, Parallel Environment, Parallel Optimization Subroutine Library, LoadLeveler, Parallel ESSL, XL High Performance Fortran & Runtime Environment





Software: Linux Cluster

- Red Hat Linux 7.3
- **System Management:** Aspen Systems Cluster Management Software, Beowatch, System Imager
- **Message Passing:** MPICH, LAM/MPI, PVM, SCA Linda
- **Scheduler:** Sun GridEngine
- Parallel Virtual File System (PVFS)





Software: Linux Cluster (cont'd)

- **Compilers:** Portland Group Fortran90, C, C++, HPF; Intel Fortran, C/C++; GNU g77, gcc, g++; NAG f95
- **Numerical Libraries:** ATLAS BLAS, LAPACK, ScaLAPACK, PLAPACK, PETSc, FFTW, etc





Software: Both

- Approved by the OSCER Board
 - IMSL numerical libraries
 - Fortran Lint & C LintPlus source code analyzer
- Final approval pending
 - TotalView debugger
- Under discussion
 - Vampir/VampirTrace performance analyzer
 - MATLAB
 - Fluent
 - NASTRAN





What Next?

- Finish configuring the machines
- Get everyone on (44 accounts so far)
- 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.





Thanks!

Join us in Tower Plaza Conference Room A for a tutorial on Performance Evaluation by Prof S. Lakshmivarahan of OU's School of Computer Science.

Thank you for your attention.

