

# Cowboy Up

Dana Brunson

Oklahoma State University

High Performance Computing Center

<http://hpc.it.okstate.edu>

NSF MRI Award # 1126330



# What is OSU HPCCC?

- Multidisciplinary Center to provide free cyberinfrastructure resources, expertise and education, not only at OSU but across OK.
- Our mission is to help our users be successful in computational research, education and outreach.
- We provide proposal support and collaboration – letters, facilities, data management, outreach, optimization etc.

# Who is OSU HPCCC?

- Director: Dana Brunson (new title April 2012)
- Manager of Operations: Jesse Schafer (new hire Feb 2012)
- Undergraduate HPC guru: Will Flanery
- Over 230 users in 24 departments in 6 Colleges as well as users from across the state.

# Who are the users?

Over 230 users so far, including:

- 45 faculty
- 14 undergraduate students
- 150 graduate students
- 5 postdocs
- 6 staff

Plus collaborators and other researchers from around the state and more added weekly.

# Who is OSU HPCC – Academic Depts.

- Ag Econ
- Animal Science
- Biochemistry & Molecular Biology
- Biosystems & Ag Engineering
- Botany
- Chemical Engineering
- Chemistry
- Civil & Env. Engineering
- Computer Science
- Elec. & Comp. Engineering
- Entomology & Plant Pathology
- Fire Protection
- Industrial Eng. & Mgmt
- Mechanical & Aerospace Eng.
- Mgmt Science & Info. Systems
- Mathematics
- Microbiology & Molecular Genetics
- Natural Resource Ecology & Mgmt
- Nutritional Science
- Physics
- Physiological Sciences
- Plant & Soil Science
- Statistics
- Zoology



# Publications Facilitated by OSUHPCC

- 2009: 3 Master's Reports
- 2010: 4 reported, including 1 dissertation
- 2011: 19 reported
- 2012: 18 reported so far, including 2 dissertations

We are getting better at collecting data.

See: <http://hpcwiki.it.okstate.edu/index.php/Publications>



# Background – prehistory

2005-6

- Faculty plan for first centrally available cluster
- Bullet: 64 nodes, 820 GFLOPs
- Administered by one person in IT





# Background: 2007-8


- November, 2007
  - 19 users
  - Hired replacement person (yours truly.)
  - Met Henry & Jeff
- In 2008 added
  - 50 new users
  - Spur: 128 GB RAM
  - Cimarron: 14 nodes, 892 GFLOPs
  - Student helper





# Background: 2009-10

- In 2009 added
  - 20 new users
  - Pistol Pete: 64 nodes, 5.4 TFLOPs
  - EOT endeavors
- In 2010, added 70 users and needed:
  - More resources
  - More staff



# Finding solutions

- The HPC community is very helpful!
- Henry Neeman & Jeff Pummill are great teachers!
- Solution: write NSF MRI proposal for hardware and use the cost share to hire a new person, with lots of help.



# The proposal highlights

- MRI research projects represented:  
44 faculty, 10 staff, 7 postdocs, 80 graduate students and 29 undergraduates
- VP for Research & Technology Transfer (VPRTT) agreed to fund the cost share.
- VPRTT & CIO committed to space, power & cooling and continued funding of new hire!
- Broader Impacts: Oklahoma Cyberinfrastructure Initiative, diversity, dissemination, CI EOT efforts.



# MRI Research Projects

- **Mechanical Behavior of Granular Materials and Molecular Dynamics (MD) Simulation.** R. Komanduri, Mechanical & Aerospace Engineering
- **Discovery genes for Canine Hip Dysplasia.** L. Zhu, Statistics
- **Improvements to performance per Watt (computing efficiency) of many-core systems.** S. Sohoni, Electrical & Computer Engineering
- **Transcriptional Profiling of determination events in adult and embryonic murine stem cell lines.** P. Hoyt, Biochemistry & Molecular Biology
- **Hoyt Collaborative Research: Muscle Adaptations Permitting Fatigue-Resistant Exercise.** M. Davis, Physiological Sciences
- **Hoyt Collaborative Research: Using comparative transcriptome and immunological analyses to “connect the dots” for bats affected by White-Nose Syndrome.** R. Van Den Bussche, Dean, Arts & Sciences:



# MRI Research Projects (cont'd)

- **Genomic, metagenomic and proteomic approaches to decipher host-pathogen interactions, complex carbohydrate metabolism and cellulosic bioenergy.** Patricia Ayoubi-Canaan, Biochemistry & Molecular Biology
- **Partial Differential Equations Modeling Geophysical Fluids.** J. Wu, Mathematics
- **Computational Chemistry.** N. Materer, Chemistry
- **Electronic Structure Simulations of Nanostructured Materials.** J. W. Mintmire, Physics
- **Dynamics and Scalable Atomistic Monte Carlo Simulation of the Growth of Carbon Nanotubes and Other Nano-Scale Structures.** S. Bukkapatnam, Industrial Engineering & Management
- **Methods and Applications on Modeling and Simulating Gas-Phase and Condensed-Phase Materials.** Y. Guo, Physics



# MRI Research Projects (cont'd)

- **Imaging Complex Geology Using High-Performance Computational Seismology.** P. Jaiswal, Geology
- **Computational Optimization.** B. Balasundaram, Industrial Engineering & Management
- **Optimal error-control coding and compressive sensing techniques for applications such as efficient data acquisition and reliable data transmissions in communication networks.** N. Rahnavard, Electrical & Computer Engineering
- **Charge, Spin and Heat Transport in Low-Dimensional Systems.** X. Xie, Physics
- **Molecular phylogeny of the Asteraceae.** L. Watson, Botany
- **Evolutionary genetics of morphological diversification and domestication in grasses (Poaceae) and mustards (Brassicaceae).** A. Doust, Botany



# MRI Research Projects (cont'd)

- **Robust Electromagnetic Field Testing and Simulation Laboratory.** J. West, Electrical & Computer Engineering
- **Phylogenomic analyses of the extremophile red alga *Galdieria sulphuraria*.** G. Schoenknecht, Botany
- **Integrating Data in Evolving Social Networks.** S. Kak, Computer Science
- **Computational and combinatorial methods in commutative algebra.** C. Francisco, Mathematics
- **Characterization of environmental microbes associated with non-cultivated and cultivated plants.** U. Melcher, Biochemistry and Molecular Biology
- **Application of ensemble based- simulated annealing algorithms(EBSA) to radio-biologically-based Intensity-Modulated Radiation Therapy (RB IMRT) planning.** F. Fondjo, Chair, Technology, Langston U
- **Phylogenomics of milkweeds (*Asclepias*, Apocynaceae).** M. Fishbein, Botany





# OK Cyberinfrastructure Initiative

- Triggered by Oklahoma's current NSF EPSCoR Track-1 grant.
- Memorandum of Understanding between OU and OSU.
- All academic institutions in Oklahoma are eligible to sign up for free use of OU's and OSU's centrally-owned CI resources.
- Other kinds of institutions (government, NGO, commercial) are eligible to use, though not necessarily for free.
- Everyone can participate in our CI education initiative.
- The Oklahoma Supercomputing Symposium, our annual conference, continues to be offered to all.
- Plus new things coming soon!



# And we got the grant!

## Next steps:

- Hire new Manager of Operations
  - Jesse started Feb. 6, 2012
- Evaluate the evolving needs of researchers
- Go out for bid for the new cluster
- Expand the data center (unexpected!)



# HPC to meet our needs

- Potential of Sandy Bridge, worth waiting?
- Variety of core counts per node, memory per node?
- Large shared memory nodes?
- Disk storage, fast/slow?
- GPUs?
- Network?

RFP went out in January, PO went out late Feb



# The UPS arrived June 5, 2012



# Cowboy arrived June 19, 2012

- Cowboy is funded by NSF grant for \$908,812.
- Cost share funded new full-time position.
- Cowboy is the largest externally funded supercomputer in state history!
- 9x the capacity of the current cluster Pistol Pete






# Jesse lining up the racks




# Cowboy's parts

- 252 compute nodes with dual 6-core CPUs and 32 GB RAM
- 2 “fat” nodes with 256 GB RAM and GPU cards
- 120 TB very fast disk storage
- QDR infiniband (high speed, low latency interconnect)
- 10G connectivity to Internet2
- 2 management and 2 login servers
- Over 48 TFLOPs





# Cooling – rear door heat exchangers



# Cowboy UP!





# Making users happy (or not)

- Making max walltimes shorter
  - Some codes have unpredictable walltimes, need reliable checkpoint/restart system
  - Short walltimes make system maintenance less intrusive
  - Short walltimes make the queue move faster
- Environment modules
  - Dynamic modification of user's env. Variables.
  - Can also log usage of applications



# In progress

- Early user phase
- Migrating users – slowly to minimize disruption of their work
- Installing user applications
- Help sessions, new user training sessions
- Wiki
- Workshops



# Workshops



2012: Bioinformatics of  
Entangled Genomes  
2009, 2011: Computational  
Chemistry



# What's next?

- Plan for next expansion
- Sustainability for both hardware and support personnel
- Condor pool
- Provide more education, outreach and training

# Thanks!

## Questions?

Dana Brunson

[dana.brunson@okstate.edu](mailto:dana.brunson@okstate.edu)

OSU High Performance Computing Center

<http://hpc.it.okstate.edu>

