# High-End Computing in New York State


**Russ Miller**
**Director, Cyberinfrastructure Lab**
**Dept of Comp Sci & Eng, SUNY-Buffalo**
**Hauptman-Woodward Inst.**



**NSF, NIH, DOE, NIMA, NYS, HP**


**www.cse.buffalo.edu/faculty/miller/CI/**

GRID

Advanced
Center for Computational Research
Data
Center

# Miller's Cyberinfrastructure Lab

- **Working Philosophy**
  - **CI sits at core of modern simulation & modeling**
  - **CI allows for new methods of investigation to address previously unsolvable problems**
- **Focus of MCIL is on development of *algorithms, portals, interfaces, middleware***
- **Goal of MCIL is to free end-users to do disciplinary work**
- **Funding (2001-pres)**
  - **NSF: ITR, CRI, MRI**
  - **NYS appropriations**
  - **Federal appropriations**

# MCIL Equipment

- Experimental Equipment (1.25 TF; 140 Cores; 22TB)
  - Clusters
    - Head Nodes: Dell 1950 (Intel)
    - Workers: Intel 8×2 ×4 & AMD 8×2×2
  - Virtual Memory Machines (2 × Intel 4×4)
  - Dell GigE Managed Switches
  - InfiniBand
  - 22 TB Dell Storage (2)
  - Condor Flock (35 Intel/AMD)
- Acquisition in Progress (30-50TF)
- Production Equipment
  - Dell Workstations; Dell 15 TB Storage
  - Access to CCR equipment (13TF Dell/Intel clusters)

# Evolution of CI Lab Projects

- Buffalo-Based Grid
  - Experimental Grid: Globus & Condor
  - Integrate Data & Compute, Monitor, Portal, Node Swapping, Predictive Scheduling/Resource Management
  - GRASE VO: Structural Biology, Groundwater Modeling, Earthquake Eng, Comp Chemistry, GIS/BioHazards
  - Buffalo, Buffalo State, Canisius, Hauptman-Woodward
- Western New York Grid
  - Heterogeneous System: Hardware, Networking, Utilization
  - Buffalo, Geneseo, Hauptman-Woodward, Niagara
- New York State Grid
  - Extension to Hardened Production-Level System State-Wide
  - Albany, Binghamton, Buffalo, Geneseo, Canisius, Columbia, HWI, Niagara, [Cornell, NYU, RIT, Rochester, Syracuse, Marist], {Stony Brook, RPI, Iona}

# NYS Grid Resources

- Albany: 8 Dual-Processor Xeon Nodes
- Binghamton: 15 Dual-Processor Xeon Nodes
- Buffalo: 1050 Dual-Processor Xeon Nodes
- Cornell: 30 Dual-Processor Xeon Nodes
- Geneseo State: Sun/AMD with 128 Compute Cores
- Hauptman-Woodward Institute: 50 Dual-Core G5 Nodes
- Marist: 9 P4 Nodes
- Niagara University: 64 Dual-Processor Xeon Nodes
- NYU: 58 Dual-Processor PowerPC Nodes
- RIT: 4 Dual-Processor Xeon Nodes
- Syracuse: 8 Dual-Processor Xeon Nodes

# Center for Computational Research (CCR)

- Founding Director (1998-2006)
- Facts & Figures
  - Top Academic HPC Center in World
  - ~25 TF of HPC
  - ~600 TB of High-End Storage
  - Significant Visualization
  - Special-Purpose Systems
  - ~30 FTEs Staff
  - 140 Projects Annually
- EOT
- [illegible]
- ROI: ~$450M to WNY

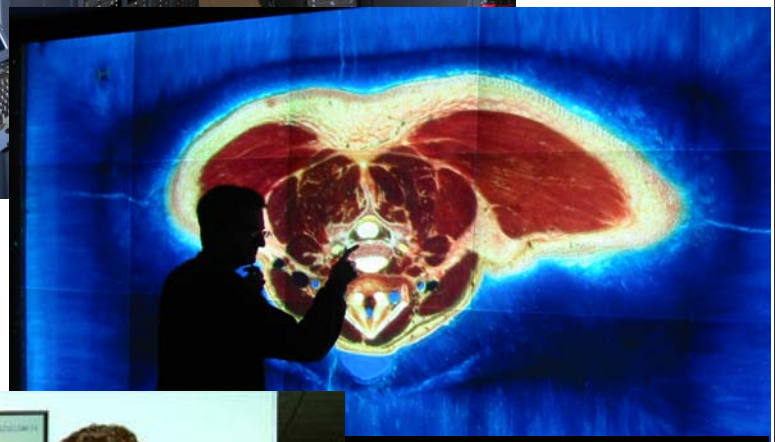

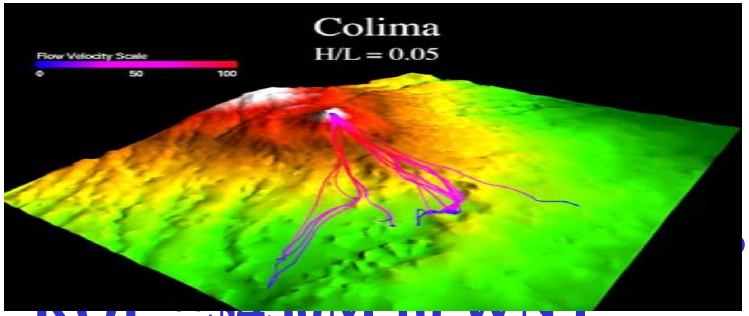


# CCR Highlights (1998-2006)

- Provide HE-Comp
- Provide HE-Vis + AGN
- Special Purpose Systems
  - Bioinformatics
  - Data Warehouse / Mining
- Support Local/National Efforts – Industry + Acad
- Create jobs in WNY
- Certificate Program
- Workshops + Tours
  - Campus, Industry
  - High-School
- Urban Planning & Design
- MTV Videos
- Peace Bridge, Med Campus
- Olmsted Parks, Thruway
- NYS Agencies
- Elected Officials
- Magnet on Campus
- Significant Funds
- Numerous Awards
- Significant Publicity

# CCR Research & Projects

- Archaeology
- Bioinformatics/Protein Folding
- Computational Chemistry
- Computational Fluid Dynamics
- Data Mining/Database
- Earthquake Engineering
- Environ Modeling & Simulation
- Grid Computing
- Molecular Structure Determination
- Physics

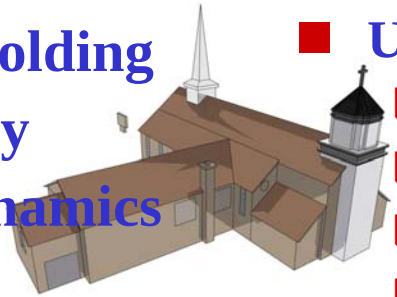

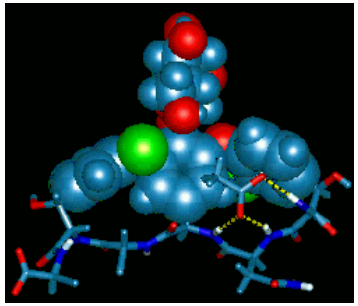

- Videos: MTV
- Urban Simulation and Viz
  - StreetScenes
  - I-90 Toll Barrier
  - Medical Campus
  - Peace Bridge
- Accident Reconstruction
- Scientific Viz
  - Dental
  - Surgery
  - MRI/CT Scan
  - Confocal Microscopy
  - Crystallization Wells
  - Collaboratories


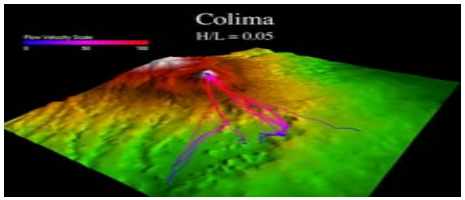


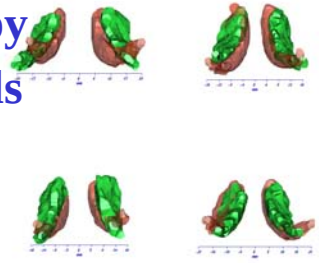

# SUNY-Buffalo 2008 (13TF; 75TB)

- Dell Linux Cluster (10TF peak)
  - 1600 Xeon EM64T Processors (3.2 GHz)
  - 2 TB RAM; 65 TB Disk
  - Myrinet / Force10
  - 30 TB EMC SAN
- Dell Linux Cluster (2.9TF peak)
  - 600 P4 Processors (2.4 GHz)
  - 600 GB RAM; 40 TB Disk; Myrinet
- HP/Compaq SAN
  - 75 TB Disk; 190 TB Tape
  - 64 Alpha Processors (400 MHz)
  - 32 GB RAM; 400 GB Disk

# RPI (110 TF)

- Partnership with NYS and IBM ($100M total including in-kind contributions)
- Computational Center for Nanotechnology Innovations (CCNI)
- Power-Based AIX Cluster
- AMD Cluster (2168 processors; 10 TF)
- IBM BlueGene/L (32K processors; 92 TF)
  - Rank #12 on 11/2007 Top 500
- Some availability outside of RPI

# Brookhaven National Laboratory (103 TF)

- Partnership (BNL/Stony Brook University) with NYS and IBM ($100M total including in-kind contributions)
- New York Center for Computational Sciences (NYCCS)
- Blue Gene/L (36864 Processors)
- 100 TF; Rank #10 on 11/2007 Top 500
- Applications include wide range of Science and Engineering
- Potential availability outside of BNL/SBU

# Rochester Institute of Technology (75TF; 67 TB)

- Campus Computational Grid
- Condor Flock with 1000 Cores
  - Windows, Macintosh, Solaris, Linux
  - Eight (8) Departments Participating
- 1.5 TB of Memory
- 67 TB of Disk
- 75 TF of Computing

# Cornell University

- Cornell University Center for Advanced Computing (new name for “The Theory Center”)
- Store several large Databases
- Part of TeraGrid
- EOT


www.cse.buffalo.edu/faculty/miller