

Enabling Discovery & Innovation via High-End Computing

Russ Miller

Cyberinfrastructure Lab, SUNY-Buffalo
Hauptman-Woodward Med Res Inst



NSF, NIH, DOE, NIMA, NYS, Dell

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

Academia in the 21st Century

- **Empower students to compete in knowledge-based economy**
- **Embrace digital data-driven society**
- **Accelerate discovery and comprehension**
- **Embrace relationships between academia and industry**
- **Provide increased Education, Outreach, and Training**
- **Enhance virtual organizations**



Academia in the 21st Century: Implementation

- **Support HPC infrastructure, research, and applications**
- **Deliver high-end cyberinfrastructure to enable efficient**
 - ❑ **Collection of data**
 - ❑ **Management/Organization of data**
 - ❑ **Distribution of data**
 - ❑ **Analysis of data**
 - ❑ **Visualization of data**
- **Create links between enabling technologists and disciplinary users**
- **Improve efficiency of knowledge-driven applications in myriad disciplines**
 - ❑ **New Techniques**
 - ❑ **New Algorithms**
 - ❑ **New Interactions (people & systems)**

Research environment that supports internet-based computation with goal of deriving novel scientific theories and generating knowledge;
Core of modern simulation and modeling;
Provides entirely new methods of investigation




Center for Computational Research (CCR): 1998-2006

■ Founding Director

■ 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





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 (Simulation & Modeling)


- 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



CCR Funding (1998-2006)

- CCR-Enabled to SUNY-Buffalo
 - ❑ \$170M External Funds
 - ❑ \$140M In-Kind Contributions
- CCR-Enabled to WNY
 - ❑ \$200M External Funds
- Federal Appropriations
- New York State Appropriations
- Local WNY Foundations
- In-Kind Contributions (Dell, SGI, Sun, etc.)
- Grants (NSF, NIH, DOE, etc.)
- Projects with Local Companies
- Government Projects
- SUNY-Buffalo: staff and space



Real-Time Visualization

StreetScenes: Real-Time 3D Traffic Simulation

- **Accurate local landmarks: Bridges, Street Signs, Business, Homes**
- **Can be viewed from driver's perspective**
- **Real-Time Navigation**
- **Works with**
 - **Corsim**
 - **Synchro**
- **Generate AVI & MOV**
- **Multiple Simultaneous**
 - **Traffic Loads**
 - **Simulation**
 - **Varying POV**



Animation & Simulation

Rendered Scenes

Peace Bridge Visualization: Animation & Simulation

■ Proposed Options

❑ Relocate US plaza

❑ Build a 3-lane companion span
& rehab existing bridge




PHOTO AND STORY BY BRUCE JACKSON



University at Buffalo *The State University of New York*

Cyberinfrastructure Laboratory

CI Lab

MTV

IBC Digital & CCR

Song: I'm OK (I Promise)

Band: Chemical Romance

Gaming Environment: Death Jr.



University at Buffalo *The State University of New York*

Cyberinfrastructure Laboratory

CI Lab

Virtual Reality

Alive on the Grid: PAAPAB

■ Networked art application for CAVE


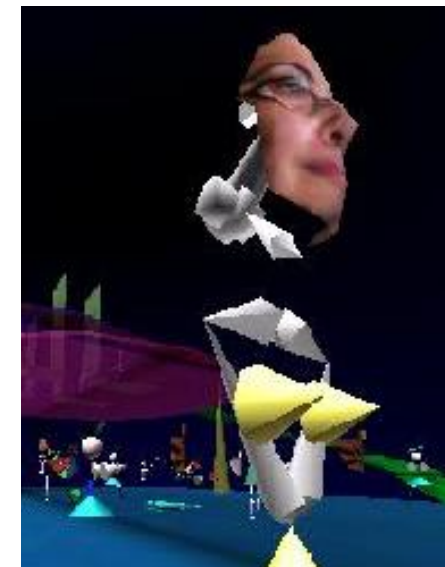

- Users from around the world
- First performance 2001

■ Dance-floor environment

- Inhabited by life-size puppets
- Dance with each other
- Synchro

■ Recording Booth

- User enters booth
- User dances
- System records dance from tracking on head and hands
- Dance mapped to Avatar





J. Anstey



VR-Fact!

- Interactive virtual factory
- Creates digital mock-up of factory
- Drag & place modular machines
- Mathematical algorithms for consistency checks



Kesh



Collaborative Visualization Environments

- Enable distributed collaboration via software developed at CCR
- Enable visualization and interaction with data across a geographically disparate network topology
- Integrate multiple data sources:
 - Scientific
 - Multimedia
- Research Topics
 - Distributed databases
 - OpenGL 3D programming
 - 3D Modeling
 - Character animation
 - User interaction
 - Virtual Reality



A. Ghadersohi, R. Miller, M. Green



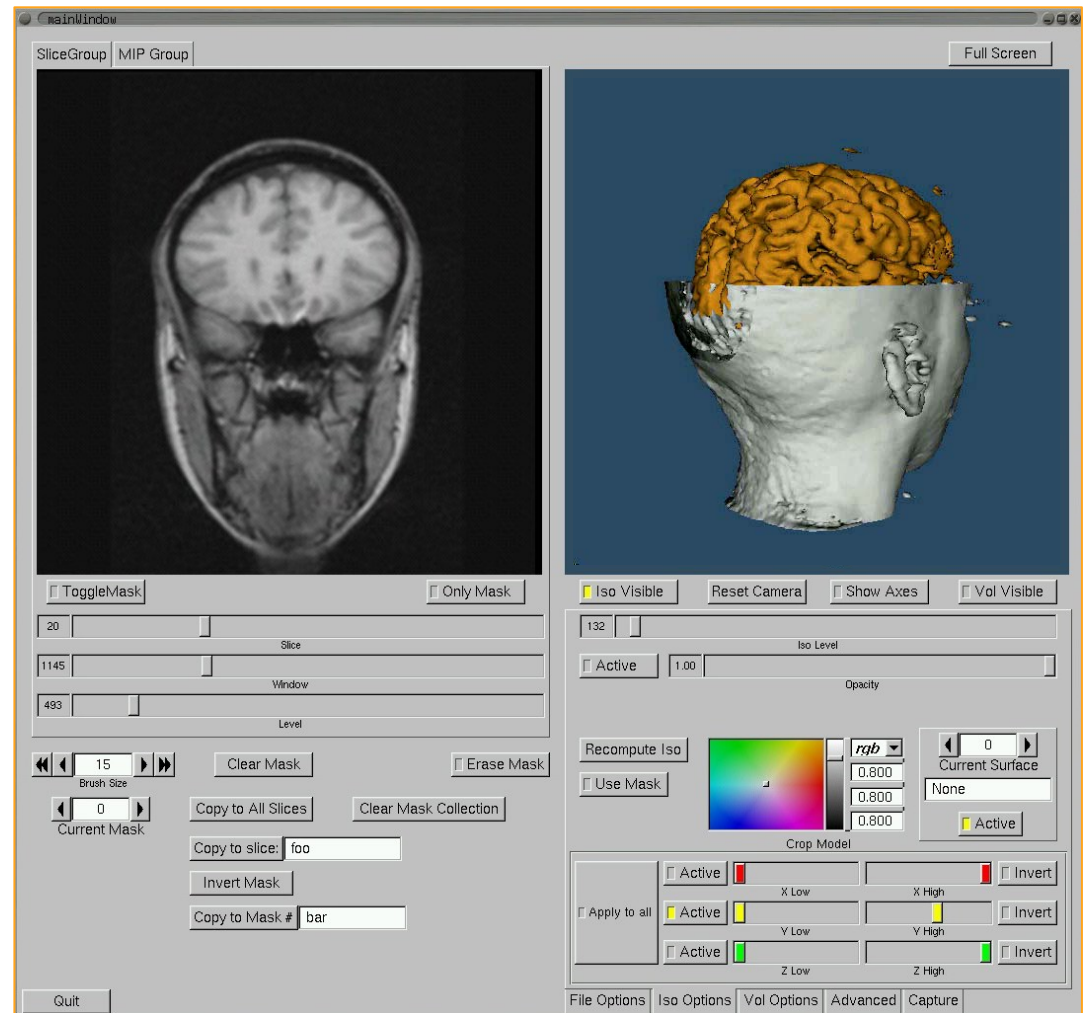
Scientific Visualization

Multiple Sclerosis Project

- Collaboration with Buffalo Neuroimaging Analysis Center (BNAC)


- Developers of Avonex, drug of choice for treatment of MS

- MS Project examines patients and compares scans to healthy volunteers



3D Medical Visualization

- Reads data output from a CT or MRI Scan
- Collaboration with Children's Hospital
- Visualize multiple surfaces and volumes
- Export images, movies or CAD file
- Pre-surgical planning
- Runs on a PC



M. Innus



University at Buffalo The State University of New York

Cyberinfrastructure Laboratory

CI Lab


Mapping Brain Activity

Positron emission tomography (PET), shows sites activated and deactivated as subjects decide whether a sound is a target or not.



Current density maps of brain surface (1–700 ms after target) show dynamic pattern of brain activity during decision-making process.

A. Lockwood

Sites Activated



Sites Deactivated




Science & Engineering

Small Subset of Projects

Groundwater Flow Modeling

- Regional scale modeling of groundwater flow and contaminant transport (Great Lakes)
- Ability to include all hydrogeologic features as independent objects
- Based on *Analytic Element Method*
- Key features:
 - Highly parallel
 - Object-oriented programming
 - Intelligent user interface
- Utilized 42 years of CPU time on CCR computers in 1 calendar year




A. Rabideau, I. Jankovic, M. Becker



Avalanches, Volcanic and Mud Flows

Geology, Math, Engineering

- Modeling of Volcanic Flows, Mud flows (flash flooding), and avalanches
- Integrate information from several sources
 - Simulation results
 - Remote sensing
 - GIS data
- Present information to decision makers using custom visualization tools local & remote
- GRID enabled for remote access
- Key Features
 - Parallel Adaptive Computation
 - Integrated with GIS System for flows on natural terrain



Flow models of Colima volcano
In Mexico – courtesy Rupp et. al.'06

A. Patra, B. Pitman, M. Sheridan, M. Jones



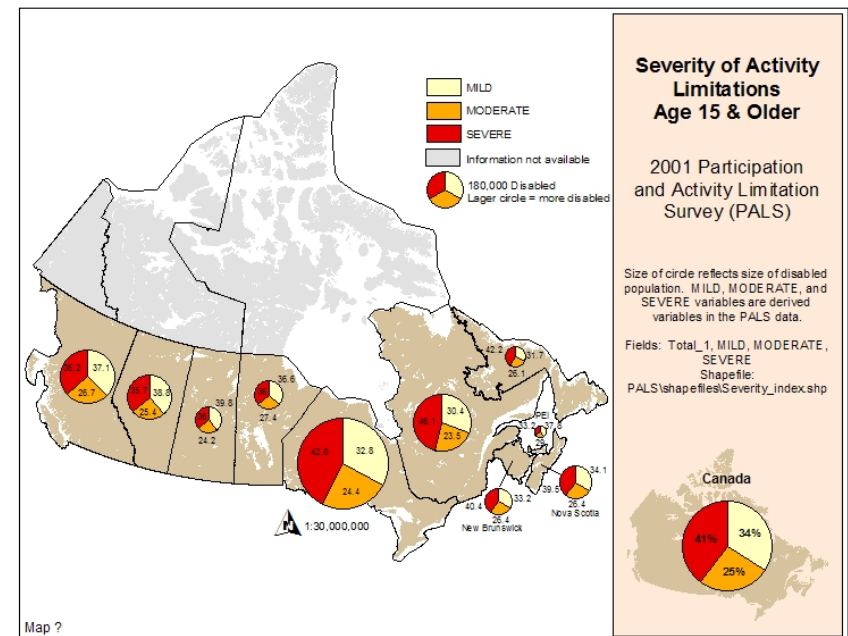
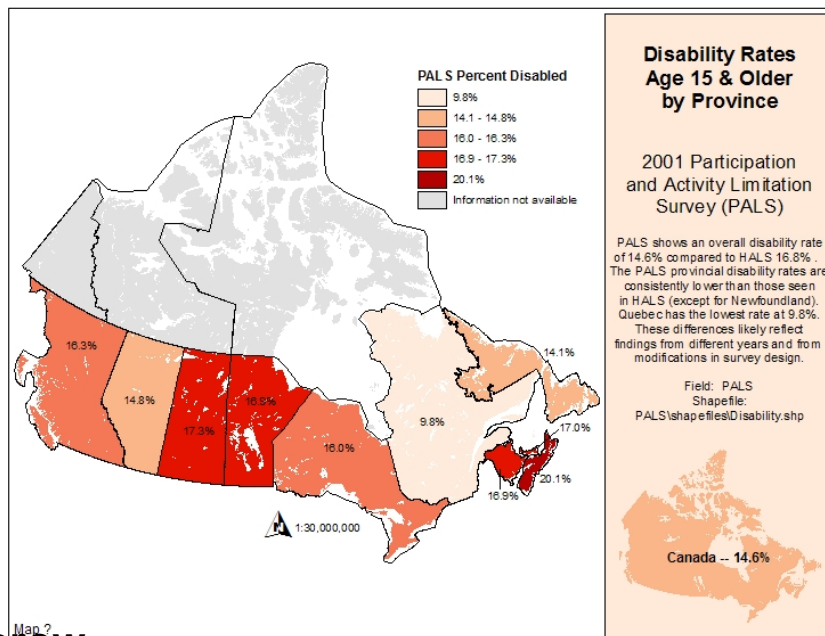
University at Buffalo The State University of New York

Cyberinfrastructure Laboratory

CI Lab

Literacy & Disability in Canada

- Exploring the relationship between illiteracy & disability across the Canadian landscape
- Social Systems GIS Lab in the Dept. of Anthropology is working with researchers from York University & the Canadian Abilities Foundation.
- Sponsored by The Adult Learning & Literacy Directorate of the Ministry of Human Resources & Social Development Canada.





E. Zubrow



Cosmological Parameter Estimation

- Wealth of new precision cosmological data
- WMAP Cosmic Microwave Background Measurement
- Sloan Digital Sky Survey: 3-D map of a million galaxies
- Interpret implications of data for models of the first trillionth of a second of the universe: *inflation*
- *Monte Carlo Markov Chain data analysis: stochastic exploration of many-dimensional parameter spaces*



W. Kinney





UB's Structural Engineering and Earthquake Simulation Laboratory (SEESL)

Structural Engineering

NEESWood:
Development of a
Performance-Based
Seismic Design for
Woodframe
Construction:



Two-story
Townhouse
on Twin
Shake Tables



2-D
Geotechnical
Laminar Box
Tests of Pile
Foundations
Subjected to
Soil
Liquefaction

M. Bruneau, A. Reinhorn, G. Lee




University at Buffalo The State University of New York

Cyberinfrastructure Laboratory

CI Lab

Understanding Combustion

- Flame-wall interaction modeling for a non-premixed flame propelled by a vortex ring.
- In this figure different time instants are shown during the interaction. White line contours and color contours represent vortex ring and flame, respectively.
- Key Features:
 - ❑ Modeling of Detailed GRI3. Mechanism for Methane Combustion
 - ❑ Parallel algorithm using mpi
 - ❑ 85-90% Parallel efficiency for up to 64 processors
- FWI study is important to determine
 - ❑ Engine Design
 - ❑ Quenching Distances
 - ❑ Flame Structure
 - ❑ Unburned hydrocarbon
 - ❑ Maximum Wall heat fluxes



C. Madnia



University at Buffalo The State University of New York

Cyberinfrastructure Laboratory

CI Lab

Miller's Cyberinfrastructure Laboratory (MCIL)



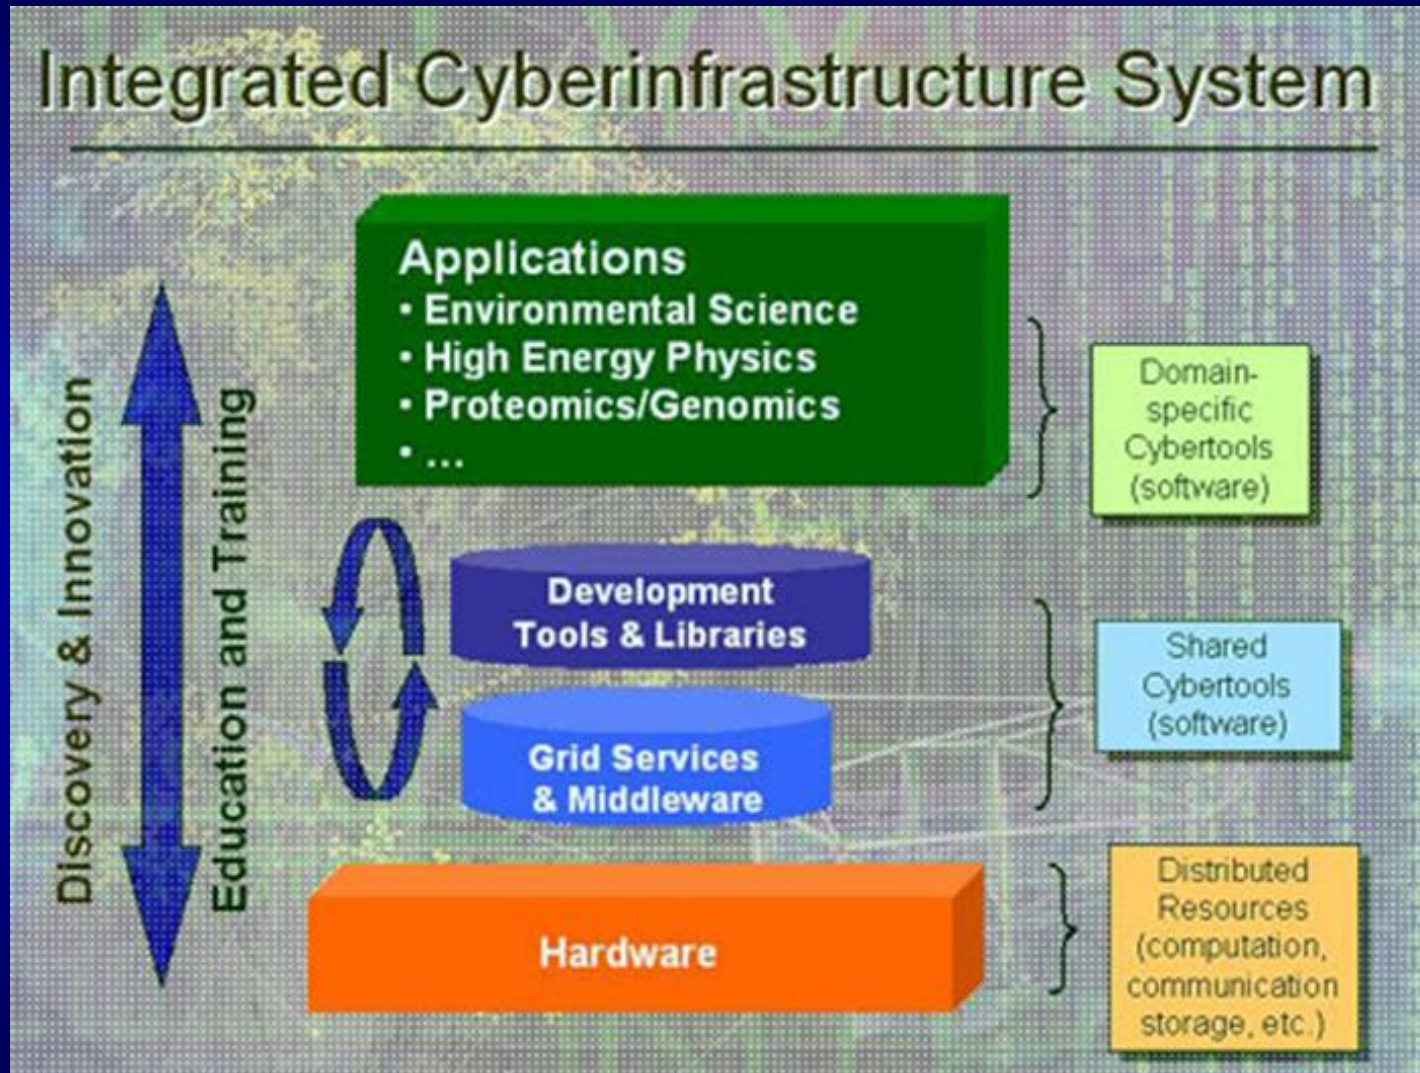
Tiled-Display Wall in CCR

MCIL Overview

- **“Cyberinfrastructure (CI) is a comprehensive phenomenon that involves creation, dissemination, preservation, and application of knowledge” (NSF)**
- **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



NSF Integrated Cyberinfrastructure



NSF Director Arden L. Bement: "leadership in cyberinfrastructure may determine America's continued ability to innovate – and thus our ability to compete successfully in the global arena."

MCIL Equipment (57.5 TF; 37 TB)

■ Production Equipment

- ❑ Magic (50+ Tops; 12,000+ cores: Largest Compute System in WNY, NYS Grid, OSG – based on GPUs)

- Dell Intel Head Node; Dell Intel Worker Nodes; 13 NVIDIA Tesla S1070s, Dell 15 TB Storage

- ❑ Dell Workstations

■ Experimental Equipment

- ❑ Clusters

- Head Nodes: Dell 1950s (Intel Dual Core Xeon)

- Workers: Dell Intel 2×4s, Intel 1×2s, & AMD 2×2s

- NVIDIA S870s

- ❑ Virtual Memory Machines (2 × Dell Intel 4×4)

- ❑ Dell GigE Managed Switches; InfiniBand Switches

- ❑ 22 TB Dell Storage (2)

- ❑ Condor Flock (35 Intel/AMD)



Evolution of MCIL 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}



MCIL Lab Projects

- **Lightweight Grid Monitor (Dashboard)**
- **Predictive Scheduler**
 - Define quality of service estimates of job completion, by better estimating job runtimes by profiling users.
- **Dynamic Resource Allocation**
 - Develop automated procedures for dynamic computational resource allocation.
- **High-Performance Grid-Enabled Data Repositories**
 - Develop automated procedures for dynamic data repository creation and deletion.
- **Integrated Data Grid**
 - Automated Data File Migration based on profiling users.
- **Grid Portal**



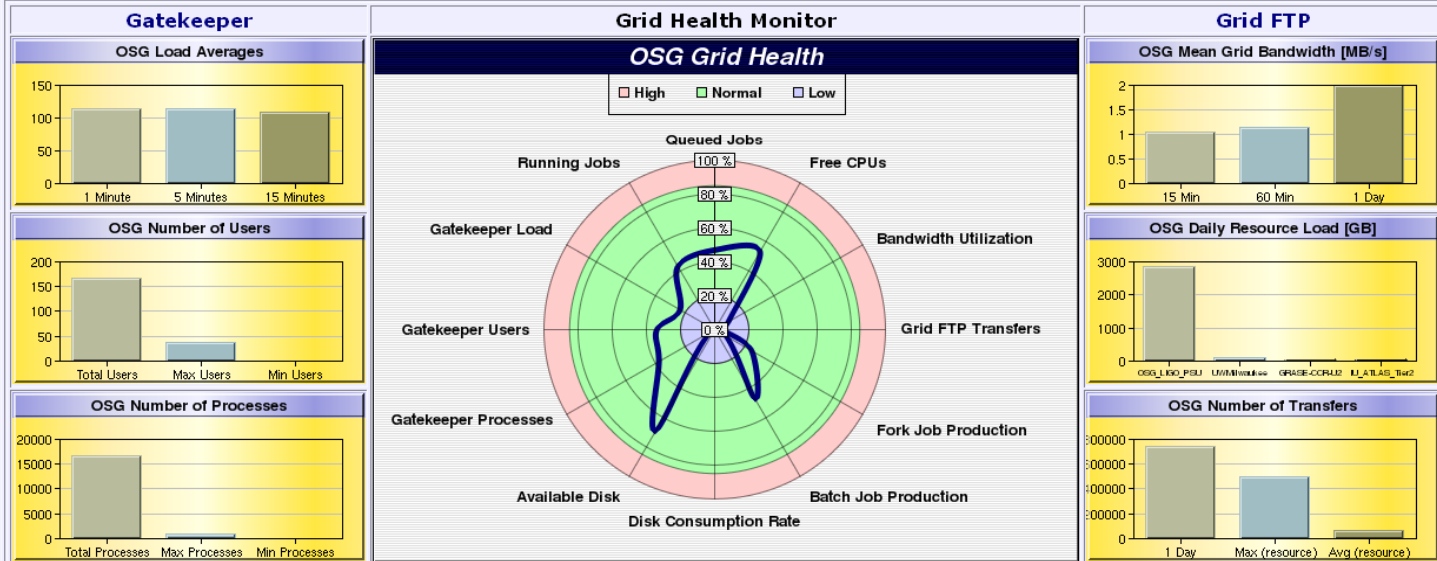



CI Lab

- CI Lab Grid Monitor
- Grid Dashboard
- Operations Dashboard
- Historical Dashboard
- Running/Queued Jobs
- Job History
- Detailed Job History
- VO Sponsor CPUs
- Free/Running/Queued CPUs
- VO Support Matrix
- Current Bandwidth Matrix
- Historical Bandwidth Matrix
- Current Latency Matrix
- Historical Latency Matrix
- Resource Queue Visualization
- Resource User Visualization
- SnB Application Demonstrator
- ACDC Grid Dashboard Site Status
- ACDC Grid Dashboard Tutorial
- GRASE VO
- Overview
- Request Membership
- Request Help
- Staff Only
- Contact Us / Staff
- CI Lab

CI LAB GRID DASHBOARD

OSG-ITB OSG ACDC TeraGrid Ad-Hoc



Supported by the National Science Foundation and the Department of Energy

CI Lab Grid Monitor

Molecular Structure Determination via *Shake-and-Bake*

■ *SnB* Software by UB/HWI

- ❑ IEEE “Top Algorithms of the Century”

■ Worldwide Utilization

■ Critical Step




- ❑ Rational Drug Design
- ❑ Structural Biology
- ❑ Systems Biology

■ Vancomycin

- ❑ “Antibiotic of Last Resort”

■ Current Efforts

- ❑ Grid
- ❑ Collaboratory
- ❑ Intelligent Learning



1. Isolate a single crystal
2. Perform the X-Ray diffraction experiment
3. Determine the crystal structure



Acknowledgments

- Mark Green
 - Cathy Ruby
 - Amin Ghadersohi
 - Naimesh Shah
 - Steve Gallo
 - Jason Rappleye
 - Jon Bednasz
 - Sam Guercio
 - Martins Innus
 - Cynthia Cornelius

 - George DeTitta
 - Herb Hauptman
 - Charles Weeks
 - Steve Potter
- Alan Rabideau
 - Igor Janckovic
 - Michael Sheridan
 - Abani Patra
 - Matt Jones

 - NSF ITR
 - NSF CRI
 - NSF MRI
 - NYS
 - CCR



