The Center for Computational Research

Russ Miller

Director, Center for Computational Research

UB Distinguished Professor, Computer Science & Engineering

Senior Research Scientist, Hauptman-Woodward Medical Inst

University at Buffalo

The State University of New York

Gordon E. Moore

- Co-Founder of Intel
- Predicted (1965/75) that transistor density would double every 12/18 months
- Processing speed doubling every 18 mos.
- Disk storage doubling every 12 mos.
- Aggregate bandwidth doubling every 9 mos.

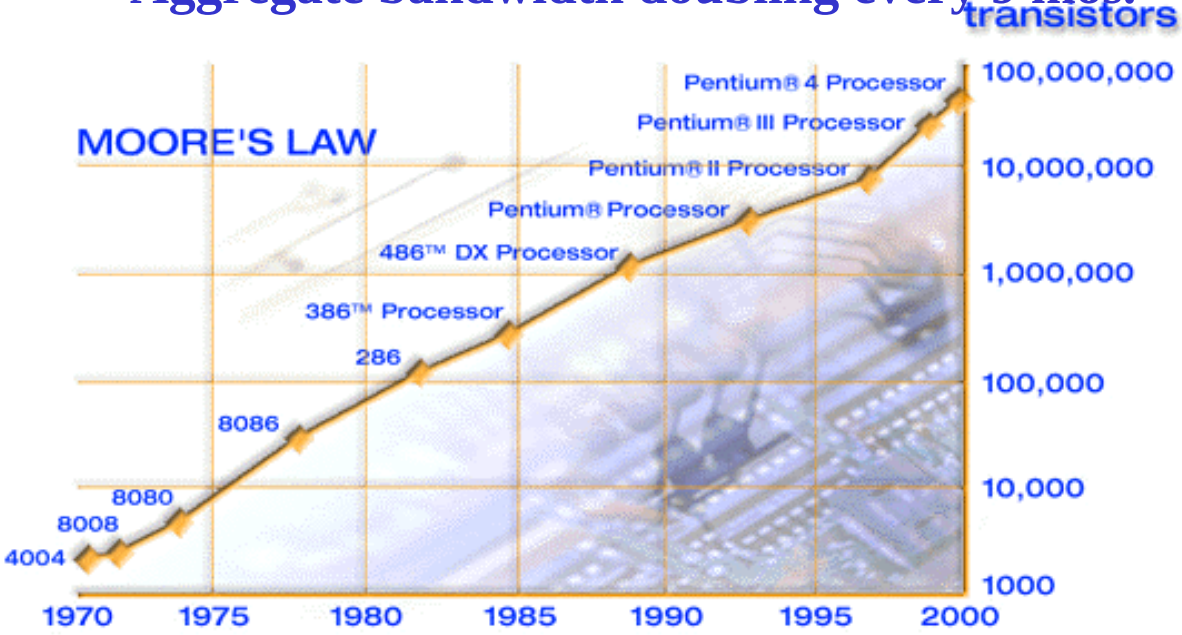

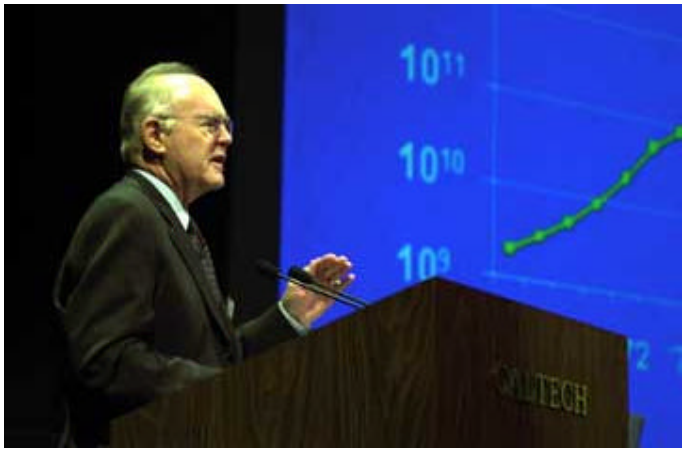

Gordon E. Moore

- A computation that took 1 year to run on a PC in 1985 would only take 5 mins to run on a PC today!
- A computation that runs in 2 hours on a PC today would have taken 24 years to run on a PC in 1985!

Supercomputers

- **Fastest computers at any point in time**
- **Used to solve large and complex problems**
- **Machines 1000 times faster than a PC**
- **Machines 10 times slower than what you need to solve the most challenging problems**

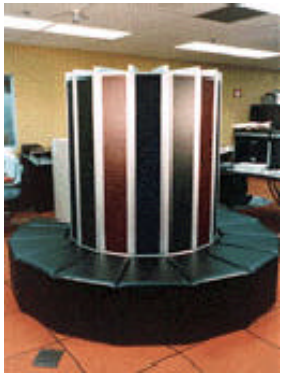

Cray1 - 1976

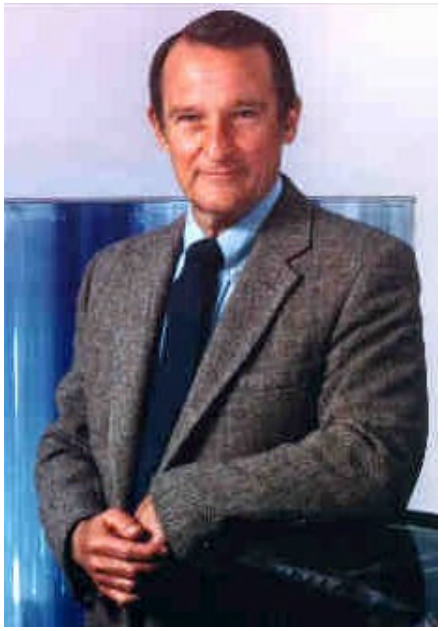

"Seymour Cray is the Thomas Edison of the supercomputing industry"
- Larry L. Smarr

Seymour Cray
1925-1996

Beowulf Clusters

- **Industry Standard Hardware and Software**
  - **PC-Based Components (Intel or AMD)**
  - **Ethernet or Myrinet**
  - **Linux, PBS, MPI**
  - **"Commodity Off-The-Shelf" (COTS)**
- **Operates as a Single System**
- **Rivals Performance of Traditional Supercomputer at a Fraction of the Price**

Thomas Sterling
Caltech

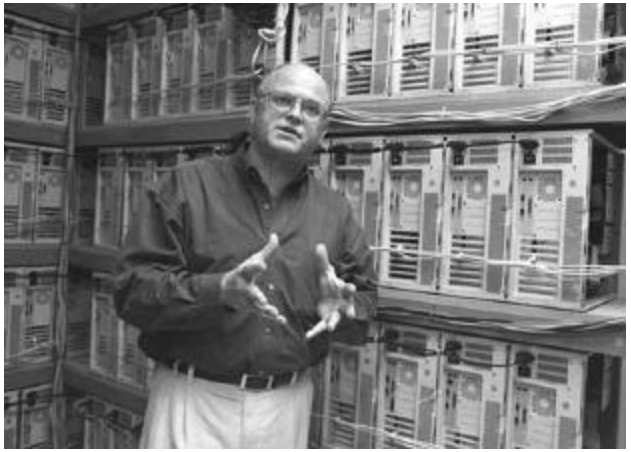

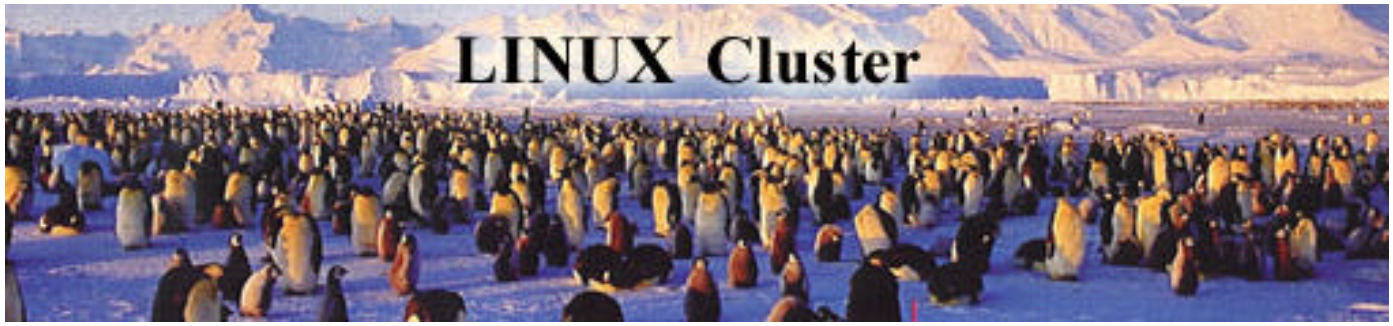

Fastest Computers

Year	Mach	Procs	GFlops
1976	Cray 1	1	0.1
1982	Cray X-MP	4	0.9
1986	Cray 2	4	2
1989	Cray Y-MP	8	2.7
1989	TMC CM-2	8192	28
1992	TMC CM-5	1024	131
A 1-year calc in 1980 = 5.4 sec today A 1990 HPC = a laptop today			

Year	Mach	Procs	GFlops
1993	Cray T3D	1024	152
1994	Fujitsu VPP	140	236
1996	Hitachi SR2	2048	368
1997	Intel ASCI-R	9152	1830
1999	SGI ASCI-BM	6144	3072
2000	IBM ASCI-W	8192	12,288
2002	NEC E.S.	5120	40,960

Growth of Peak Performance

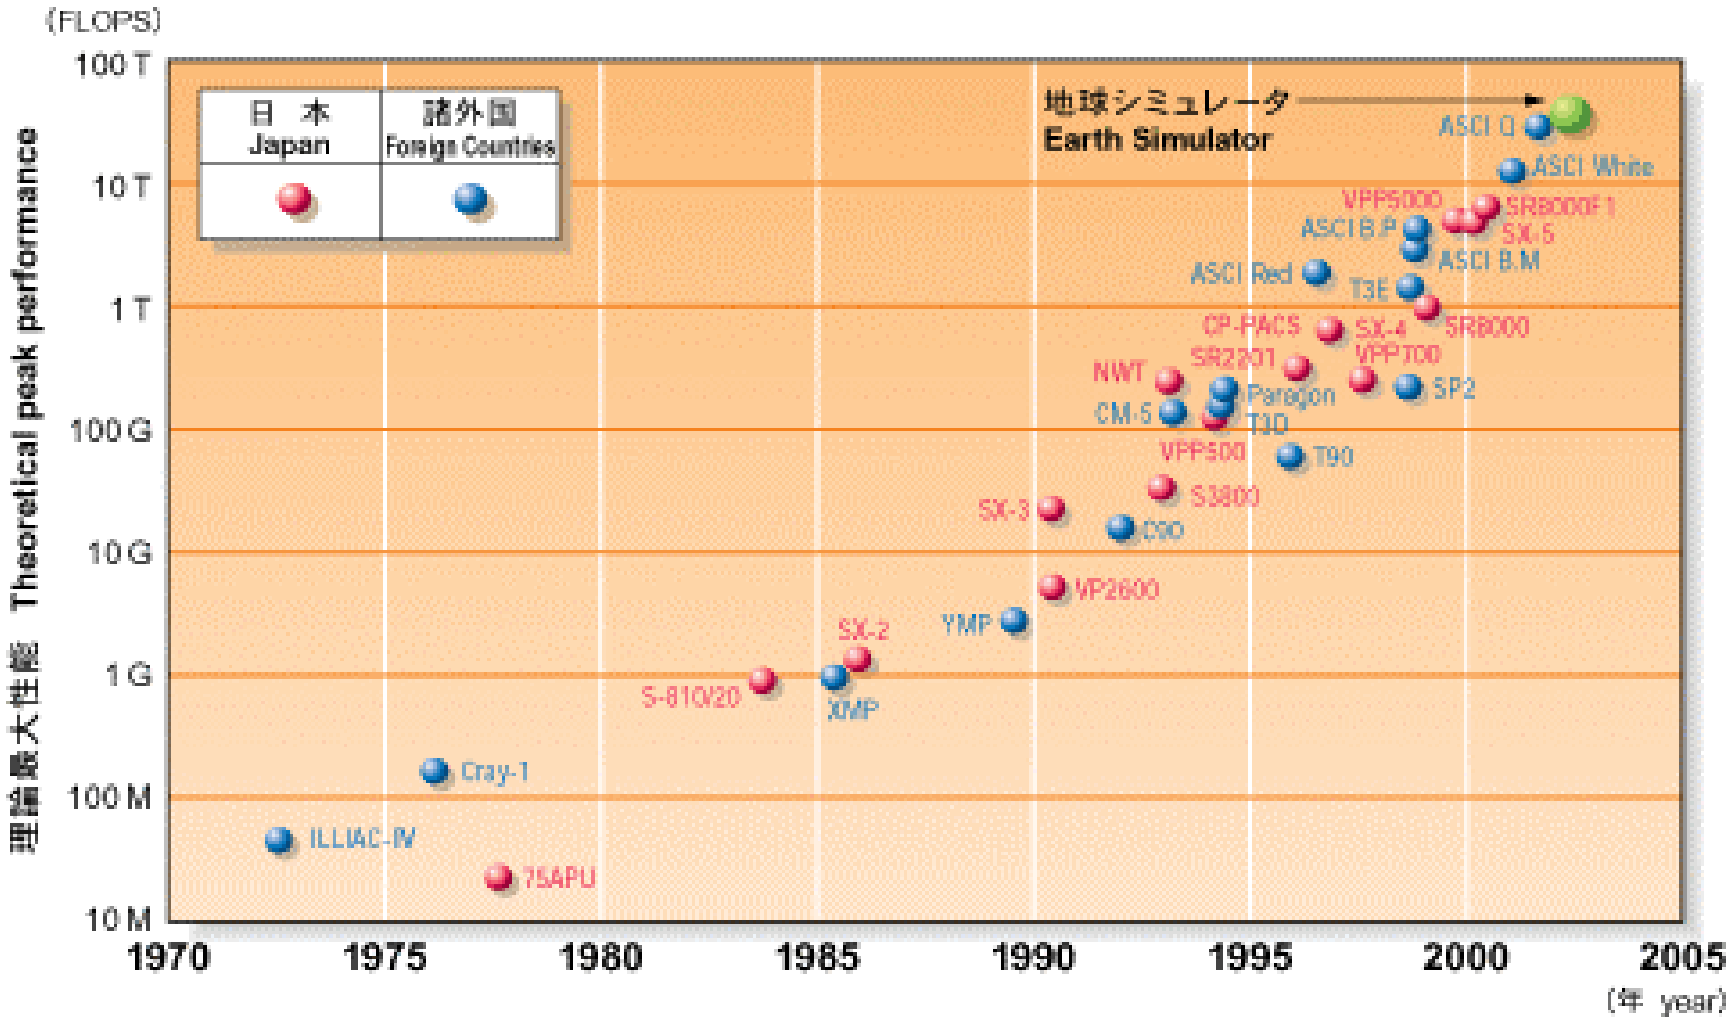

Earth Simulator

- 40TFlops Peak
- Homogeneous, Centralized,

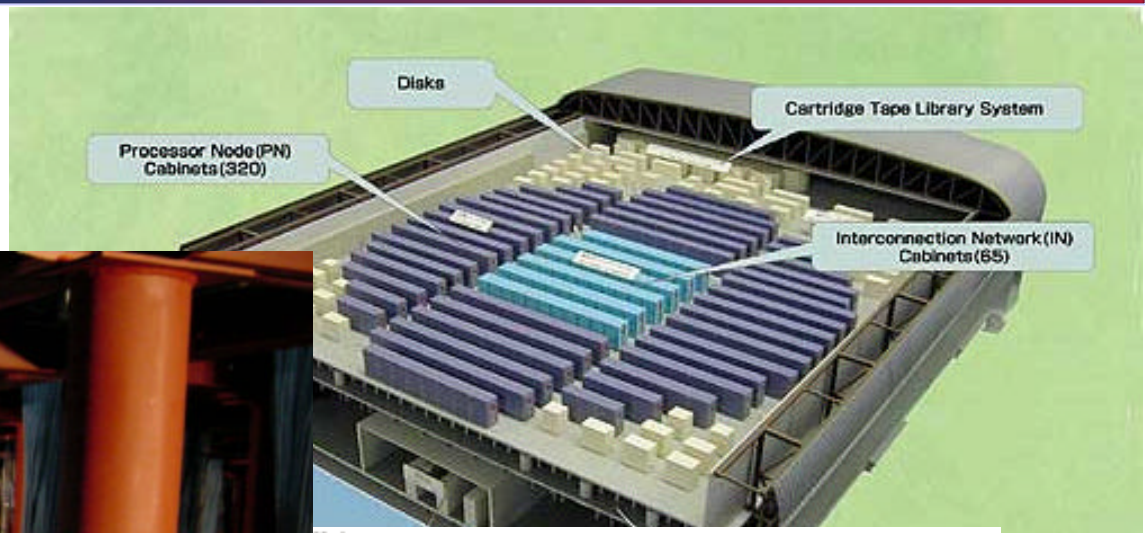

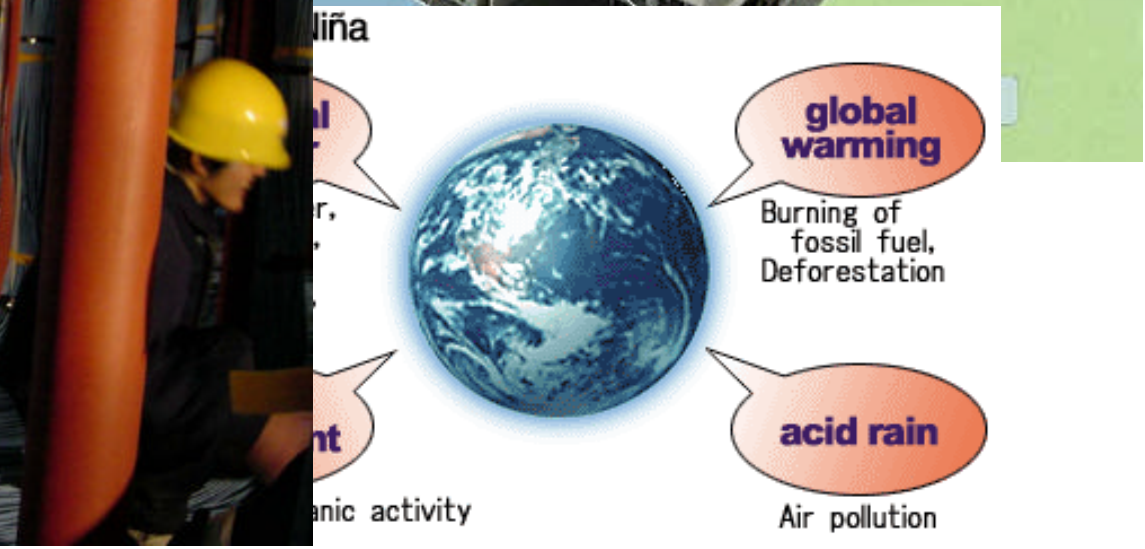

Center for Computational Research 1999-2003 Snapshot

- **High-Performance Computing and High-End Visualization**
  - 110 Research Groups in 27 Depts
  - 13 Local Companies
  - 10 Local Institutions
- **External Funding**
  - $111M External Funding
    - $13.5M as lead
    - $97.5M in support
  - $41.8M Vendor Donations
- **Deliverables**
  - 350+ Publications
  - Software, Media, Algorithms, Consulting, Training, CPU Cycles…

Major CCR Resources

- Dell Linux Cluster: #22 ® #25 ® #38
  - 600 P4 Processors (2.4 GHz)
  - 600 GB RAM; 40 TB Disk; Myrinet
- Dell Linux Cluster: #187 ® #368 ® off
  - 4036 Processors (PIII 1.2 GHz)
  - 2TB RAM; 160TB Disk; 16TB SN
  - Restricted Use (Skolnick)

- SGI Origin3800
  - 64 Processors (400 MHz)
  - 32 GB RAM; 400 GB Disk
- IBM RS/6000 SP
  - 78 Processors
  - 26 GB RAM; 640 GB Disk
- Sun Microsystems Cluster
  - 48 Sun Ultra 5s (333MHz)
  - 16 Dual Sunblades (750MHz)
  - 30 GB RAM, Myrinet
- SGI Intel Linux Cluster
  - 150 PIII Processors (1 GHz)
  - 75 GB RAM, 2.5 TB Disk Storage
- Apex Bioinformatics System
  - Sun V880 (3), 6800, 280R (2), PIIIs
  - Sun 3960: 7 TB Disk Storage
- HP/Compaq SAN (4Q03)
  - 75 TB Disk; 200 TB Tape

Advanced CCR Data Center (ACDC) Computational Grid Overview

Nash: Compute Cluster

75 Dual Processor
1 GHz Pentium III
RedHat Linux 7.3
1.8 TB Scratch Space

Joplin: Compute Cluster

300 Dual Processor
2.4 GHz Intel Xeon
RedHat Linux 7.3
38.7 TB Scratch Space

Mama: Compute Cluster

9 Dual Processor
1 GHz Pentium III
RedHat Linux 7.3
315 GB Scratch Space

ACDC: Grid Portal

4 Processor Dell 6650
1.6 GHz Intel Xeon
RedHat Linux 9.0
66 GB Scratch Space

Young: Compute Cluster

16 Dual Sun Blades
47 Sun Ultra5
Solaris 8
770 GB Scratch Space

Crosby: Compute Cluster

SGI Origin 3800
64 - 400 MHz IP35
IRIX 6.5.14m
360 GB Scratch Space

Fogerty: Condor Flock Master

1 Dual Processor
250 MHz IP30
IRIX 6.5

Expanding

RedHat, IRIX, Solaris, WINNT, etc

CCR

19 IRIX, RedHat, & WINNT Processors

T1 Connection

Computer Science & Engineering

25 Single Processor Sun Ultra5s

School of Dental Medicine

9 Single Processor Dell P4 Desktops

Hauptman-Woodward Institute

13 Various SGI IRIX Processors

Note: Network connections are 100 Mbps unless otherwise noted.

Network Connections

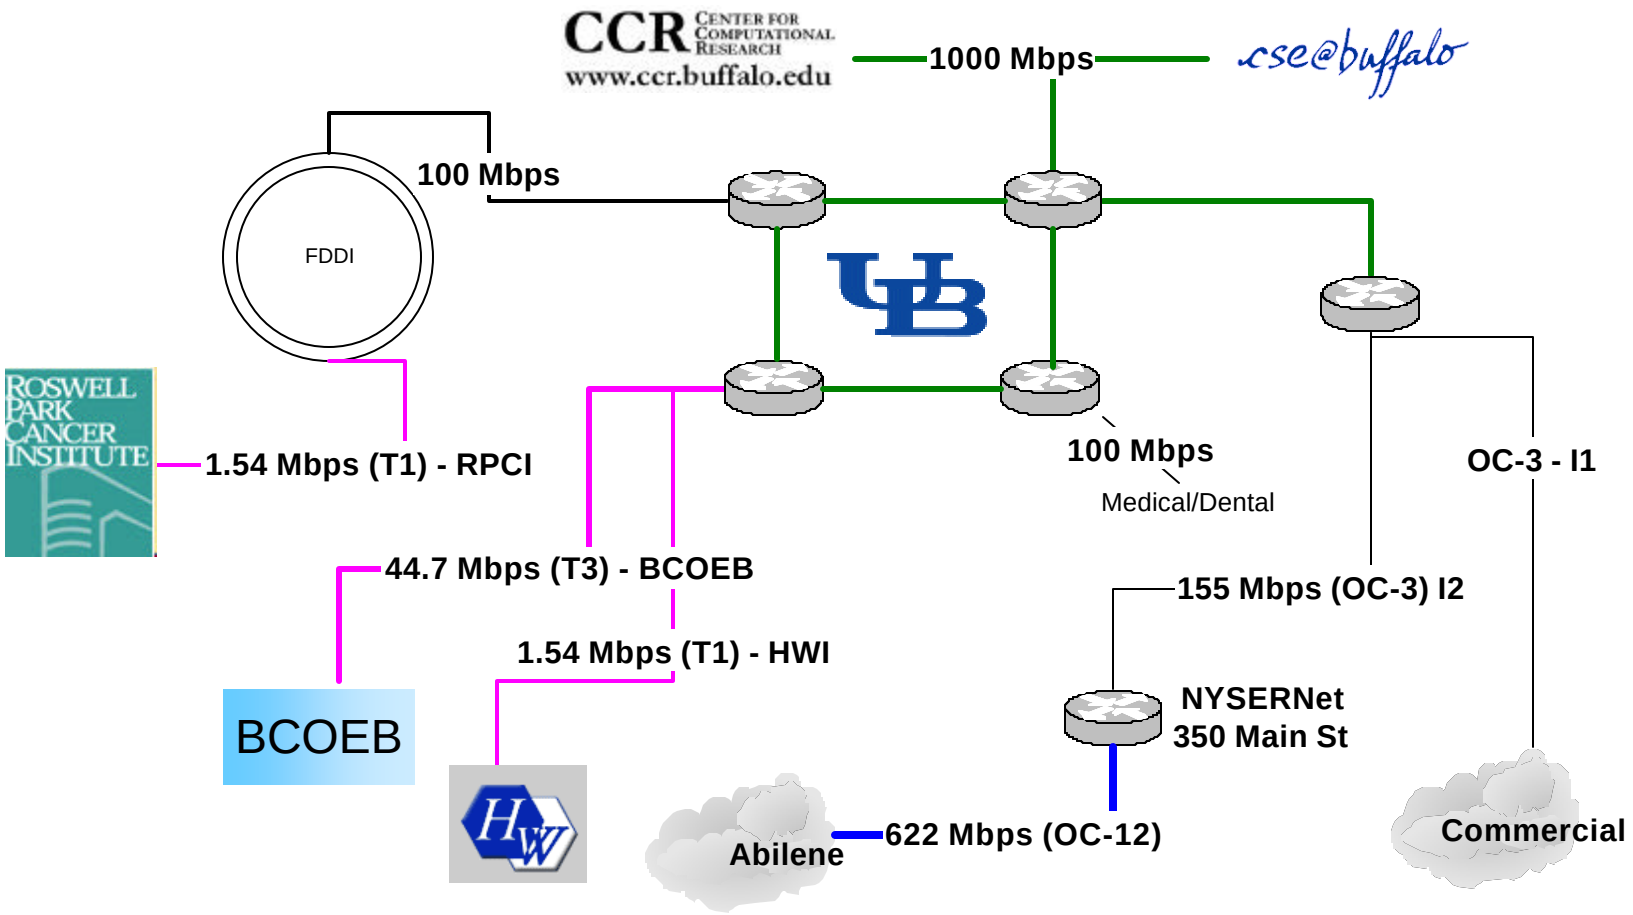

X-Ray Crystallography

- **Objective: Provide a 3-D mapping of the atoms in a crystal.**

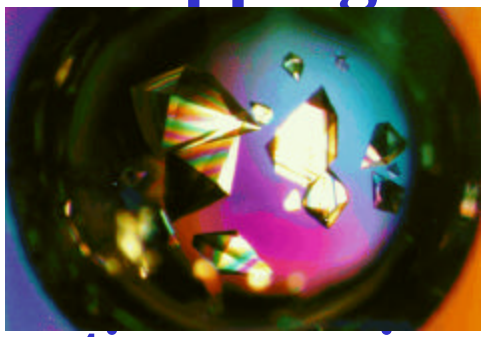

- **Procedure:**
  1. **Isolate a single crystal.**
  2. **Perform the X-Ray diffraction experiment.**

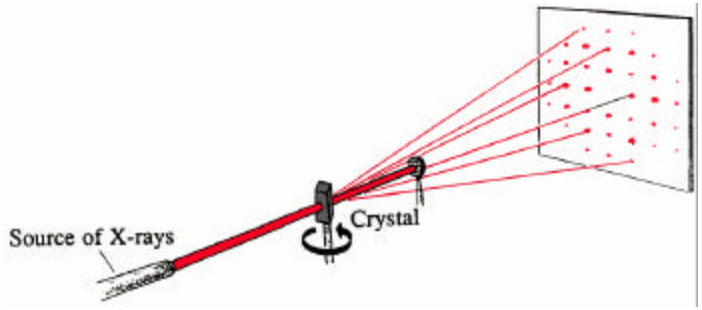

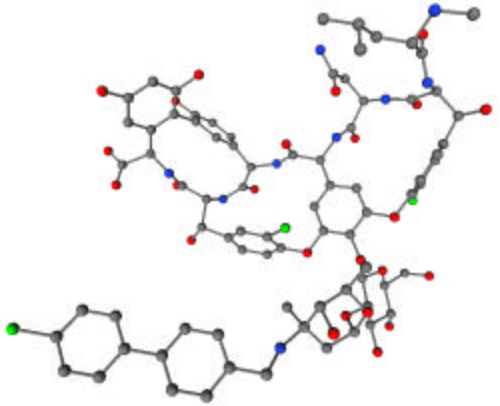

  3. **Determine molecular structure that agrees with diffration data.**

Shake-and-Bake Method: Dual-Space Refinement

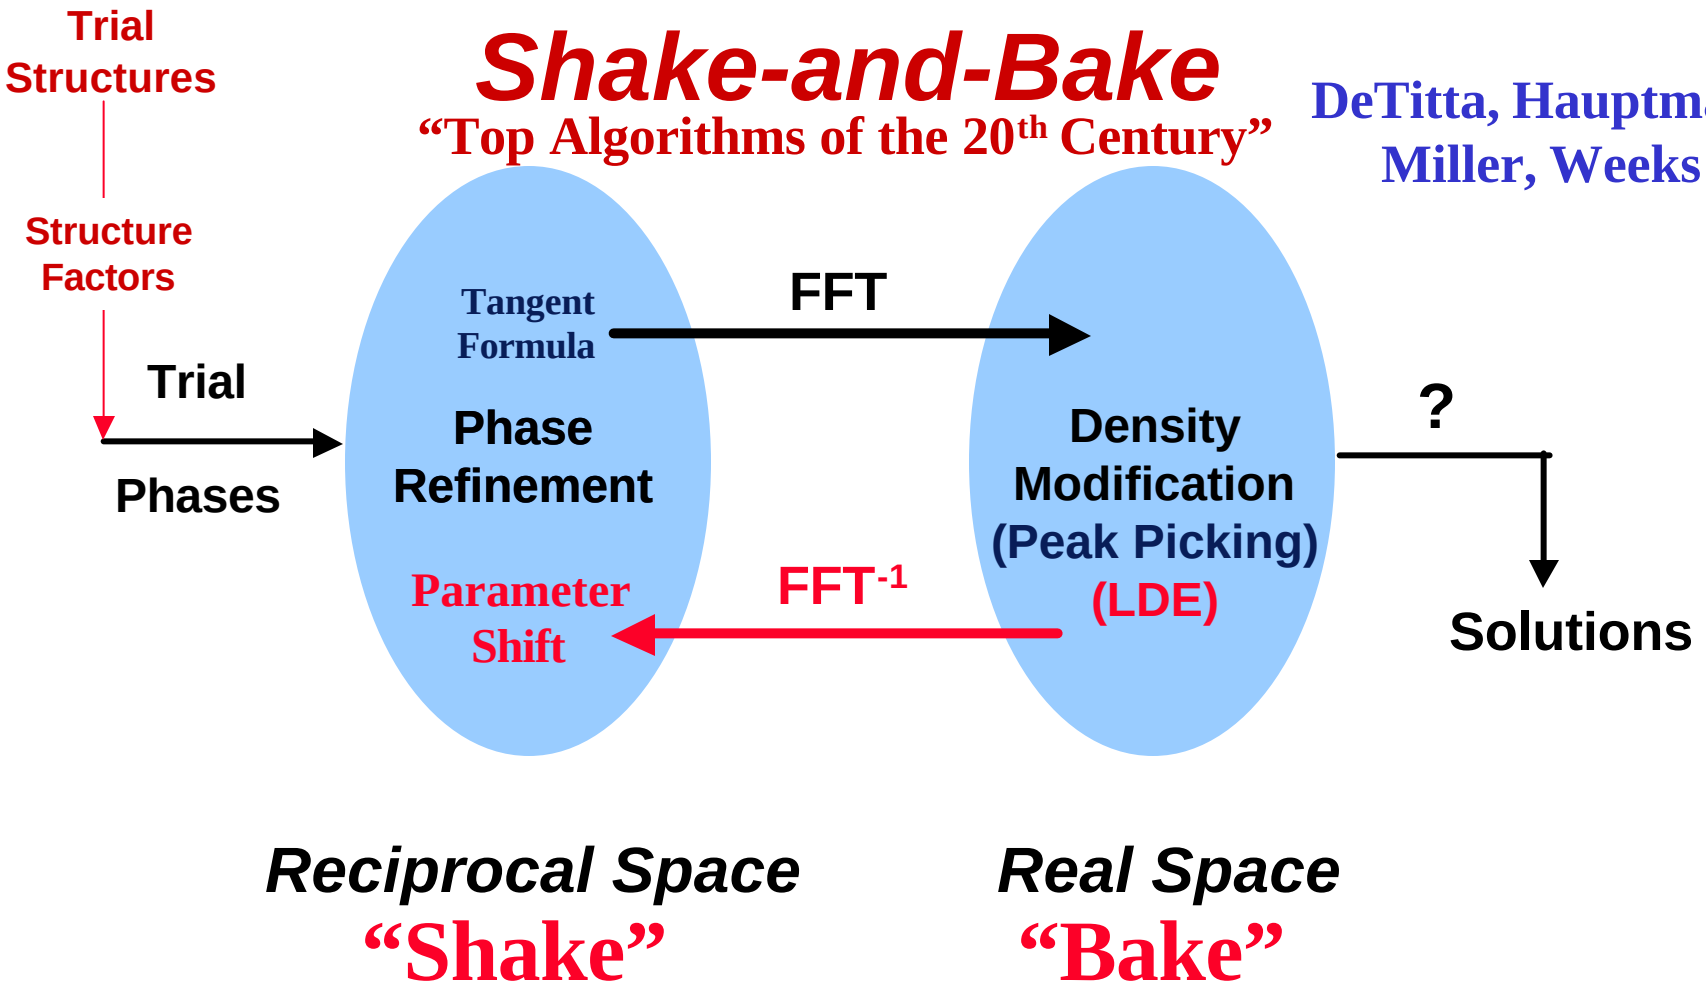

Phasing and Structure Size

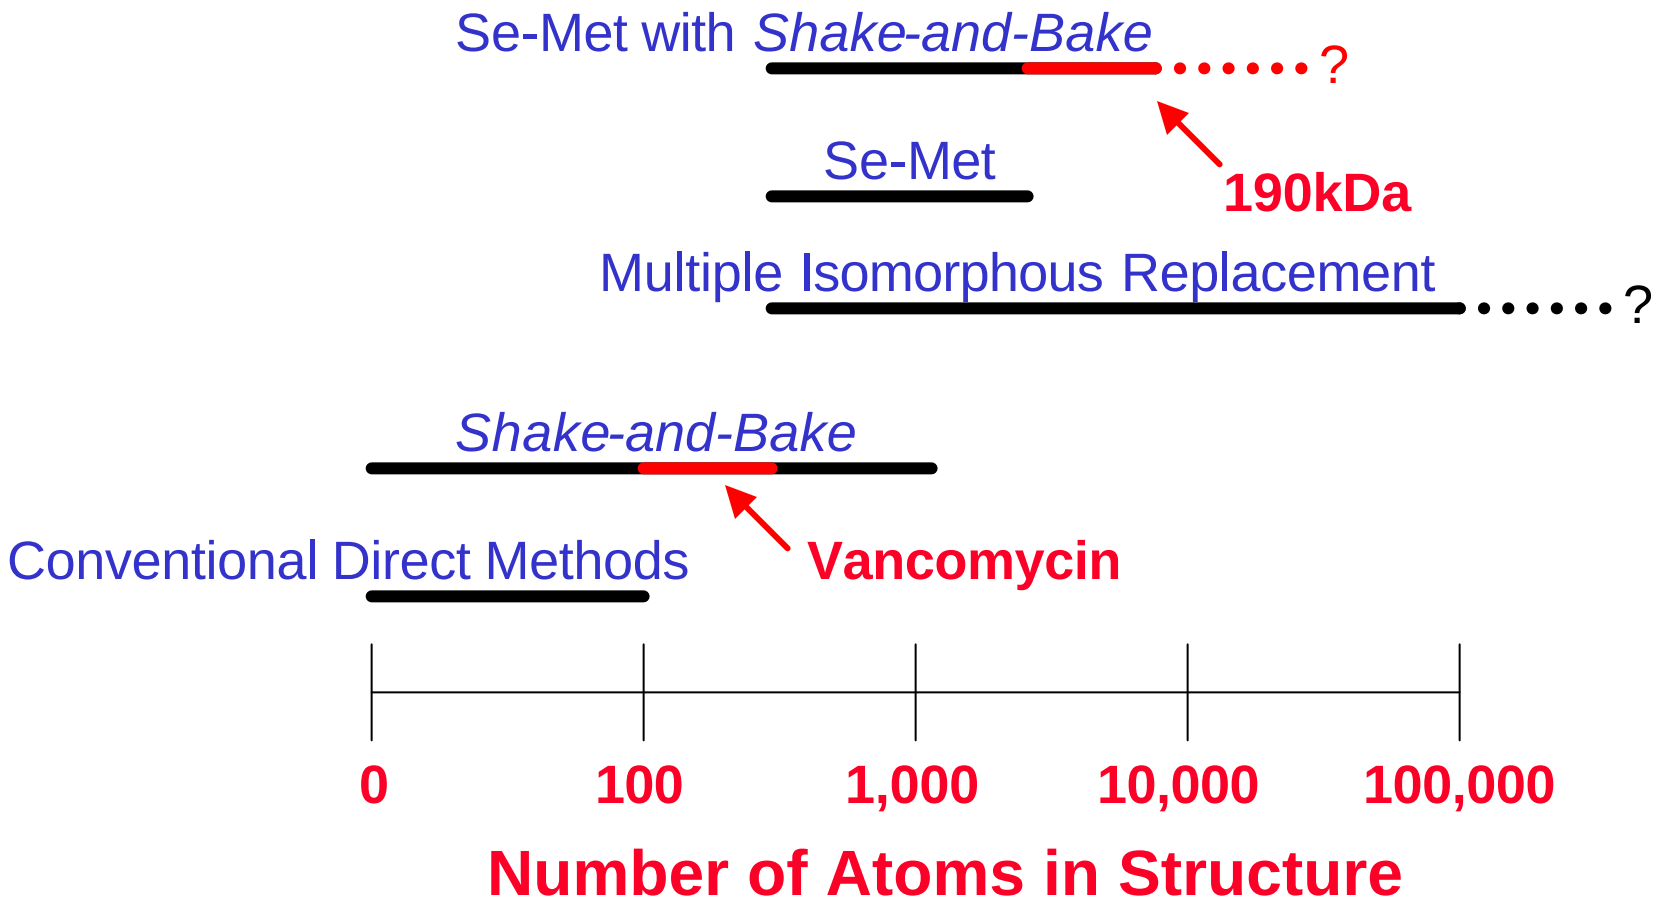

Grid-Based *SnB* Objectives

- **Install Grid-Enabled Version of *SnB***
- **Job Submission and Monitoring over Internet**
- ***SnB* Output Stored in Database**
- ***SnB* Output Mined through Internet-Based Integrated Querying Tool**

- **Serve as Template for Chem-Grid & Bio-Grid**
- **Experience with Globus and Related Tools**

Grid Services and Applications

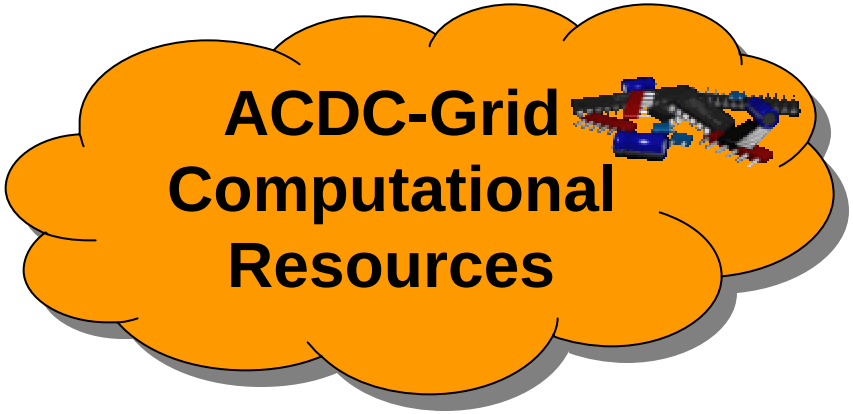

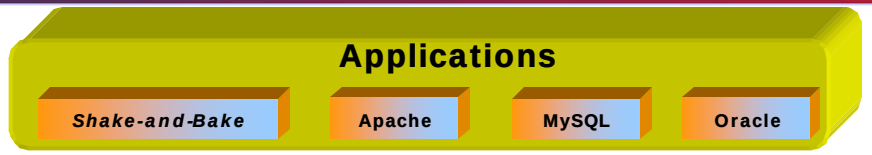

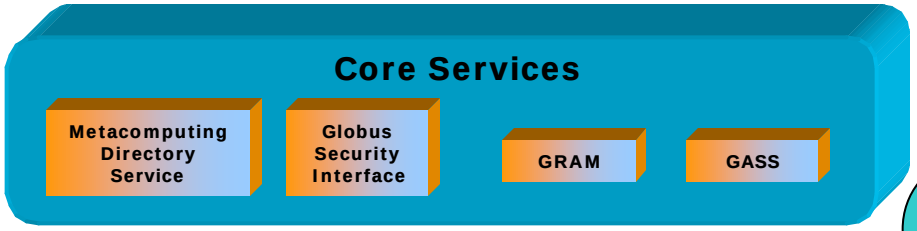

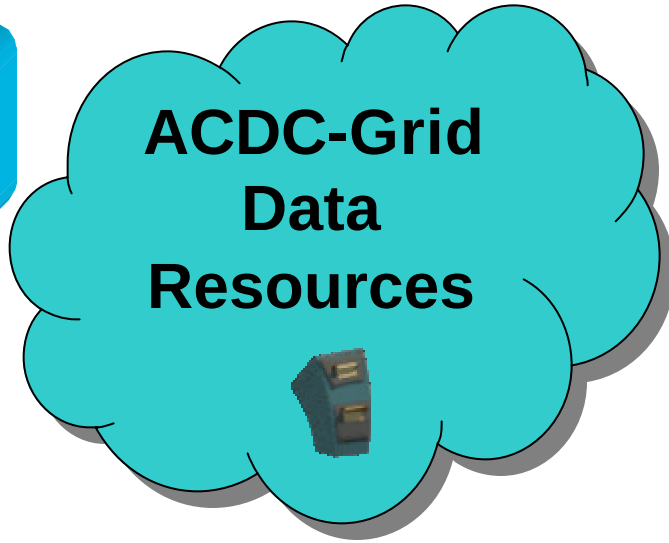

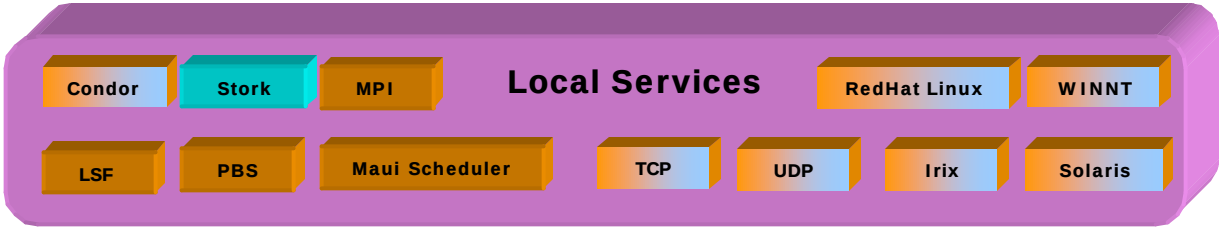

Adapted from Ian Foster and Carl Kesselman

ACDC-Grid Portal

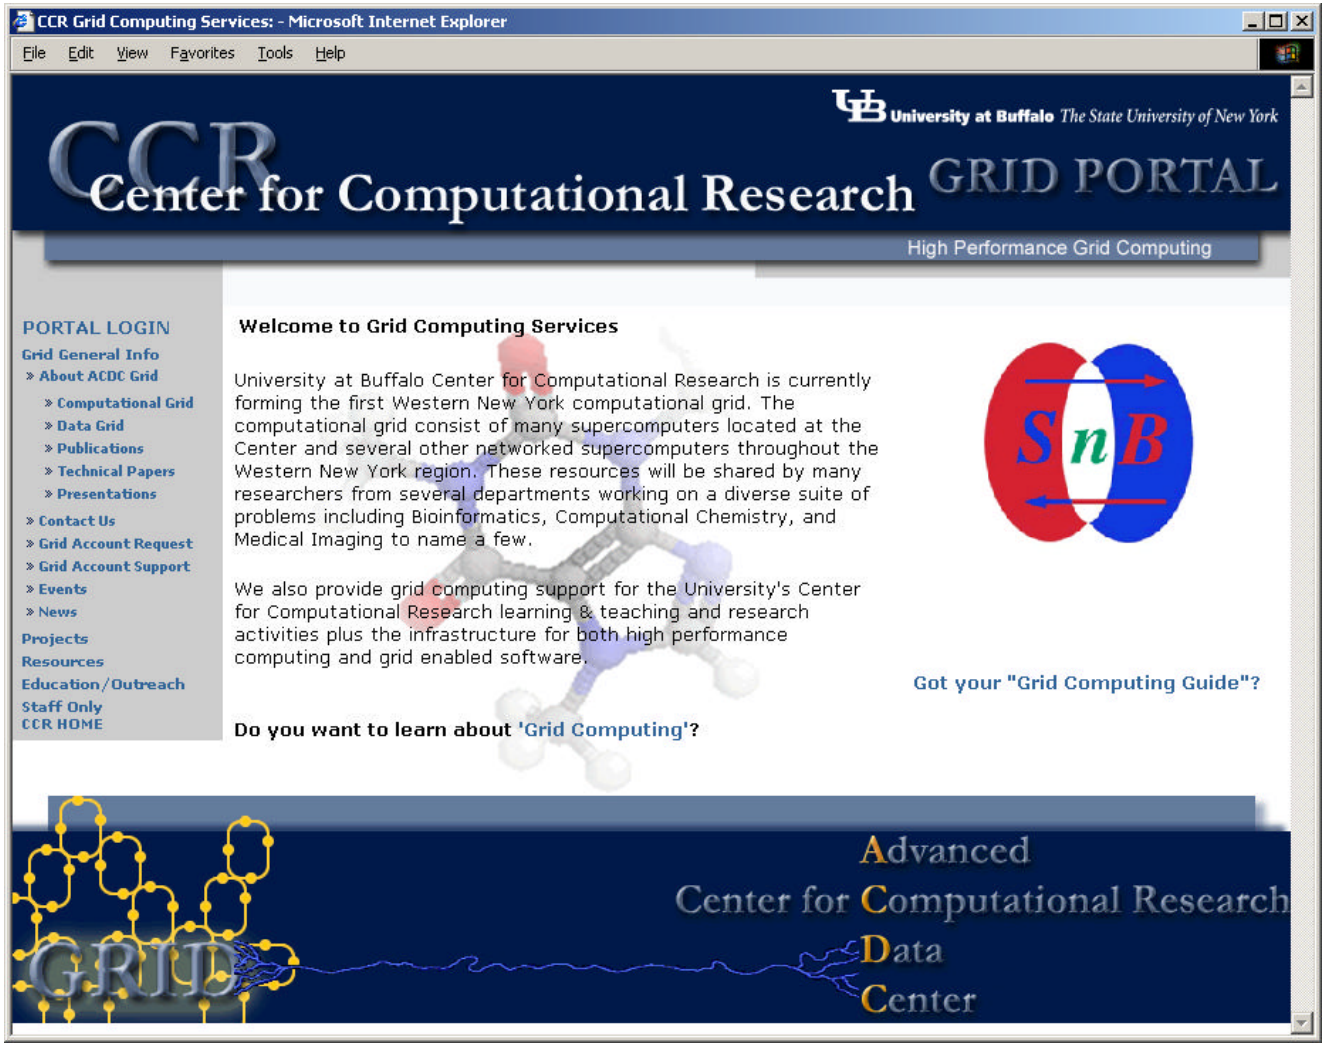

Data Grid Capabilities

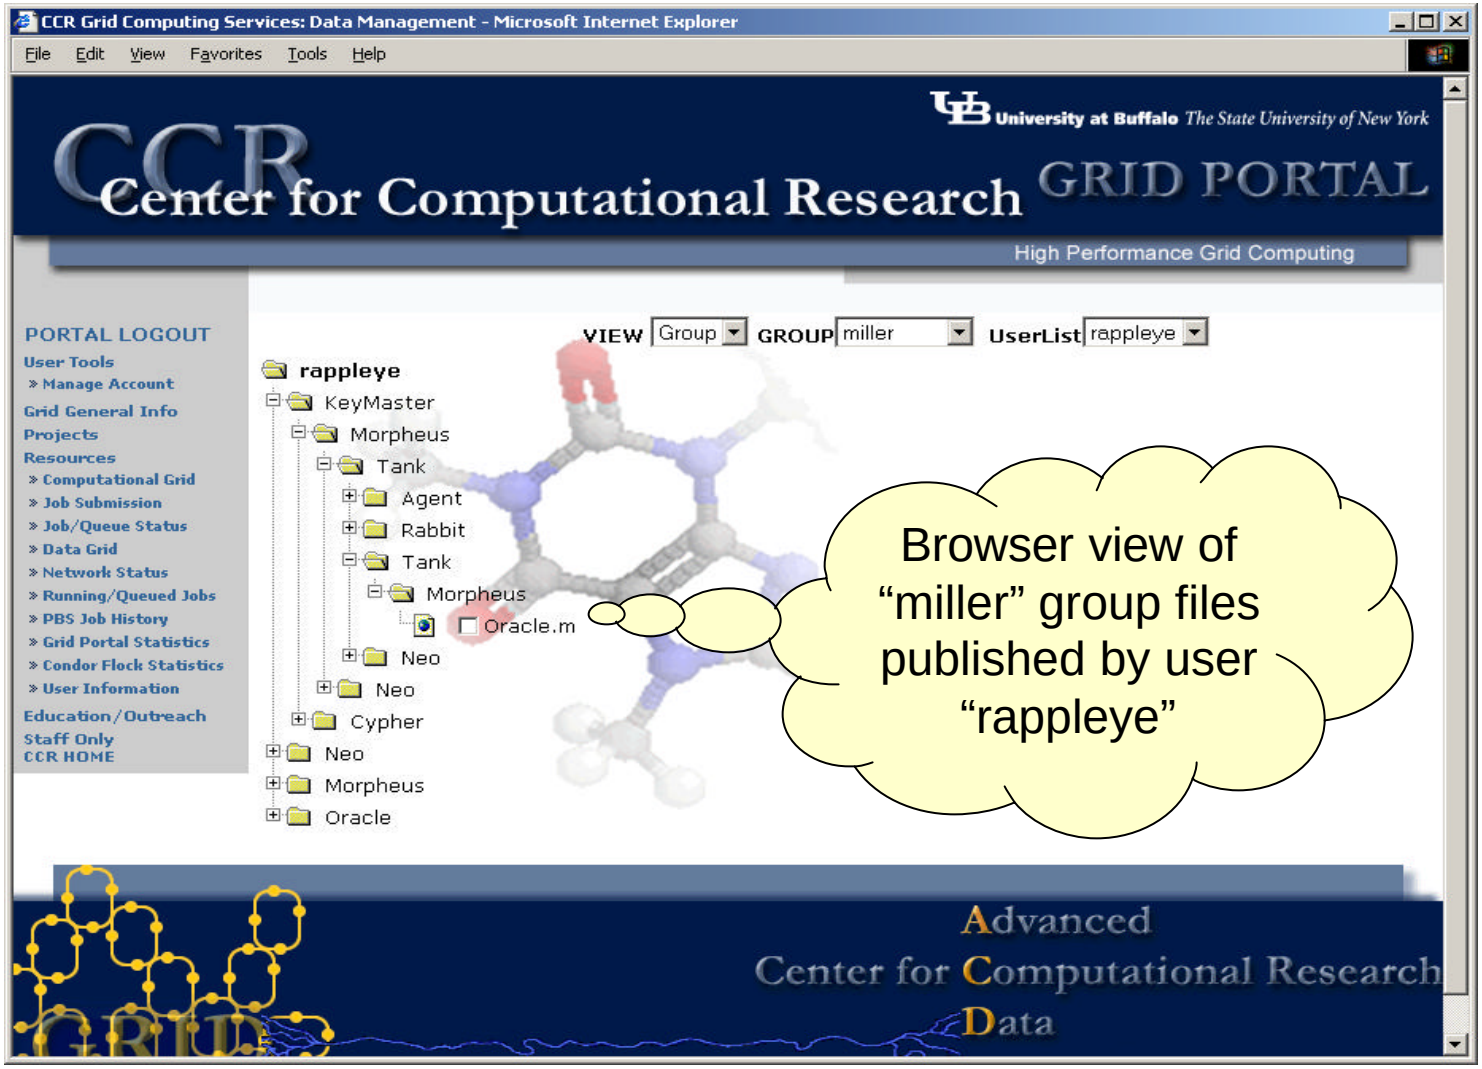

Grid Portal Job Status

- **Grid-enabled jobs can be monitored using the Grid Portal web interface dynamically.**
  - **Charts are based on:**
    - **total CPU hours, or**
    - **total jobs, or**
    - **total runtime.**
  - **Usage data for:**
    - **running jobs, or**
    - **queued jobs.**
  - **Individual or all resources.**
  - **Grouped by:**
    - **group, or**
    - **user, or**
    - **queue.**

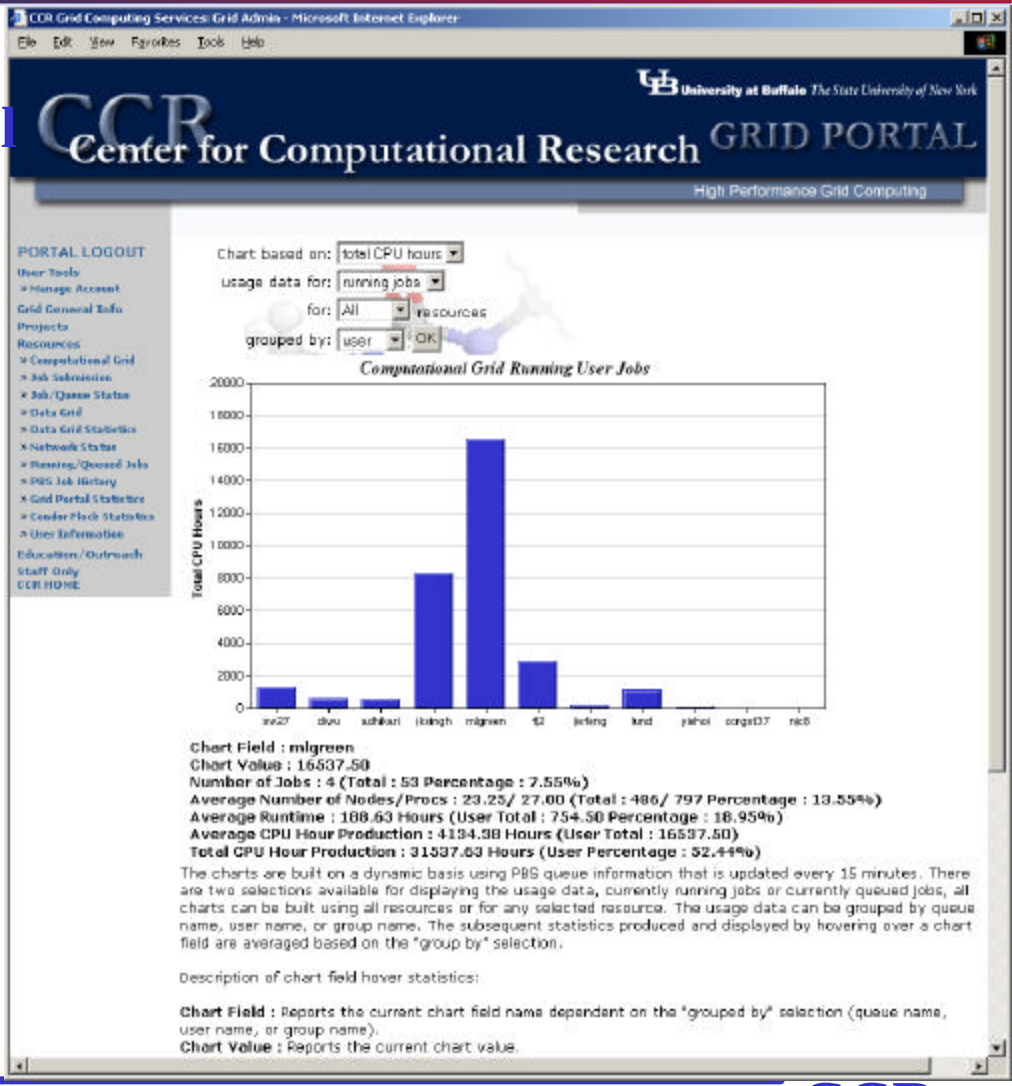

Grid Portal Job Status

CCR Grid Computing Services: Grid Admin - Microsoft Internet Explorer

File Edit View Favorites Tools Help

University at Buffalo The State University of New York

CCR Center for Computational Research GRID PORTAL

High Performance Grid Computing

PORTAL LOGOUT

User Tools
- Manage Account

Grid General Info

Projects

Resources
- Computational Grid
- Job Submission
- Job/Queue Status
- Data Grid
- Data Grid Statistics
- Network Status
- Running/Queued Jobs
- PBS Job History
- Grid Portal Statistics
- Condor Flock Statistics
- User Information

Education/Outreach

Staff Only

CCR HOME

Chart based on: total CPU hours

usage data for: queued jobs

for: All resources

grouped by: user OK

Computational Grid Queued User Jobs

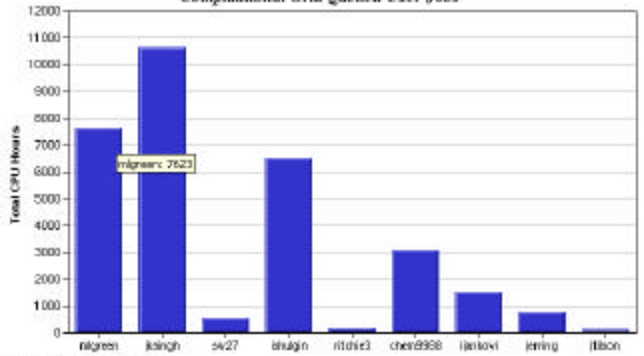

Chart Field : mlgreen
Chart Value : 7623.17
Number of Jobs : 27 (Total : 60 Percentage : 45.00%)
Average Number of Nodes/Procs : 1.96/ 3.93 (Total : 1641/ 3206 Percentage : 3.31%)
Average Runtime : 71.92 Hours (User Total : 1941.75 Percentage : 71.90%)
Average CPU Hour Production : 282.34 Hours (User Total : 7623.17)
Total CPU Hour Production : 31015.17 Hours (User Percentage : 24.58%)

The charts are built on a dynamic basis using PBS queue information that is updated every 15 minut
are two selections available for displaying the usage data, currently running jobs or currently queue
charts can be built using all resources or for any selected resource. The usage data can be groupe
name, user name, or group name. The subsequent statistics produced and displayed by hovering ov
field are averaged based on the "group by" selection.

Description of chart field hover statistics:

Chart Field : Reports the current chart field name dependant on the "grouped by" selection (queue
user name, or group name).
Chart Value : Reports the current chart value.

CCR Grid Computing Services: Grid Admin - Microsoft Internet Explorer

File Edit View Favorites Tools Help

University at Buffalo The State University of New York

CCR Center for Computational Research GRID PORTAL

High Performance Grid Computing

PORTAL LOGOUT

User Tools
- Manage Account

Grid General Info

Projects

Resources
- Computational Grid
- Job Submission
- Job/Queue Status
- Data Grid
- Data Grid Statistics
- Network Status
- Running/Queued Jobs
- PBS Job History
- Grid Portal Statistics
- Condor Flock Statistics
- User Information

Education/Outreach

Staff Only

CCR HOME

Chart based on: total CPU hours
usage data for: queued jobs
for: all resources
grouped by: user mlgreen

Job_Num	Job_Id	Username	Groupname	Queue	Node_Cnt	Proc_Cnt	Rank	Request_Time(hr)	Start_Time(hr)
1	23364	mlgreen	miler	grid	1	2	1	71.916666666667	queued
2	23365	mlgreen	miler	grid	1	2	2	71.916666666667	queued
3	23366	mlgreen	miler	grid	2	4	3	71.916666666667	queued
4	23367	mlgreen	miler	grid	2	4	4	71.916666666667	queued
5	23368	mlgreen	miler	grid	2	4	5	71.916666666667	queued
6	23369	mlgreen	miler	grid	3	6	6	71.916666666667	queued
7	23370	mlgreen	miler	grid	2	4	7	71.916666666667	queued
8	23371	mlgreen	miler	grid	2	4	8	71.916666666667	queued
9	23372	mlgreen	miler	grid	2	4	9	71.916666666667	queued
10	23373	mlgreen	miler	grid	2	4	10	71.916666666667	queued
11	23374	mlgreen	miler	grid	2	4	11	71.916666666667	queued
12	23375	mlgreen	miler	grid	2	4	12	71.916666666667	queued
13	23376	mlgreen	miler	grid	2	4	13	71.916666666667	queued
14	23377	mlgreen	miler	grid	2	4	14	71.916666666667	queued
15	23378	mlgreen	miler	grid	2	4	15	71.916666666667	queued
16	23379	mlgreen	miler	grid	2	4	16	71.916666666667	queued
17	23380	mlgreen	miler	grid	2	4	17	71.916666666667	queued
18	23381	mlgreen	miler	grid	2	4	18	71.916666666667	queued
19	23382	mlgreen	miler	grid	2	4	19	71.916666666667	queued
20	23383	mlgreen	miler	grid	2	4	20	71.916666666667	queued
21	23384	mlgreen	miler	grid	2	4	21	71.916666666667	queued
22	23385	mlgreen	miler	grid	2	4	22	71.916666666667	queued
23	23386	mlgreen	miler	grid	2	4	23	71.916666666667	queued
24	23387	mlgreen	miler	grid	2	4	24	71.916666666667	queued
25	23388	mlgreen	miler	grid	2	4	25	71.916666666667	queued
26	23389	mlgreen	miler	grid	2	4	26	71.916666666667	queued
27	23393	mlgreen	miler	grid	2	4	27	71.916666666667	queued

Description of table fields:

Job_Num : Sequential counter for row number.
Job_Id : PBS job identification number corresponding to local PBS queue job number.
Username : PBS job owner local username.
Groupname : PBS job owner local primary group name.

ACDC-Grid Portal Condor Flock

- **CondorView integrated into ACDC-Grid Portal**

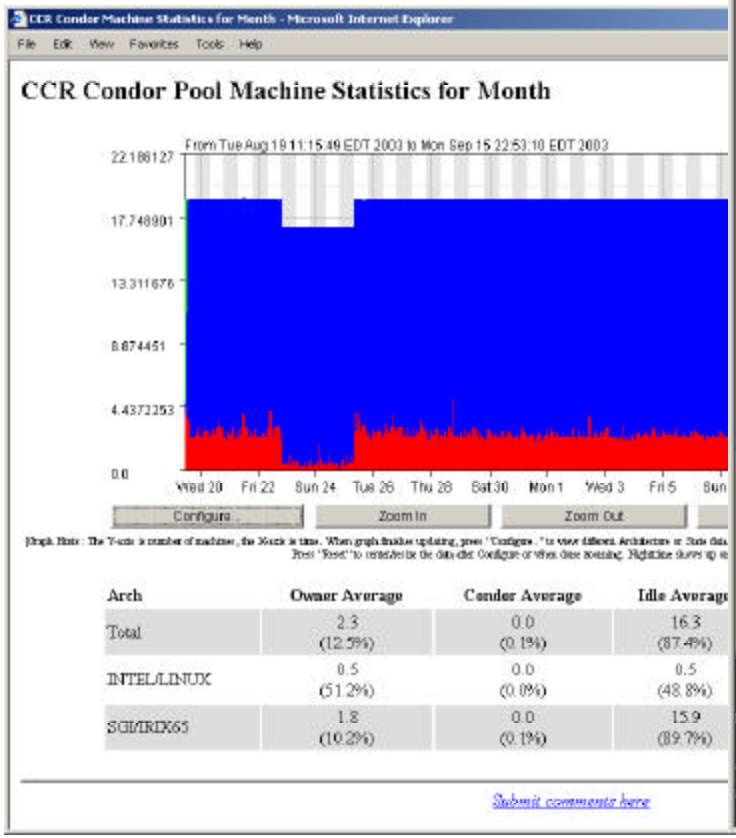

Groundwater Flow Modeling

- **Regional-scale modeling of groundwater flow and contaminant transport (Great Lakes Region)**
- **Ability to include all hydrogeologic features as independent objects**
- **Current work is based on *Analytic Element Method***
- **Key features:**
  - **High precision**
  - **Highly parallel**
  - **Object-oriented programming**
  - **Intelligent user interface**
  - **GIS facilitates large-scale regional applications**
- **Utilized 10,661 CPU days (32 CPU years) of computing in past year on CCR's commodity clusters**

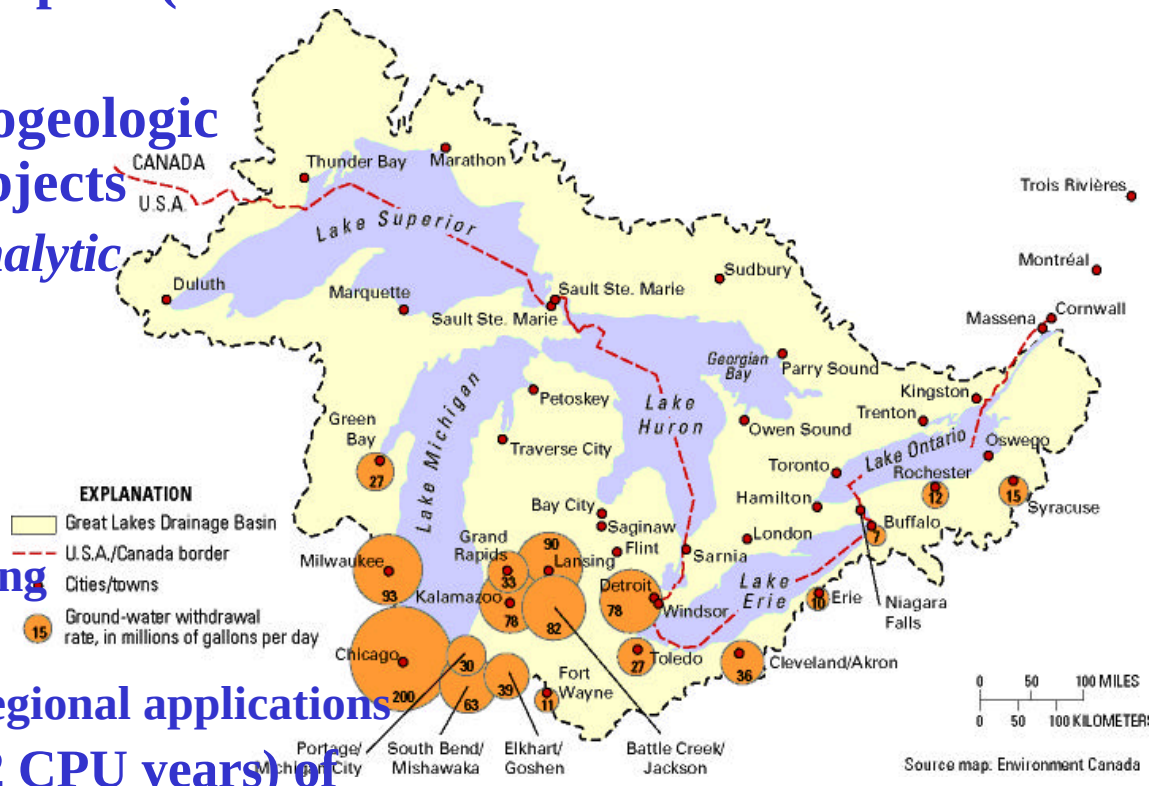

Risk Mitigation

- Integrate information from several sources
  - Simulation results
  - Remote sensing
  - GIS data
- Develop realistic 3D models of geophysical mass flows
- Present information at user appropriate resolutions

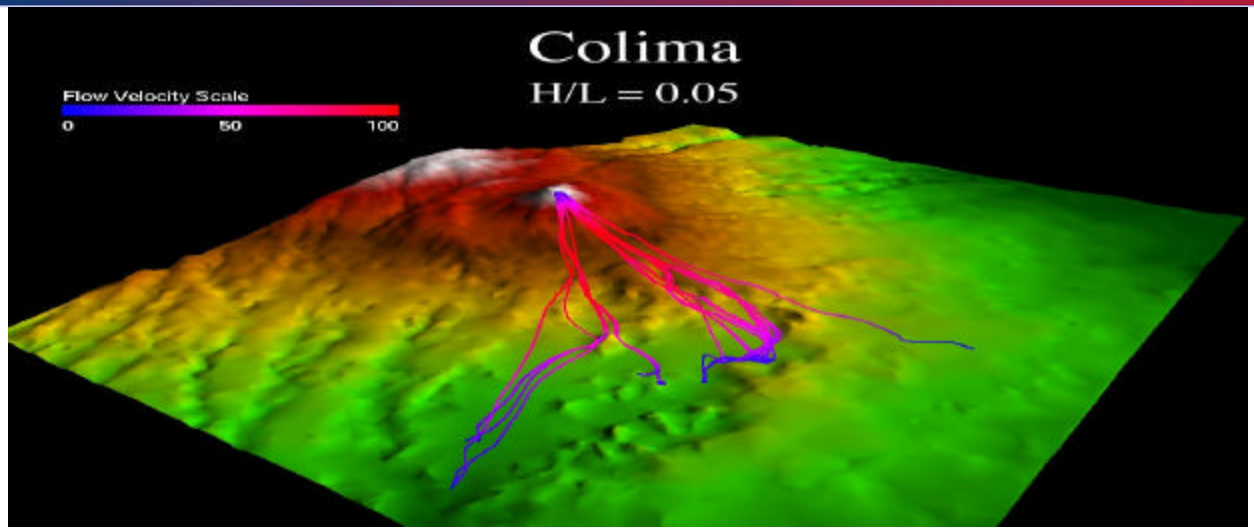

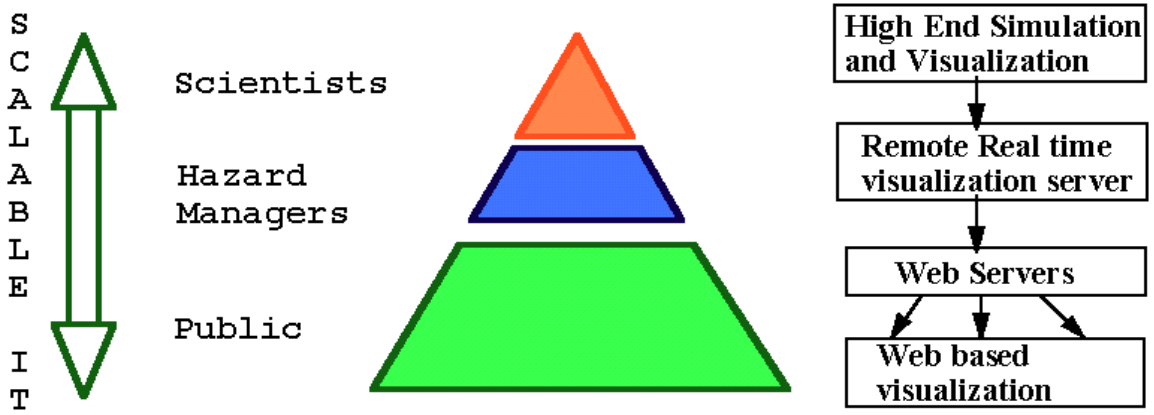

Protein Folding

- **Ability of proteins to perform biological function is attributed to their 3-D structure.**
- **Protein folding problem refers to the challenge of predicting 3-D structure from amino-acid sequence.**
- **Solving the protein folding problem will impact drug design.**

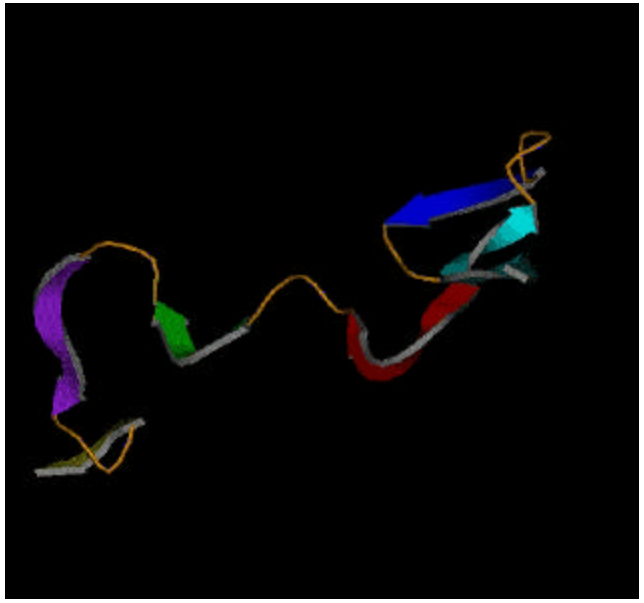

CCR Visualization Resources

- **Fakespace ImmersaDesk R2**
  - Portable 3D Device
- **Tiled-Display Wall**
  - 20 NEC projectors: 15.7M pixels
  - Screen is 11’ ´ 7’
  - Dell PCs with Myrinet2000
- **Access Grid Node**
  - Group-to-Group Communication
  - Commodity components
- **SGI Reality Center 3300W**
  - Dual Barco’s on 8’ ´ 4’ screen
- **VREX VR-4200 Stereo Imaging Projector**
  - Portable projector works with PC

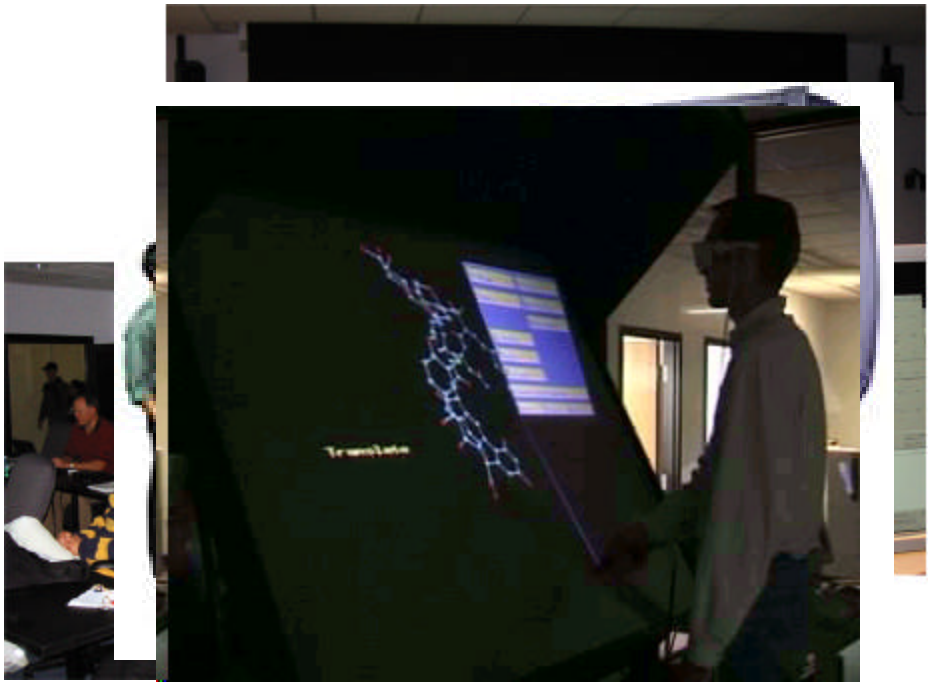

CCR Visualization Projects

- **Peace Bridge Visualization Project**
  - CCR, Niagara College of Canada, Bergmann, IBC, eMedia Inc, Parsons Engineering
- ***StreetScenes*®**
  - Virtual Reality (VR) software solution for 3D visualization of surface traffic
- **Buffalo Niagara Medical Campus**
  - CCR, IBC Digital Inc.
- **Emergency response / GIS – Earthquake**
  - CCR
- **Williamsville Toll Barrier**
  - CCR, TVGA
- **Accident Reconstruction**
  - CCR, TVGA
- **Biomedical Imaging**
  - CCR

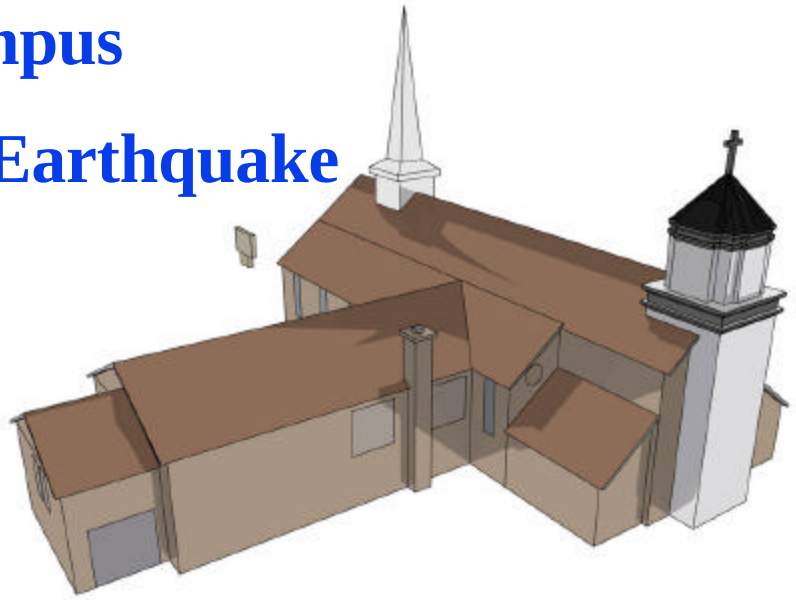

Visualization in Planning Studies

Bergmann Associates: Indian Street Bridge - Florida

Williamsville Toll Barrier Improvement Project

Initial Photo Match incorporating real and computer-generated components

Real-time Simulation

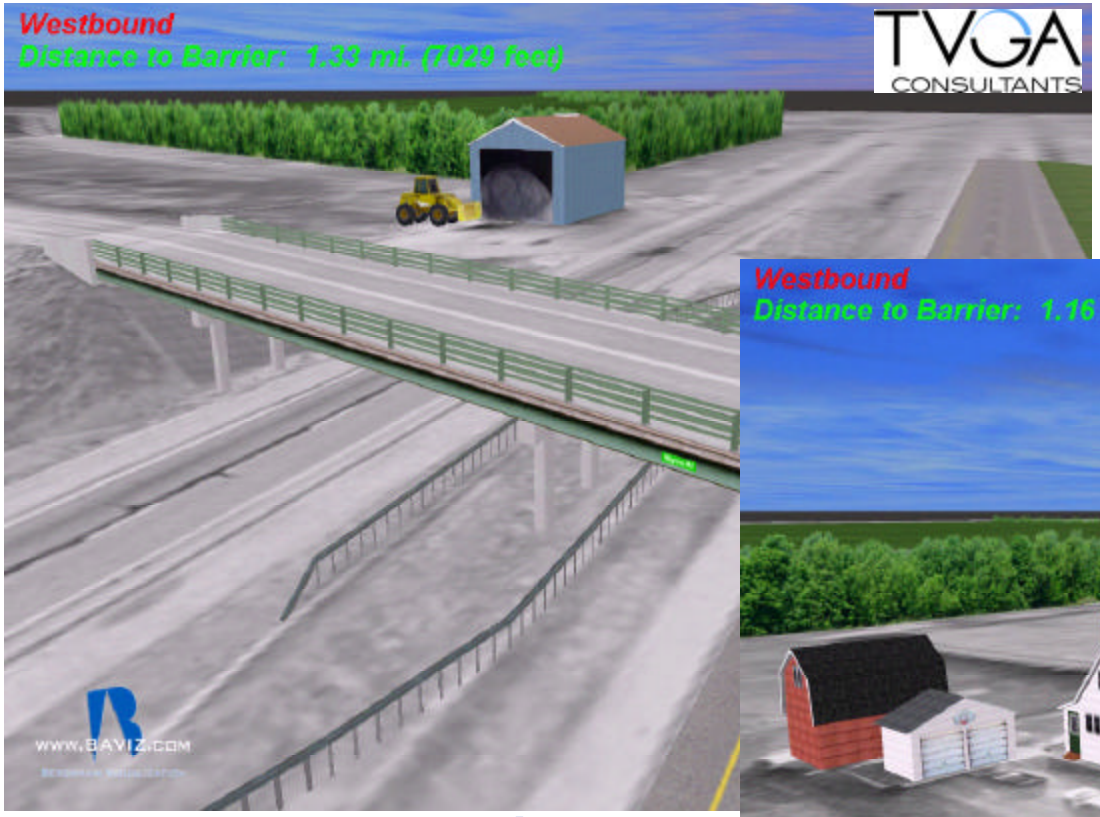

- Key Receptor Sites
- Multiple Viewpoints
- Fully Interactive
- Aerial Photography

Accident Reconstruction

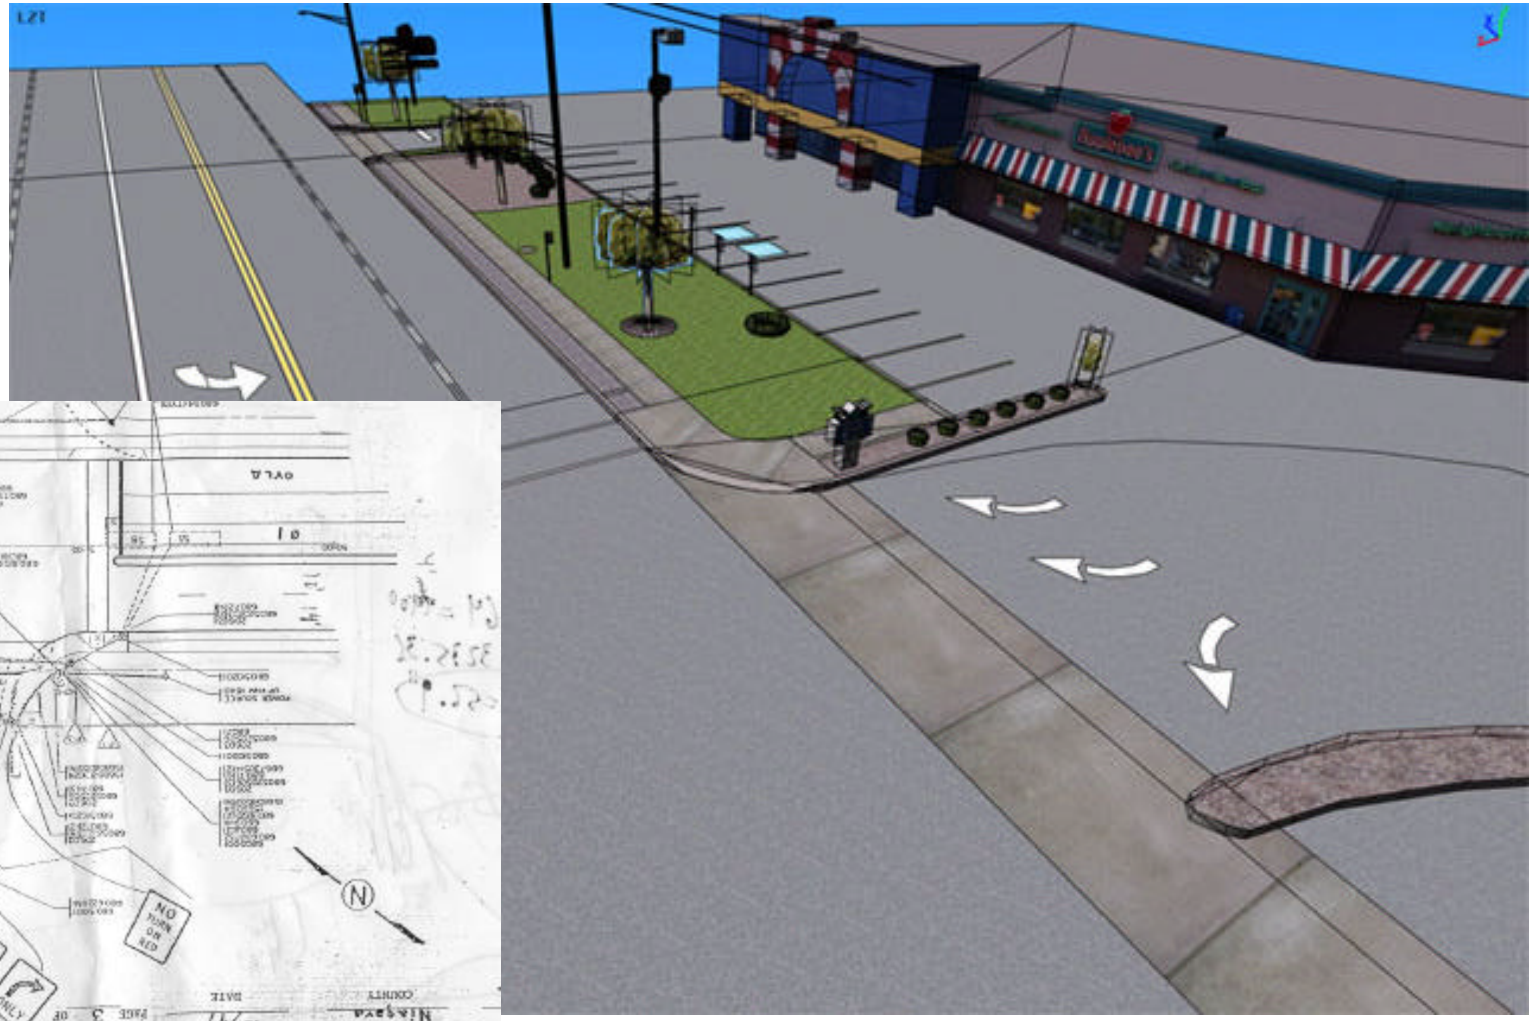

The Accident

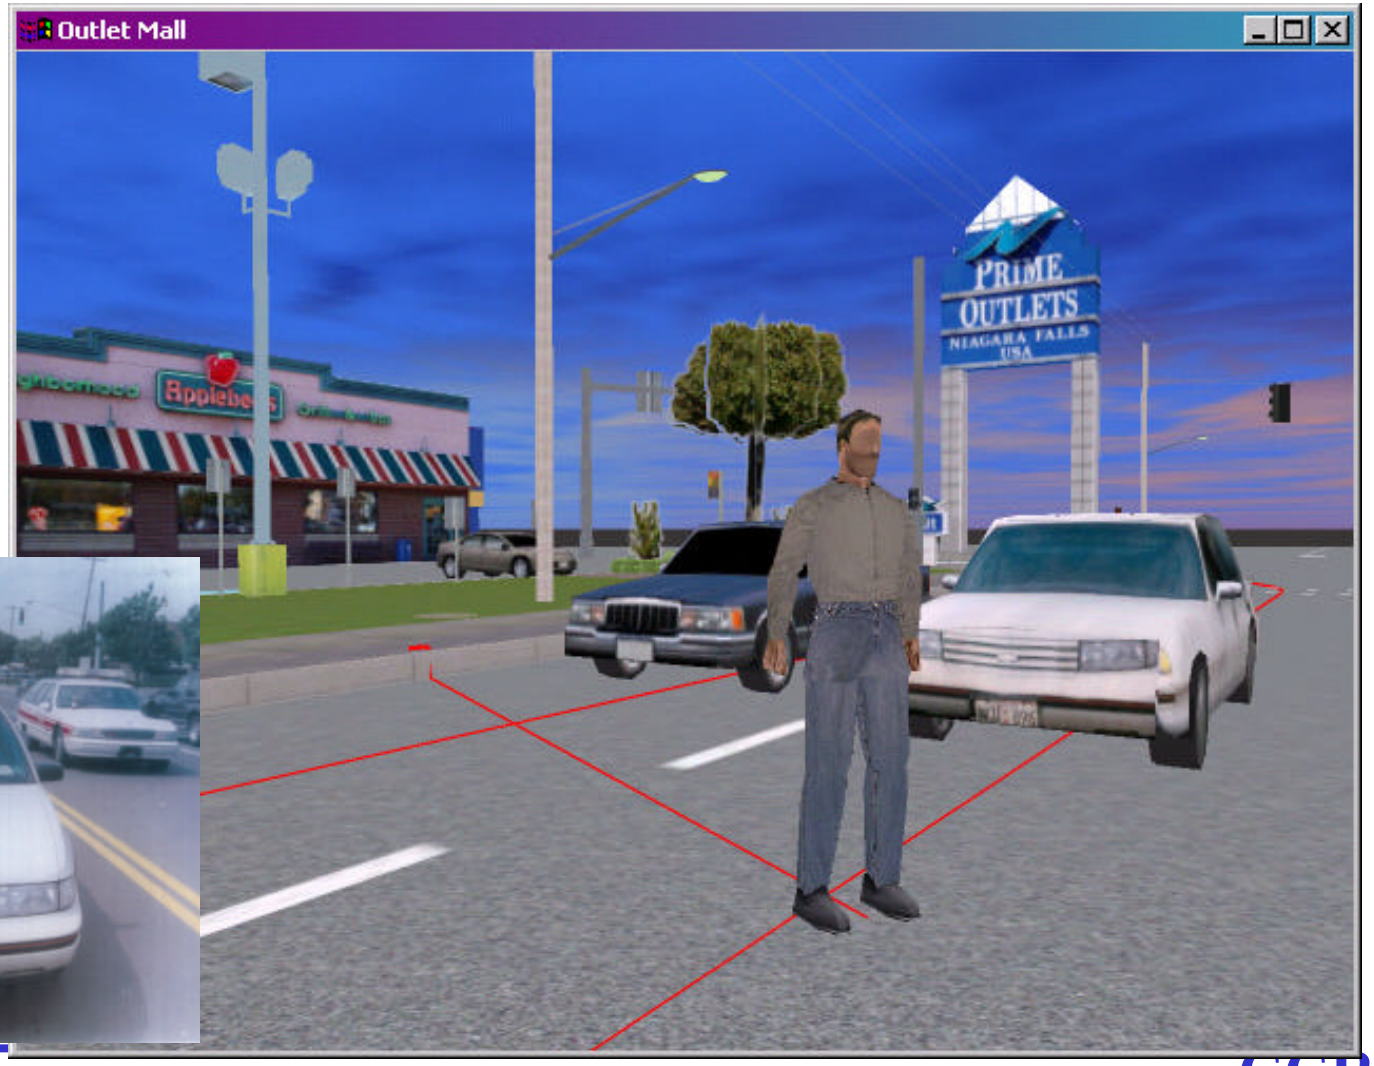

Accident Animation (Driver's View)

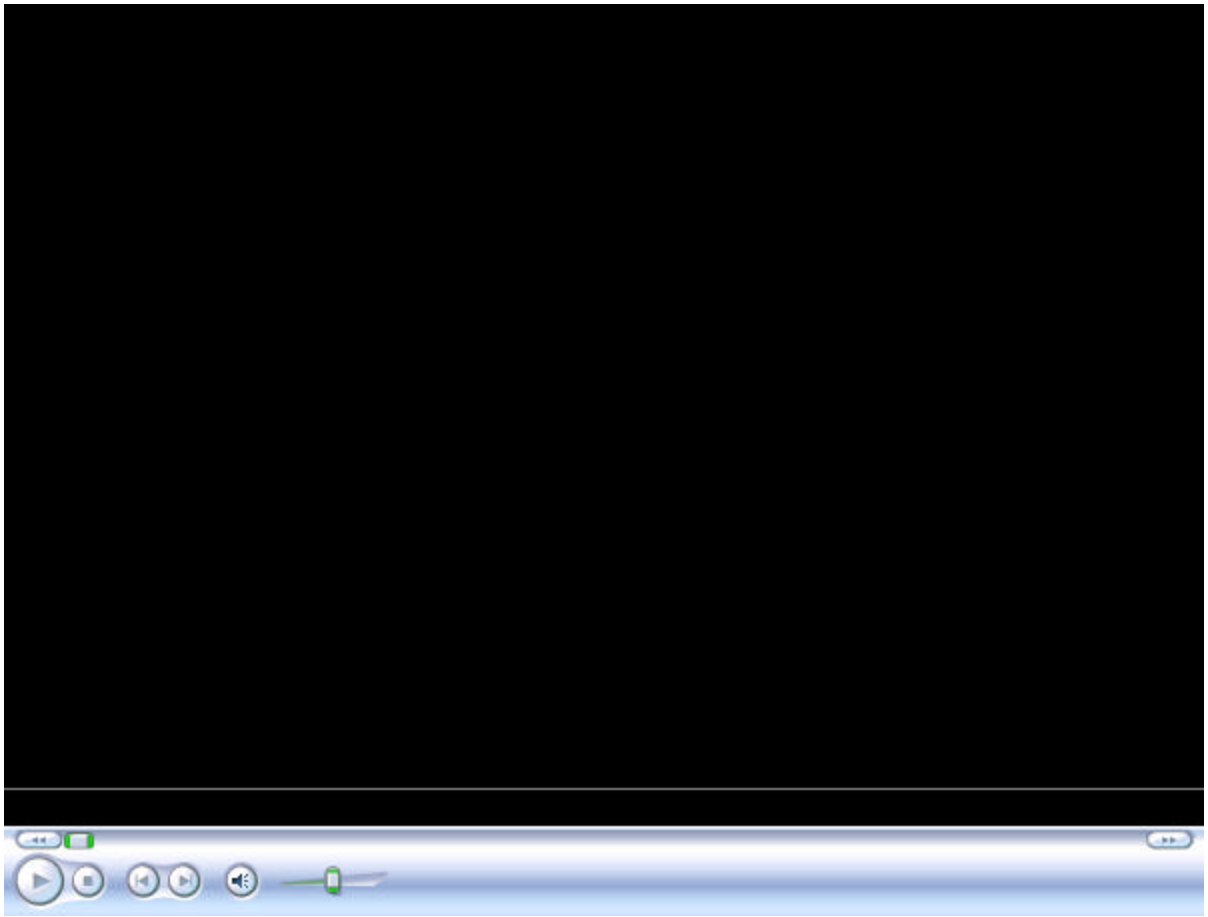

StreetScenes® Demo

- ***StreetScenes***® **is a Virtual Reality (VR) software solution for 3D visualization of surface traffic**
- **3D model of proposed soccer stadium in Rochester**
- **Used *StreetScenes*® to import output file from Synchro traffic simulation**

3D Medical Visualization App

- Collaboration with Children's Hospital
  - Leading miniature access surgery center
- Application reads data output from a CT Scan
- Visualize multiple surfaces and volumes
- Export images, movies or CAD representation of model

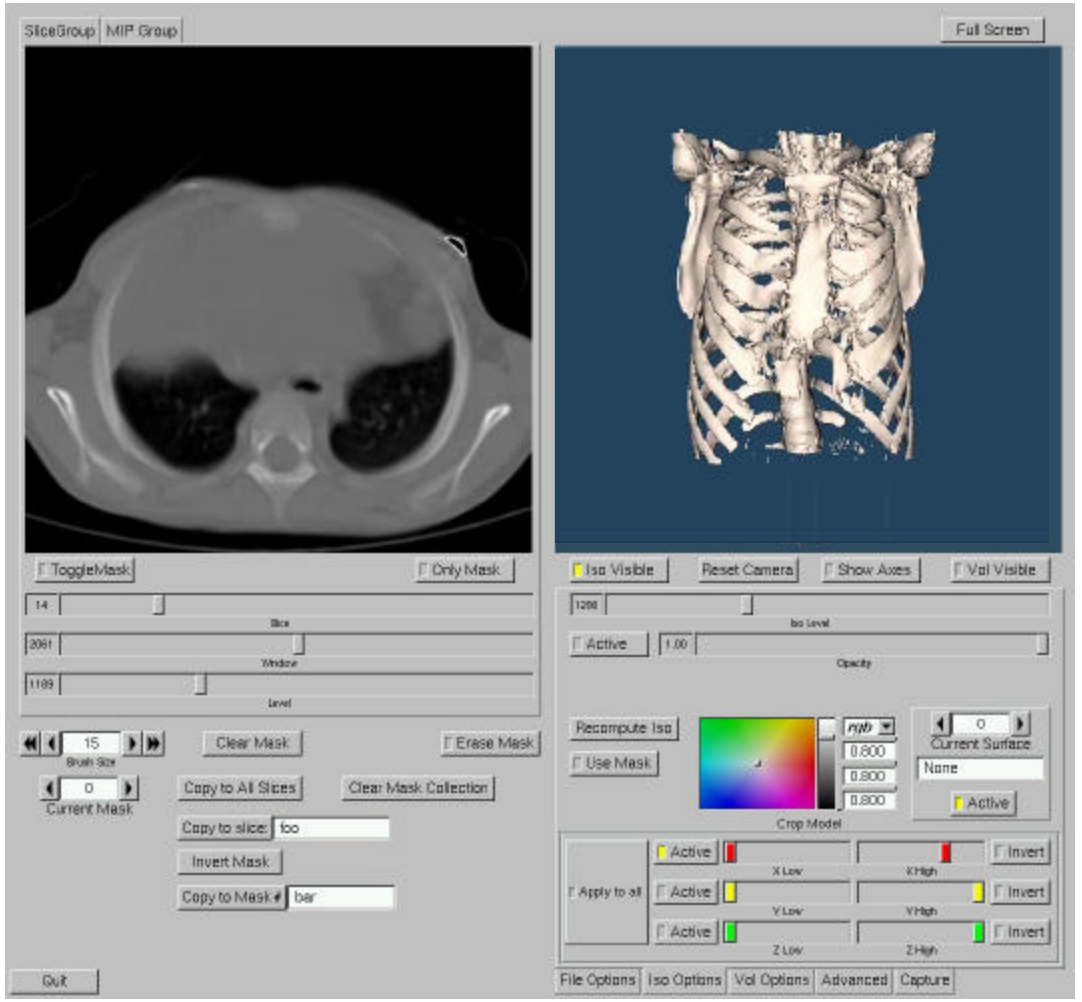

Multiple Sclerosis Project

- **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**

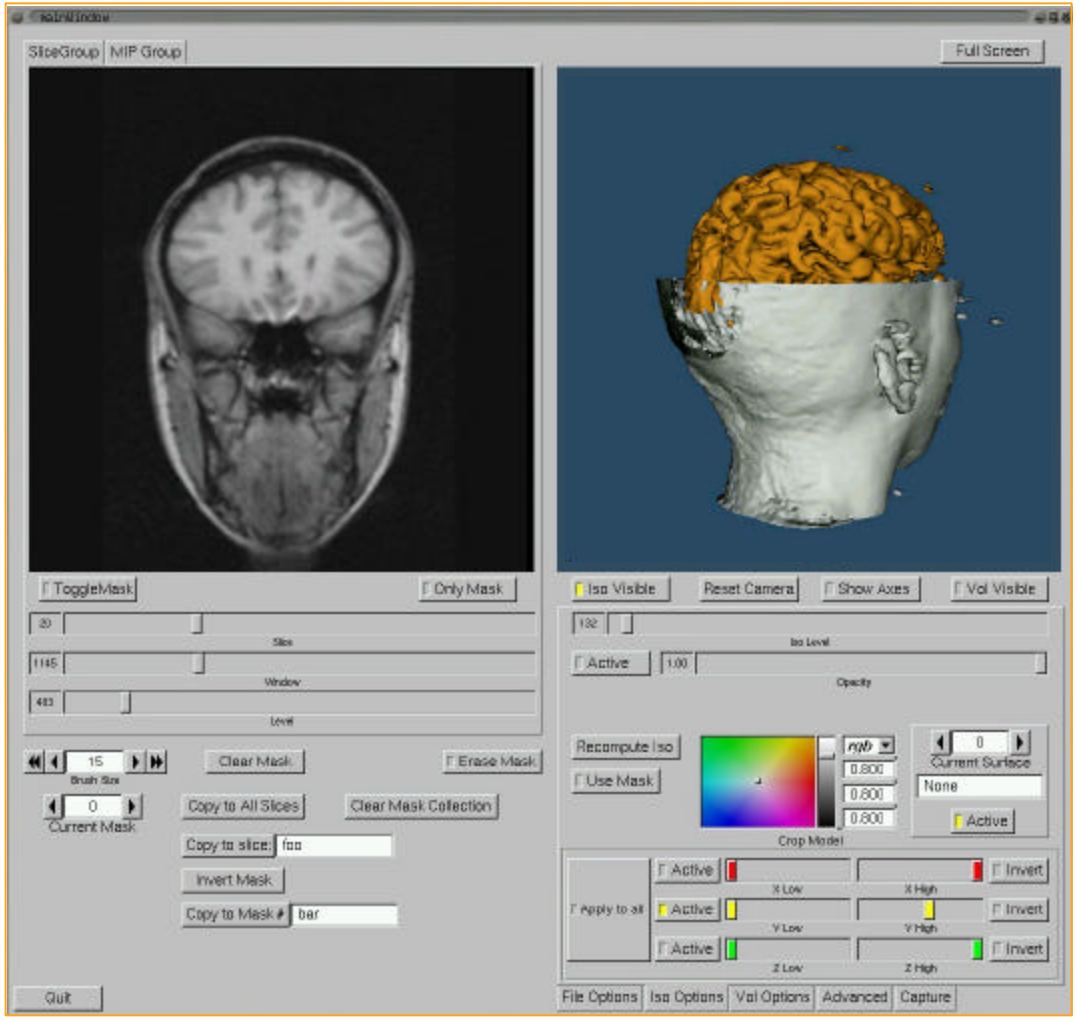

Select WNY Synergies

- **IBC Digital**
  - Gov. Pataki Visit
  - Peace Bridge (Early & Current)
  - Buffalo-Niagara Medical Campus
  - Compute Cycles for Animation
- **Bergmann Associates**
  - Peace Bridge (Current)
  - NYS Thruway Toll Plaza
- **Azar & More**
  - Reenactment of 1901 Pan Am Exhibition
  - PHSCologram & Courses
  - Avid Digital Editing
- **Niagara College**
  - Start up
  - Peace Bridge (Current)
- **Hauptman-Woodward Medical Research Institute**
  - Computing
  - Collaboratory
- **The Children's Hospital of Buffalo**
  - Medical Visualization
- **Veridian**
  - Battlespace Management

Outreach

- **HS Summer Workshops in Computational Science**
  - **Chemistry, Bioinformatics, Visualization**
  - **10-14 HS Students Participate Each Summer for 2 weeks**
  - **Project-Based Program**

Outreach

- **Pilot HS Program in Computational Science**
  - **Year long extracurricular activity at Mount St. Mary's, City Honors, and Orchard Park HS**
  - **Produce next generation scientists and engineers**
  - **Students learn Perl, SQL, Bioinformatics**
  - **$50,000 startup funding from Verizon, PC's from HP**

Media Coverage

BUFFALO
Business First
Western New York's Business Newspaper

HMOs cut Medicare premiums

Patio home development proposed for Town of Aurora

Grammy: Designs on Buffalo

Ani DiFranco, art director up for award

UB brings bioinformatics to a younger generation

THE BUFFALO NEWS

EDUCATION

University at Buffalo undergraduate David Walla works with Jaclyn Shaw, right, to demonstrate the "Next Generation Scientists" program. At left is Shannon D'Arcy.

An early look at bioinformatics

By EMMA D. SAPONG
News Northtowns Bureau


Outreach
CCR

Contact Information

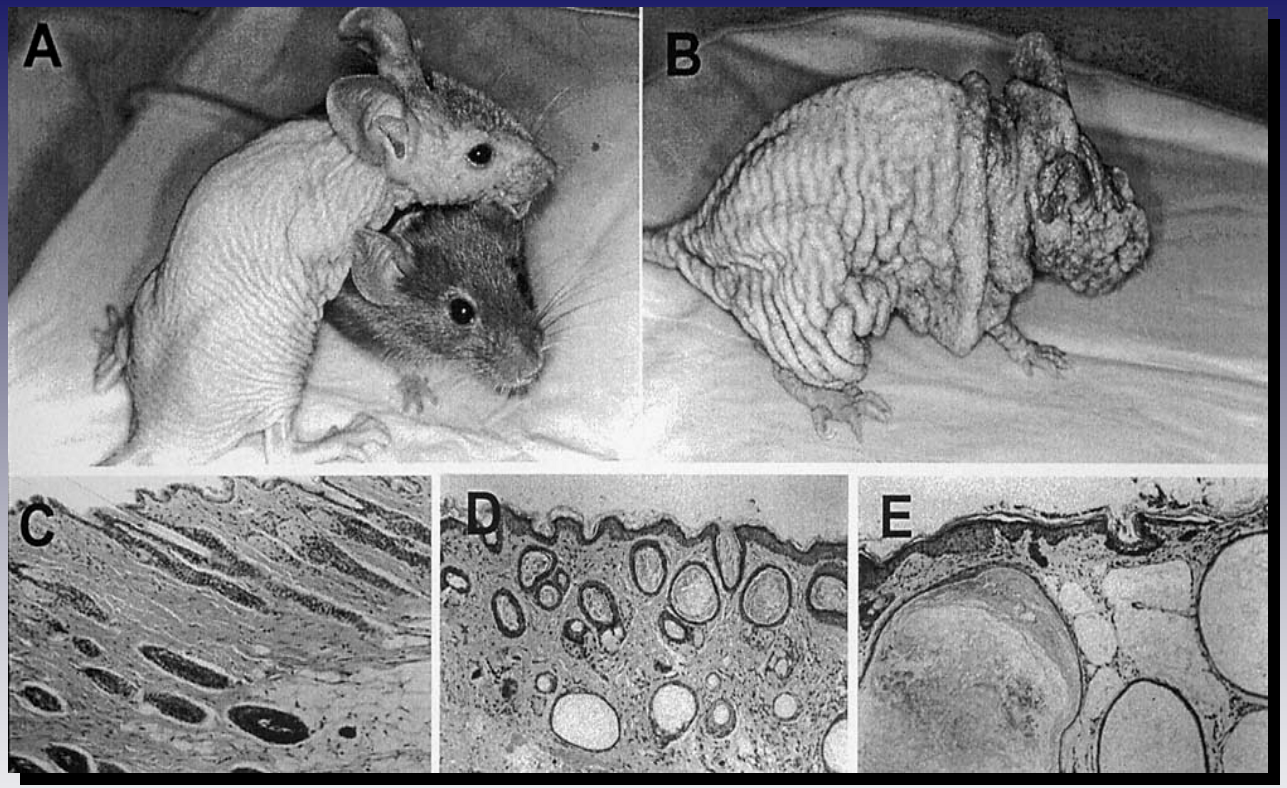

miller@buffalo.edu
www.ccr.buffalo.edu