A Global Petascale Computing Environment for Fundamental Science

Paul Messina

Center for Advanced Computing Research

California Institute of Technology

March 14, 2002

Outline

- A brief history of grids
- Why grids are getting so much attention
- Some examples of current and planned grid-based projects
- Grids in 2010
- Summary

The Grid Concept

- Grid R&D has its origins in high-end computing & metacomputing, but…
- In practice, the "Grid problem" is about *resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations*
- ***Data*** is often the focus
  - as opposed to classical numerically intensive simulations

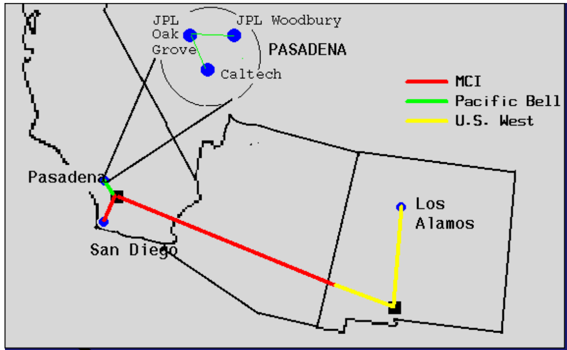

A Brief History

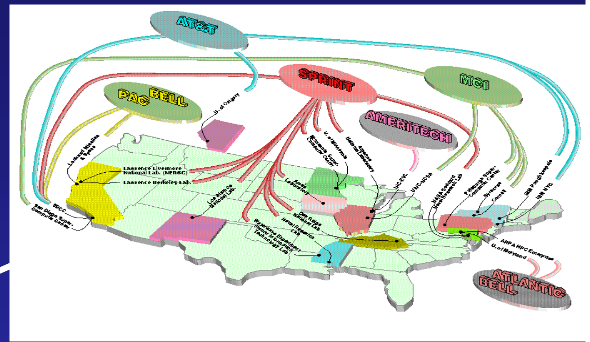

- Early 90s
  - Gigabit testbeds, metacomputing
- Mid to late 90s
  - Early experiments (e.g., I-WAY), academic software projects (e.g., Globus, Legion), application experiments
- 2001
  - Major application communities emerging
  - Major infrastructure deployments
  - Growing technology base
  - Global Grid Forum: ~500 people, >90 orgs, 20 countries

Walking down a street in St.Louis, Missouri (the "Show Me" state)

A closer look

Paul Messina

and closer...

A Different Grid Vision

- Analogy with the electrical power grid
- “Unlimited” ubiquitous distributed computing
- Transparent access to multi-petabyte distributed data bases
- Easy to plug resources into
- Complexity of the infrastructure is hidden

e-Science and Information Utilities? *(John Taylor)*

e-Science

- science increasingly done through ***distributed global collaborations*** between people, enabled by the internet
- using very large data collections, terascale computing resources and high performance visualisation
- derived from instruments and facilities controlled and shared via the infrastructure
- Scaling X1000 in processing power, data, bandwidth

Cyberinfrastructure

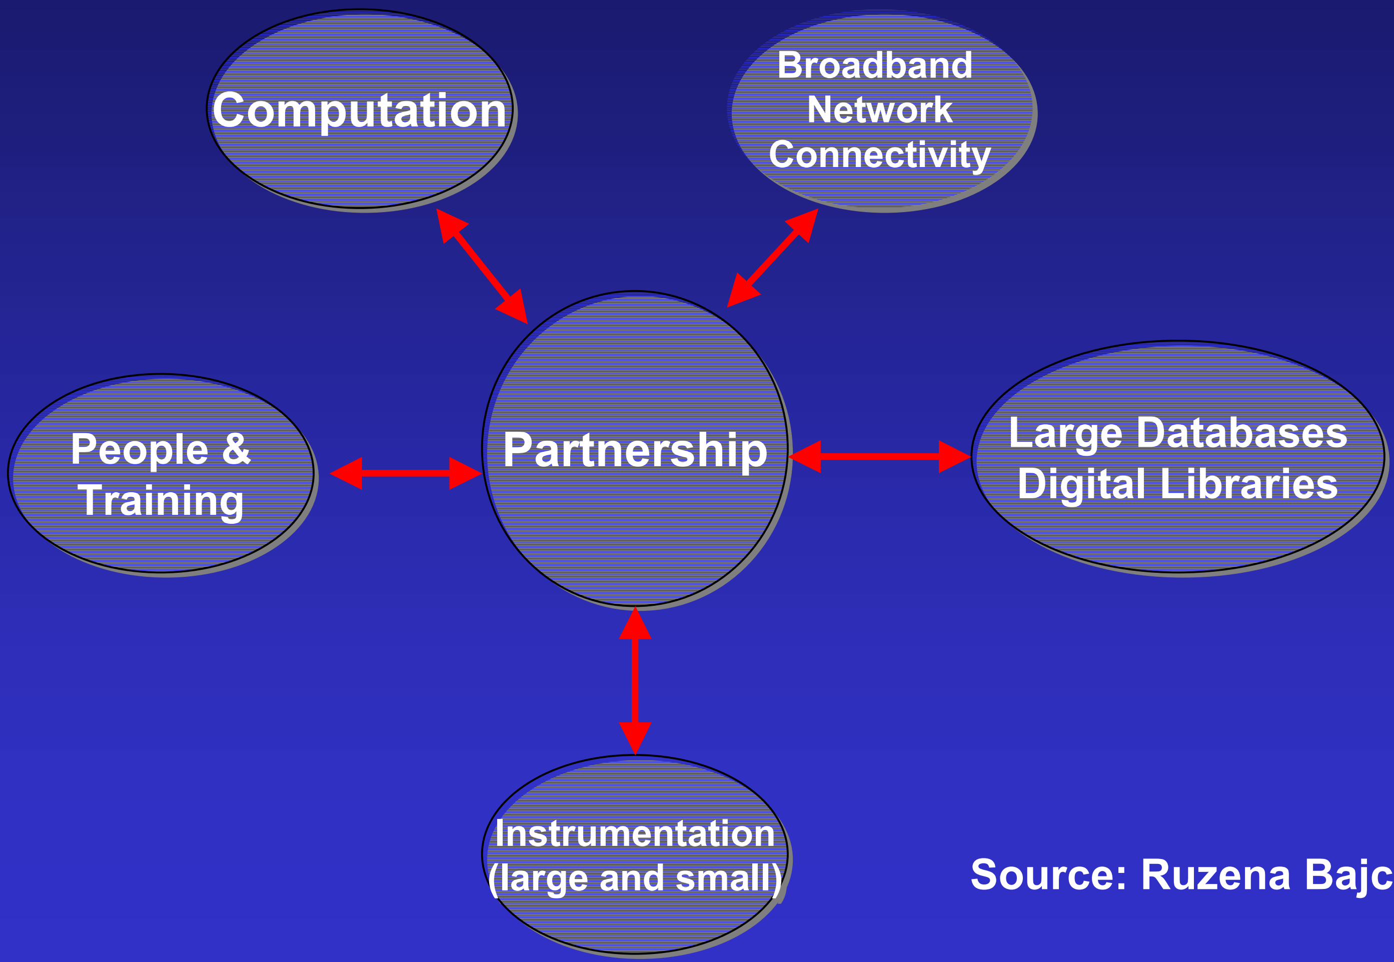

Source: Ruzena Bajcsy, NSF

Why such a fast emergence of Grids as important for science?

Technology Trends

- Very fast networks can be built and are becoming *much* cheaper
- Archival and disk storage becoming *much* cheaper
- Sensors and instruments becoming ubiquitous and have much higher resolution => oceans of data
- Computing power continues to increase
  - so many groups can analyze lots of data
- Middleware, while still immature, is not *vaporware*
- Some demonstration projects have shown the benefits
- Everyone uses the web, so concept is not foreign

Larger Scale Computational Science and Engineering

- Larger computations on remote systems are feasible
  - the large I/O and storage needs can be handled
  - more resources can be used, pipelined or concurrently
- Harvesting of computer cycles from thousands of network-connected systems is made easier
- Telecollaboration is also facilitated by grid-based tools
  - larger distributed teams are realistic

Grids are (Mis)Perceived as saving money on computing

- Funding agencies often get the (erroneous) idea that they will be able spend less on high-end computing
  - not for hardware
    - major scientific systems are chronically overloaded, so there are few if any spare cycles
    - of course one can harvest cycles on desktops, small clusters but those "free" cycles are a "bonus," unlikely to eliminate the need for the high-end systems
  - not for software and operation
    - grid systems are substantially more complex
- But grids do make it much easier to aggregate resources for a shared goal

Instruments are expensive

(notice I use a telescope as an example, instead of a particle accelerator)

Grids Leverage Major Investments in Experimental Research Facilities

- Research instruments, satellites, particle accelerators, etc., cost a great deal
- Data from those devices can be accessed and analyzed by many more scientists
  - not just the team that gathered the data
- More productive use of instruments
  - calibration, data sampling during a run, via on-demand real-time processing

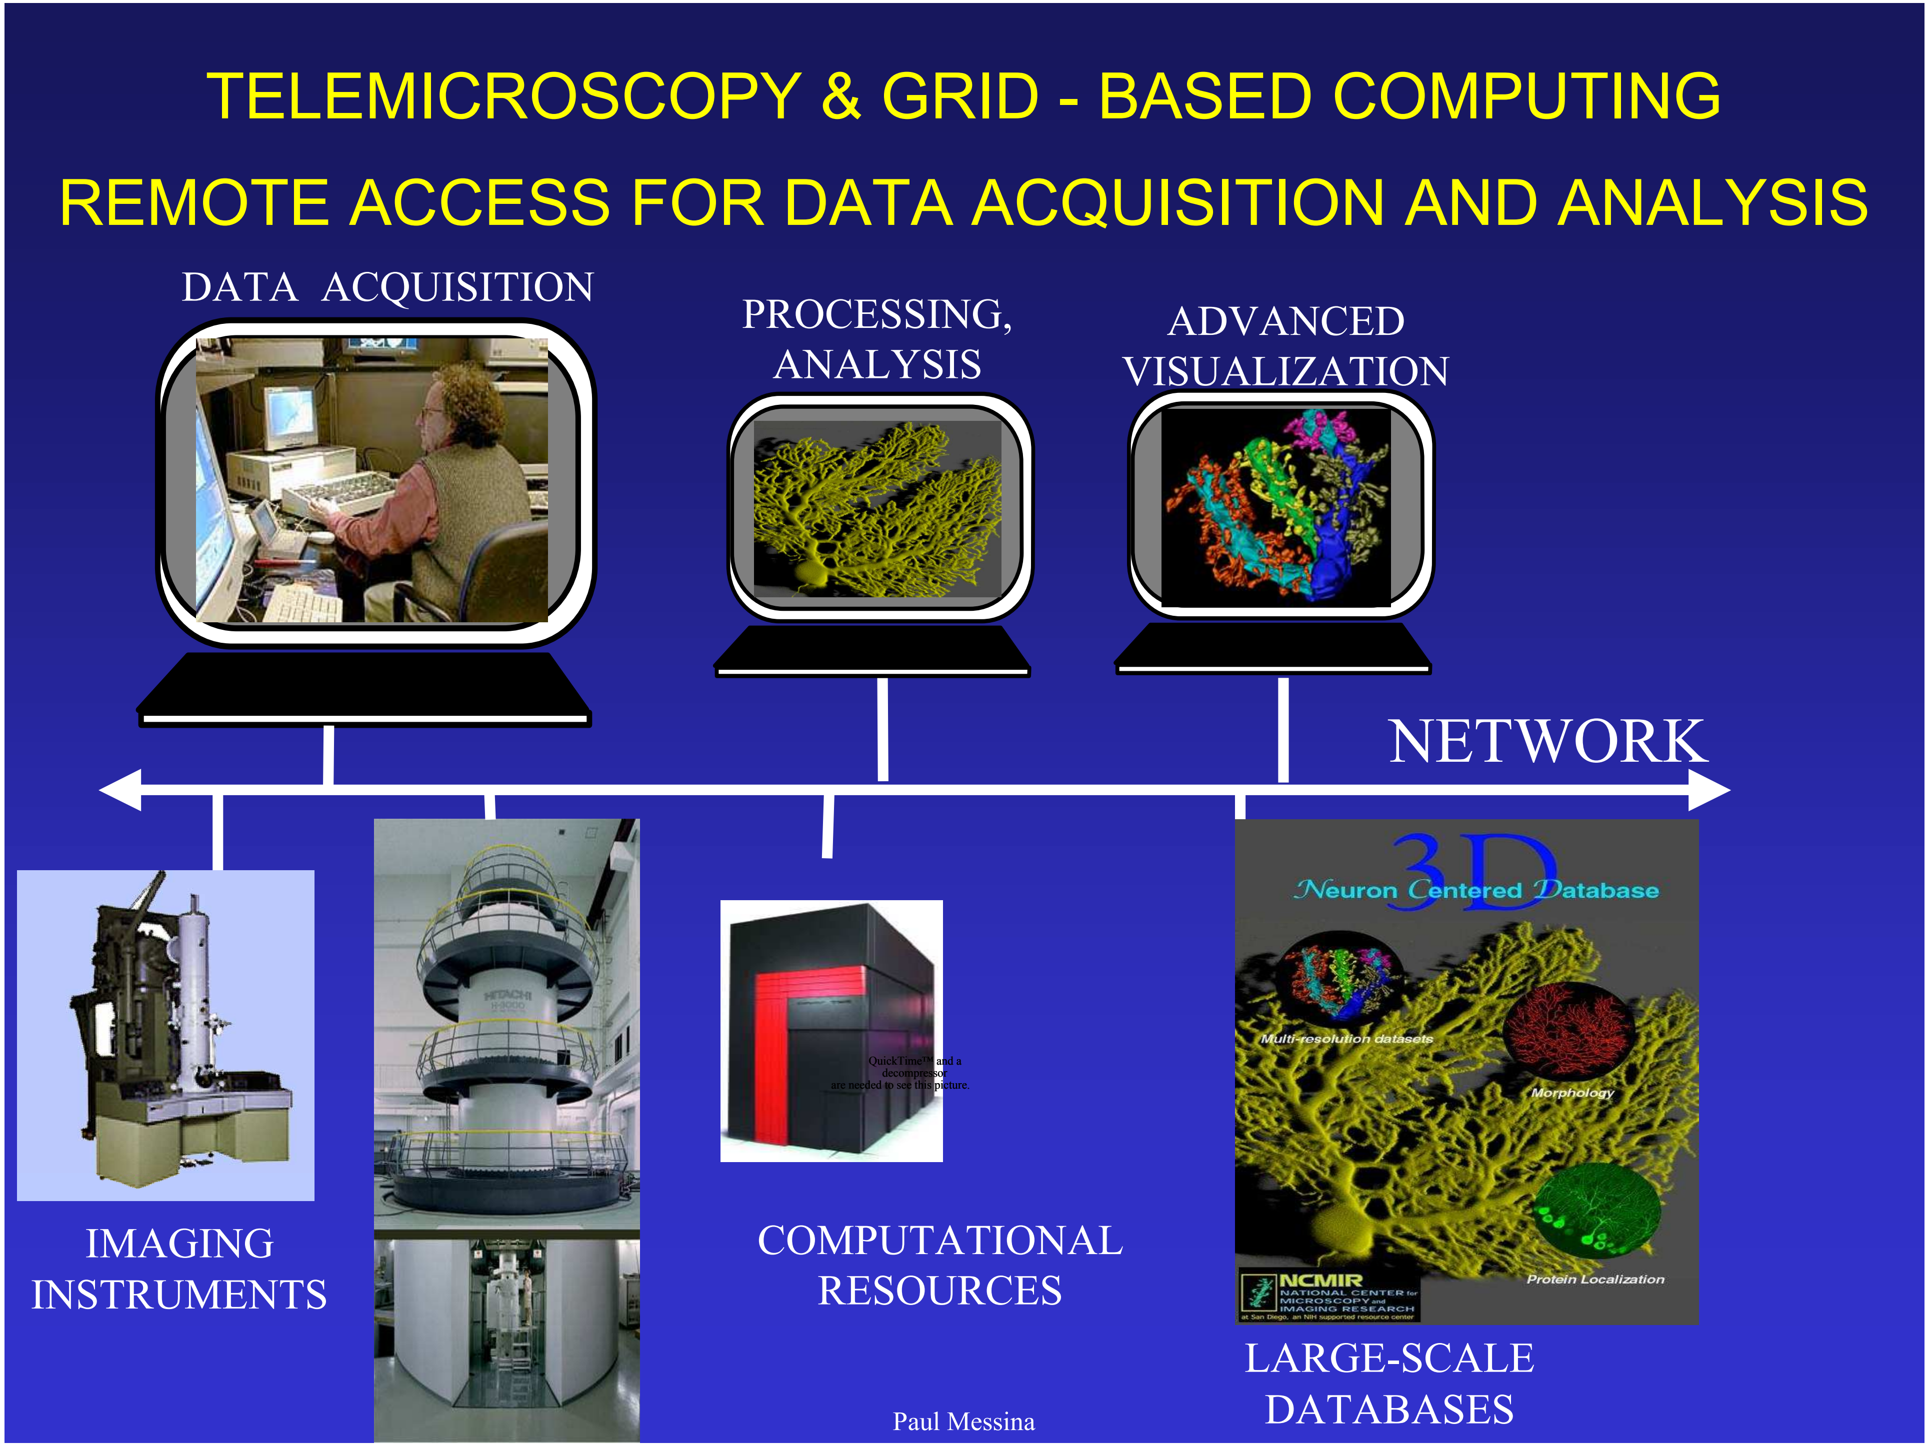

TELEMICROSCOPY & GRID - BASED COMPUTING
REMOTE ACCESS FOR DATA ACQUISITION AND ANALYSIS
DATA ACQUISITION
PROCESSING, ANALYSIS
ADVANCED VISUALIZATION
NETWORK
HITACHI
QuickTime™ and a decompressor are needed to see this picture.
3D Neuron Centered Database
Multi-resolution datasets
Morphology
Protein Localization
NCMIR NATIONAL CENTER for MICROSCOPY and IMAGING RESEARCH
IMAGING INSTRUMENTS
COMPUTATIONAL RESOURCES
LARGE-SCALE DATABASES
Paul Messina

And Grids enable new applications

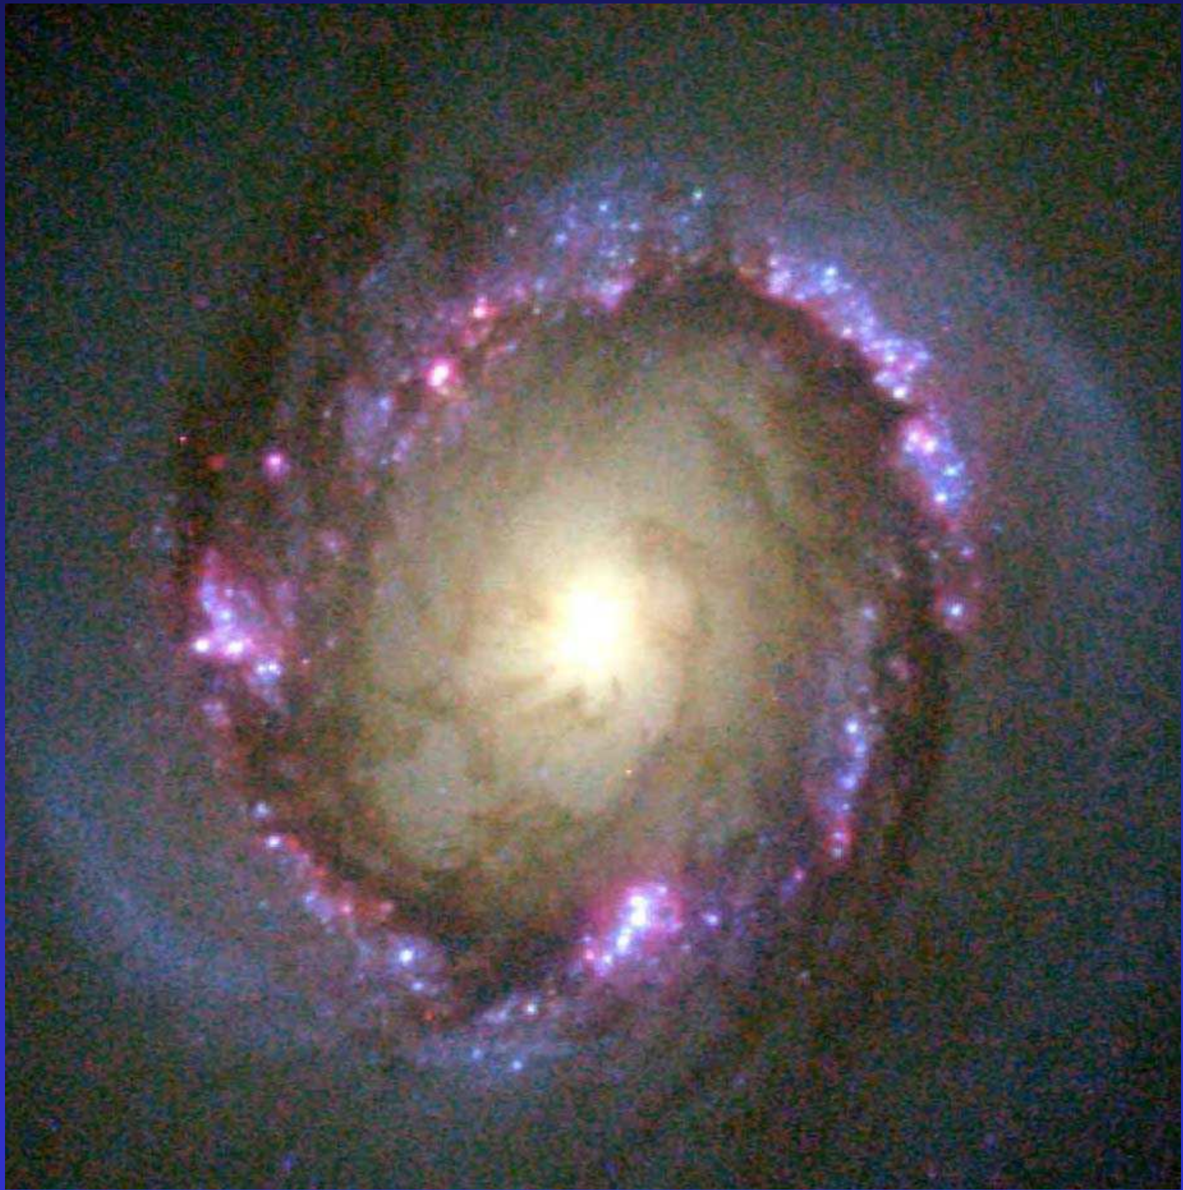

Virtual Observatory is an exemplar of the scientific opportunities

Digital Sky: A Federation of Sky Surveys

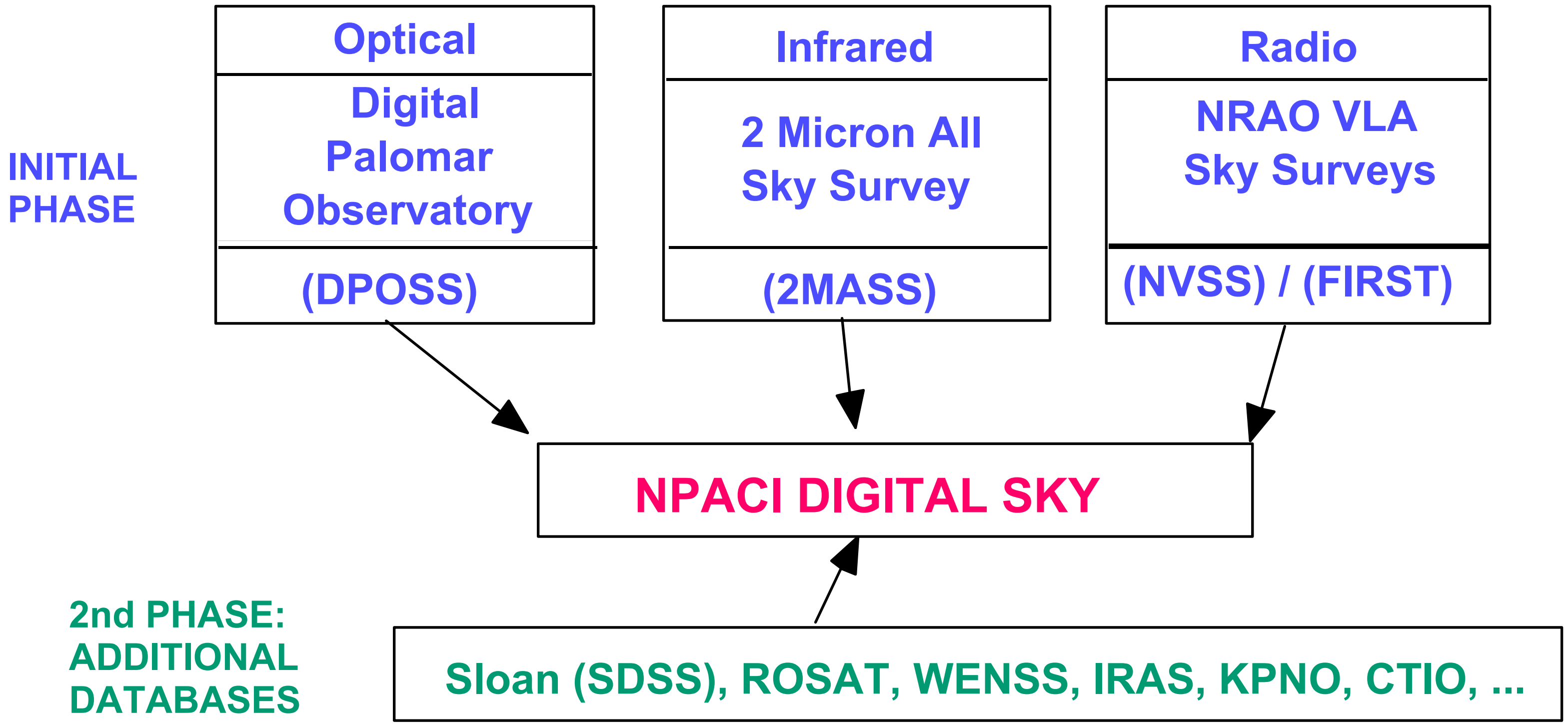

Digital Palomar Observatory Sky Survey (DPOSS)

- Optical northern sky survey (50% of sky)
  - Palomar Schmidt Telescope
  - 1 arc second pixels
  - 3 wavelength bands (G,R,I)
- Data
  - Catalog of 2 billion sources
  - Images (900 fields):
    - 1 Gigabyte/field x 3 bands
    - 3 Terabytes total

Globular Cluster

http://astro.caltech.edu/~george/dposs/dposs.html

2 Micron All Sky Survey (2MASS)

- Infrared full-sky survey
  - Two telescopes:
    - Mt. Hopkins & CTIO
  - 3 Bands - J,H,K
  - Pixel size: 2” x 2”
- Data
  - 20 Gigabytes/night
  - Catalog of 100 million sources
  - 10 Terabytes of image data

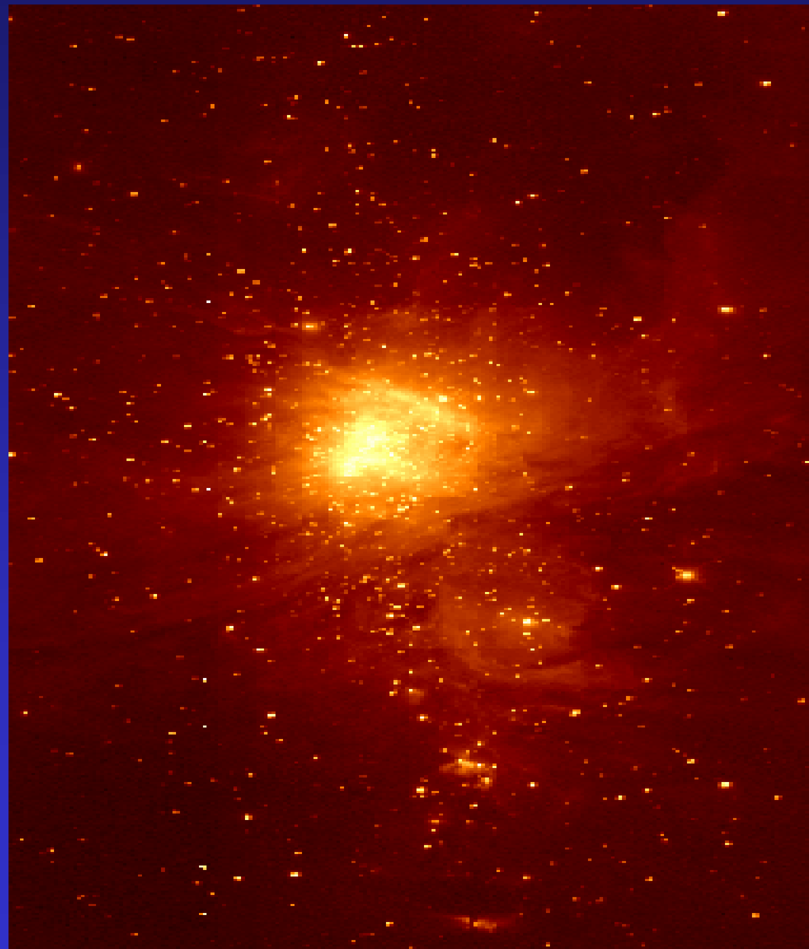

http://pegasus.phast.umass/edu/2mass/overview.html

2MASS

- The Two-Micron All-Sky Survey recently posted on its web site images of half a million galaxies and 162 million stars, enough data to fill 6,000 CD-ROMs
- This survey is carried out by 2 1.3 meter ground-based telescopes, one in Arizona and one in Chile

VLA Sky Surveys (NVSS + FIRST)

- NVSS
  - 21 cm
  - 82% of sky
  - 15” pixel
  - 2 million sources
  - Images: Stokes I, Q, and U
  - Processed (u,v - interferometric data)

http://info.cv.nrao.edu/~jcondon/nvss.html

(NVSS = NRAO VLA Sky Survey)

- **FIRST**
  - 21 cm
  - 10,000 degrees (complement to Sloan)
  - 1.8” pixel
  - 1 million sources

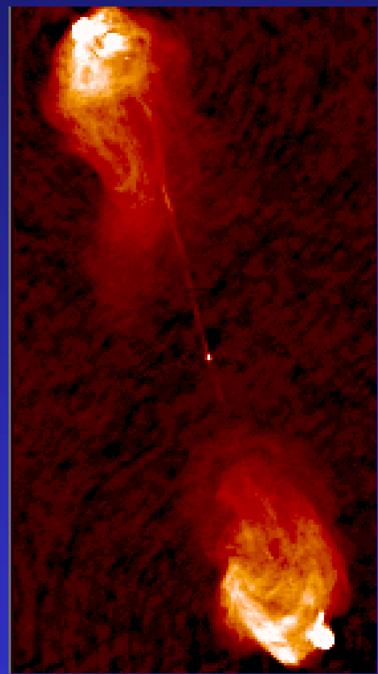

Cyg A in Radio

http://sundog.stsci.edu/top.html

(FIRST = Faint Images of the Radio Sky at 20 cm)

Ongoing Mega-Surveys

- Large number of new surveys
  - Multi-terabyte in size, 100 million objects or larg
  - In databases
  - Individual archives planned and under way
- Multi-wavelength view of the sky
  - More than 13 wavelength coverage within 5 years
- Impressive early discoveries
  - Finding exotic objects by unusual colors
    L,T dwarfs, high redshift quasars
  - Finding objects by time variability
    gravitational micro-lensing

MACHO
2MASS
SDSS
DPOSS
GSC-II
COBE
MAP
NVSS
FIRST
GALEX
ROSAT
OGLE, ...

Crab Nebula in 4 spectral regions

X-ray, optical, infrared, radio

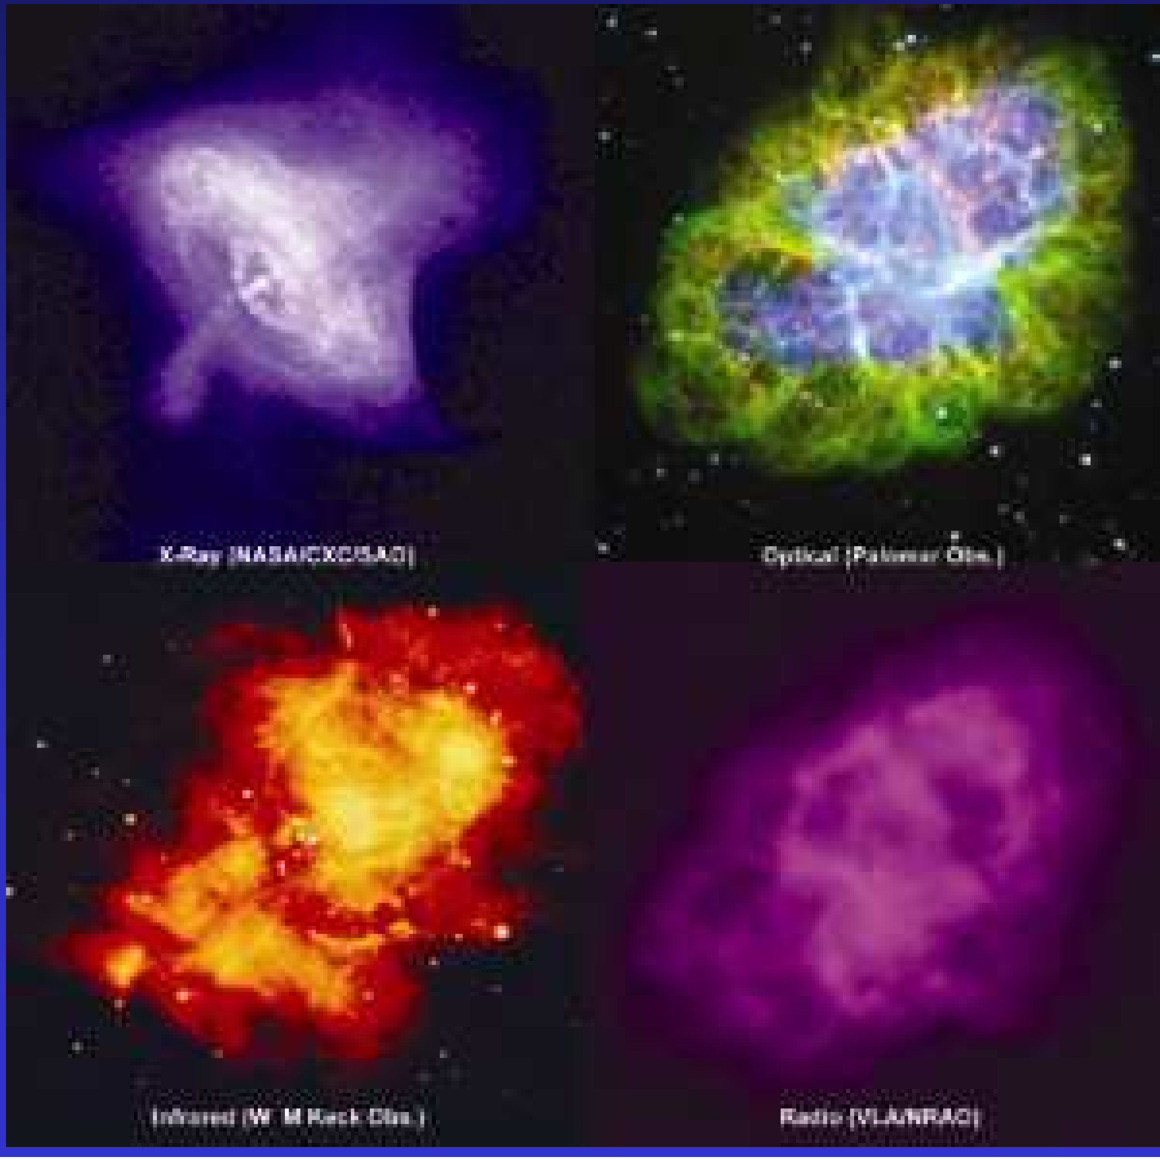

The Digital Sky concept captured the imagination

- The **National Academy of Sciences** Astronomy and Astrophysics Survey Committee, in its decadal survey entitled *Astronomy and Astrophysics in the New Millennium*, recommends, as a first priority, the establishment of a **National Virtual Observatory**
  - please excuse the "National" bit
  - see [NAS99] National Academy of Science, Astronomy and Astrophysics Survey Committee, *Astronomy and Astrophysics in the New Millennium (Decadal Survey)* http://www.nap.edu/books/0309070317/html/

Coming Floods of Astronomy Data
not just HEP data

- The planned Large Synoptic Survey Telescope will produce over 10 petabytes per year by 2008!
  - All-sky survey every few days, so will have fine-grain time series for the first time

Data Federation

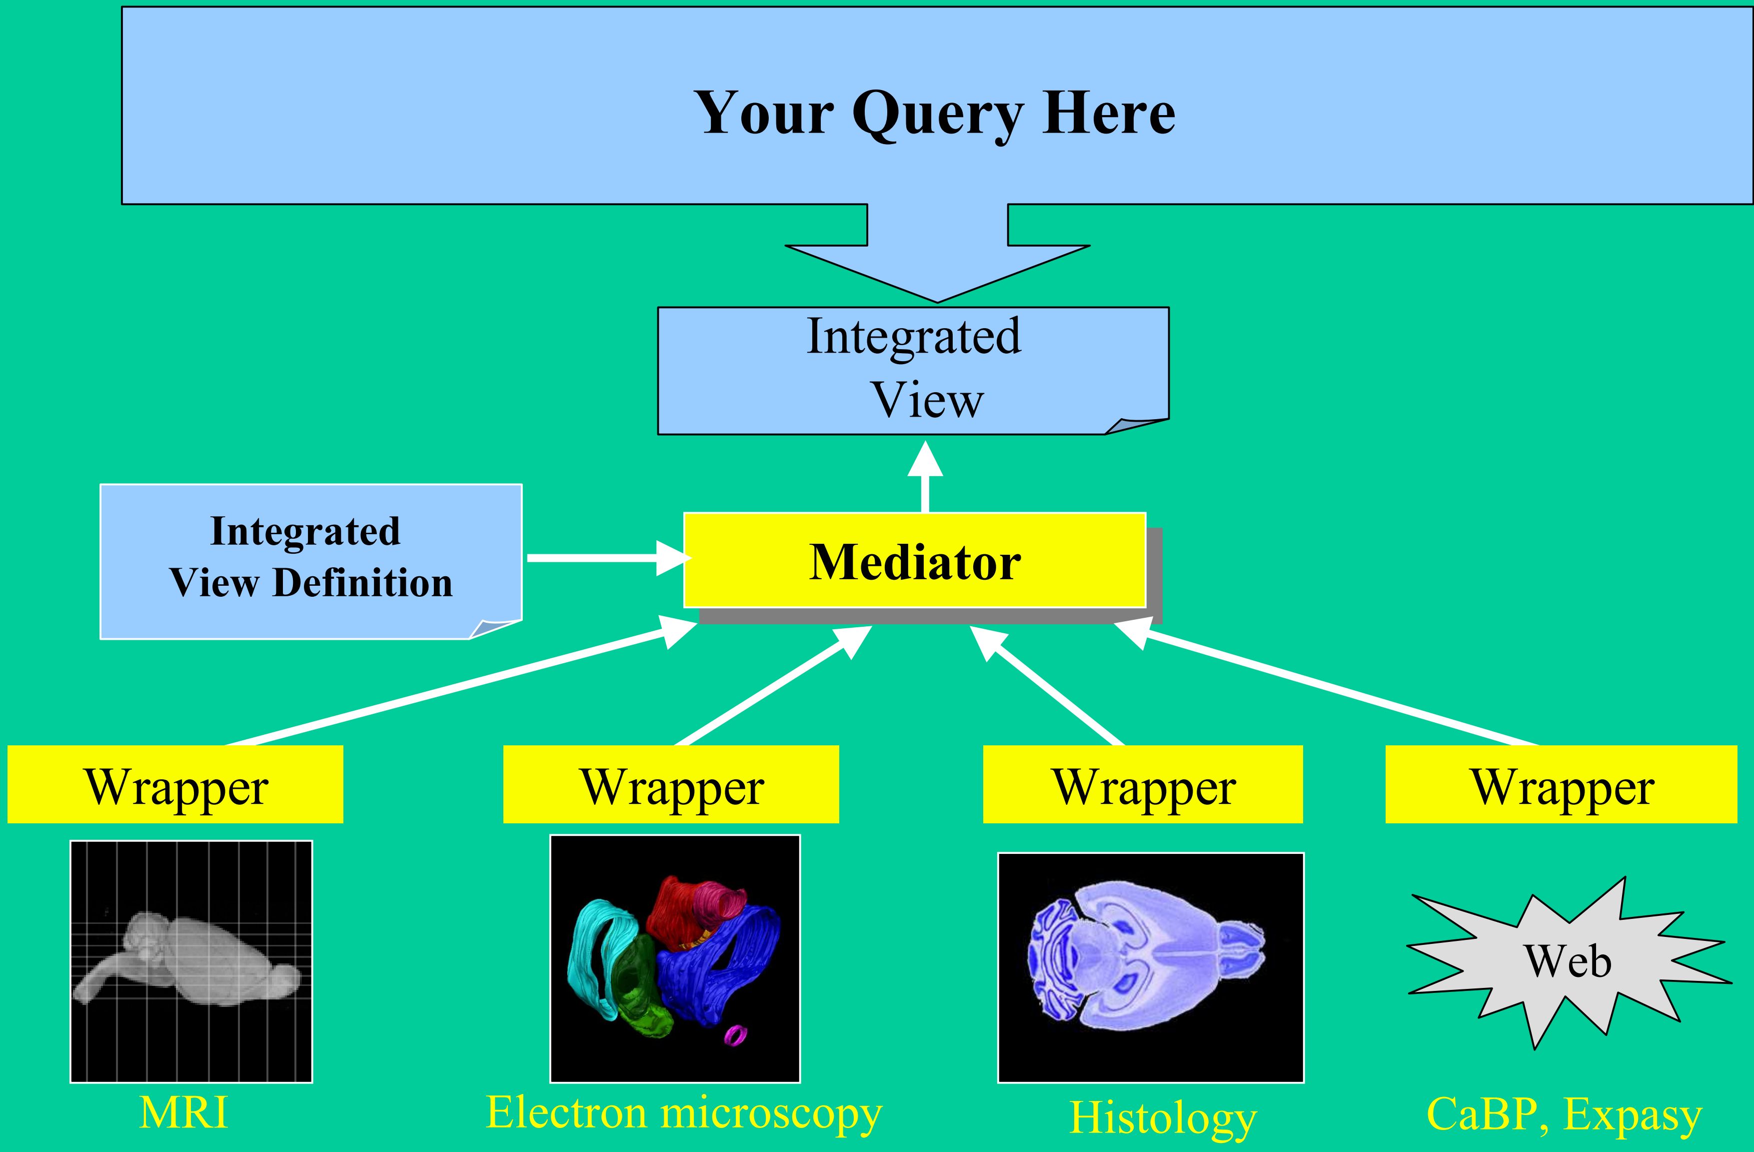

A selection of current and planned grid-based projects

NSF PACI 13.6 TF TeraGrid

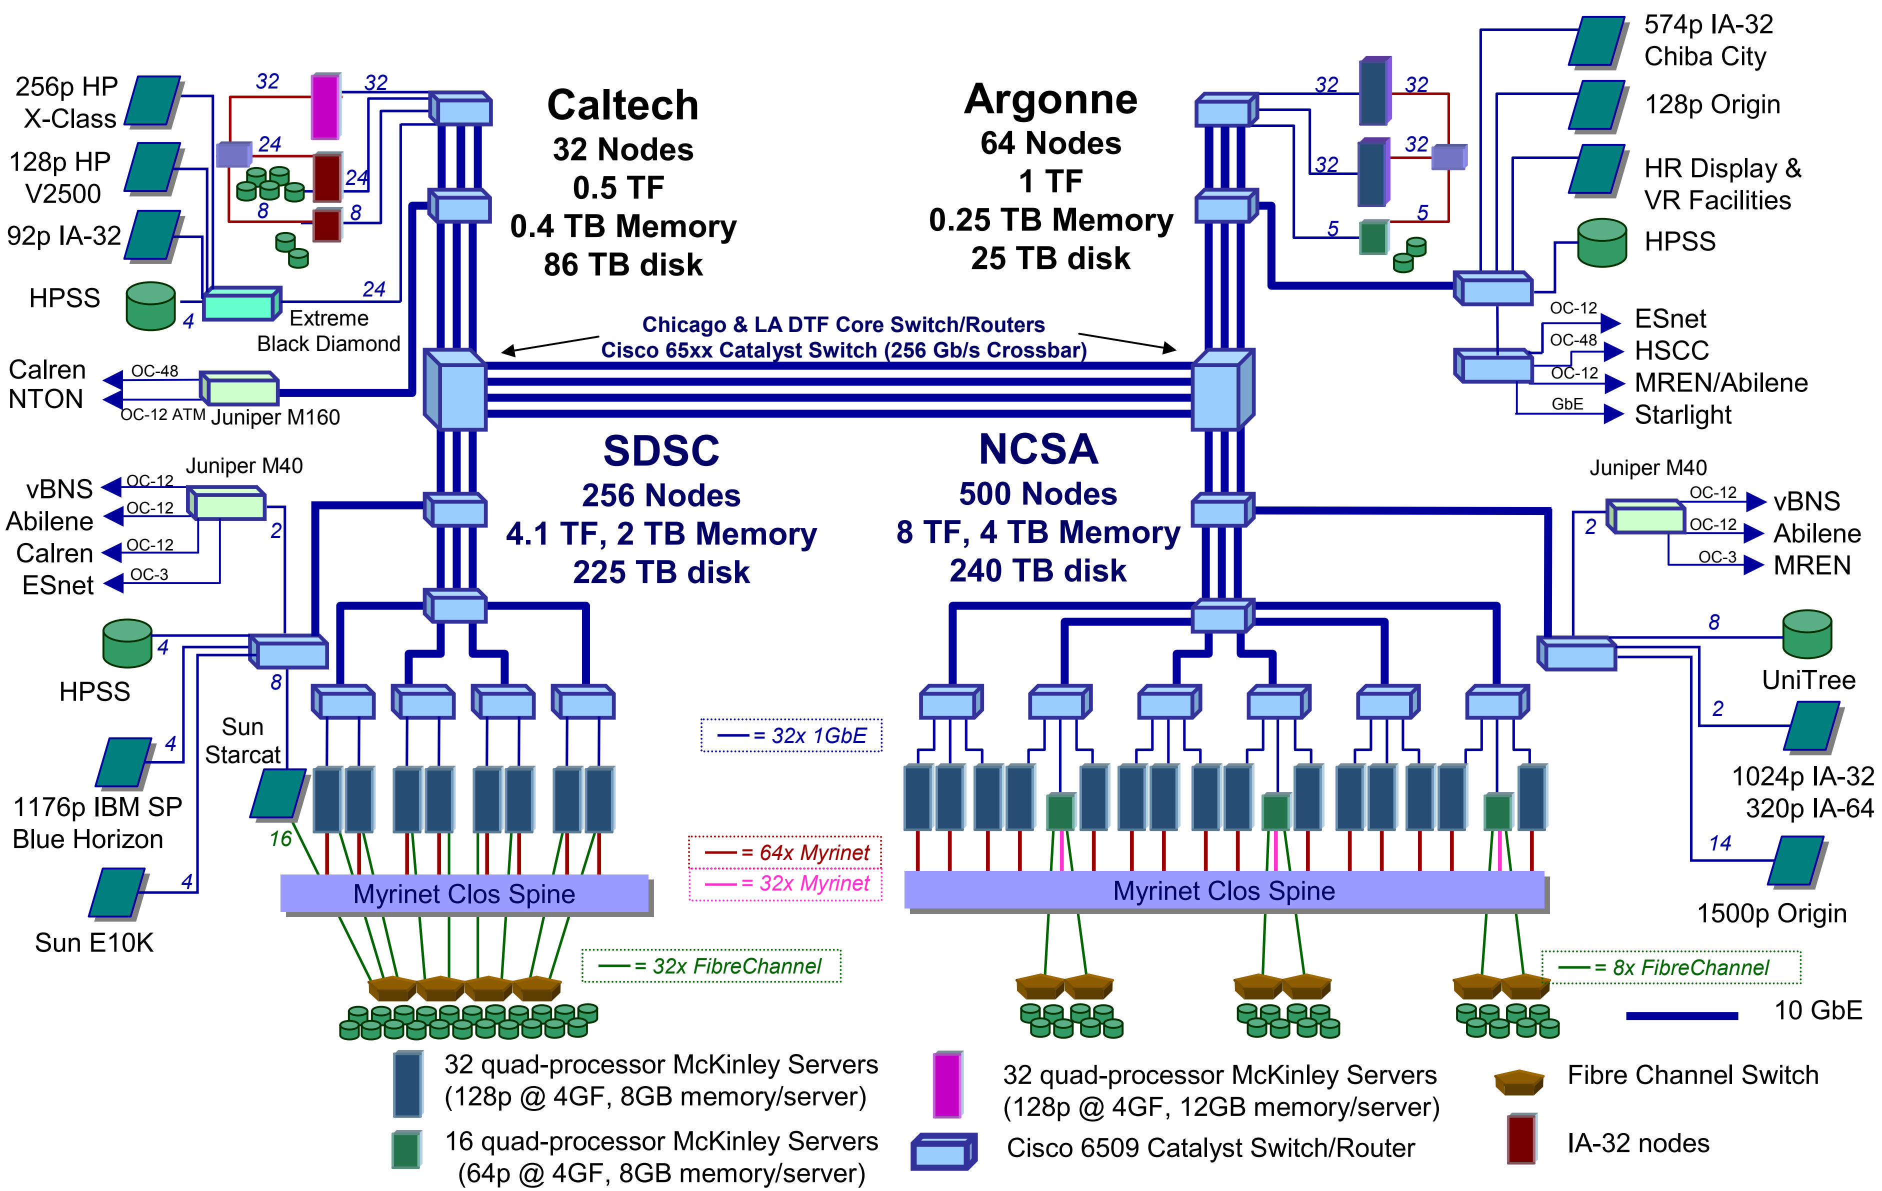

New Results Possible on TeraGrid
(due to more powerful computer ensemble)

Modeling Cell Structures

- **Pre-Blue Horizon:**
  - Possible to model electrostatic forces of a structure with up to **50,000** atoms -- a single protein or small assembly
- **Pre-TeraGrid:**
  - Possible to model **one million** atoms **–** enough to simulate drawing a drug molecule through a microtubule or tugging RNA into a ribosome
- **TeraGrid:**
  - Models of **10 million** atoms will make it possible to model function, structure movement, and interaction at the cellular level **for drug design and to understand disease**

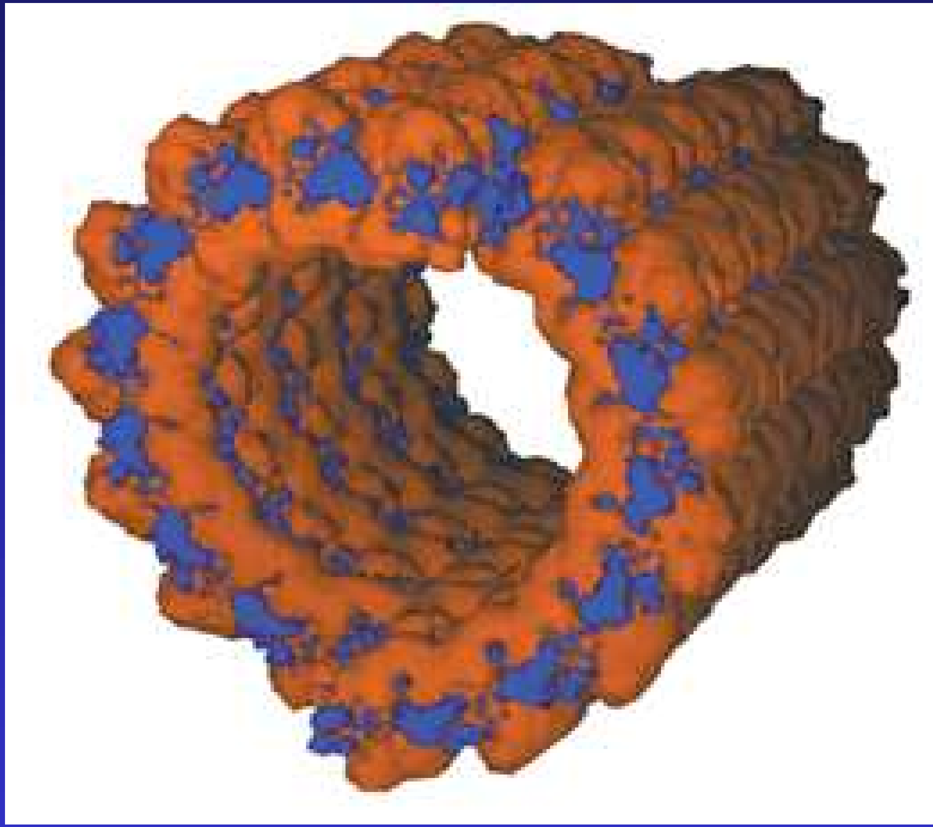

Baker, N., Sept, D., Joseph, S., Holst, M., and McCammon, J. A. *PNAS* **98**: 10037-10040 (2001).

New Results Possible on TeraGrid

(due to more storage, faster network)

- Biomedical Informatics Research Network [BIRN] (NIH grant):
  - Evolving reference set of brains provides essential data for developing therapies for neurological disorders (Multiple Sclerosis, Alzheimer's, etc.).
- Pre-TeraGrid:
  - One lab
  - Small patient base
  - 4 TB collection
- Post-TeraGrid:
  - Tens of collaborating labs
  - Larger population sample
  - 400 TB data collection: more brains, higher resolution
  - Multiple scale data integration and analysis

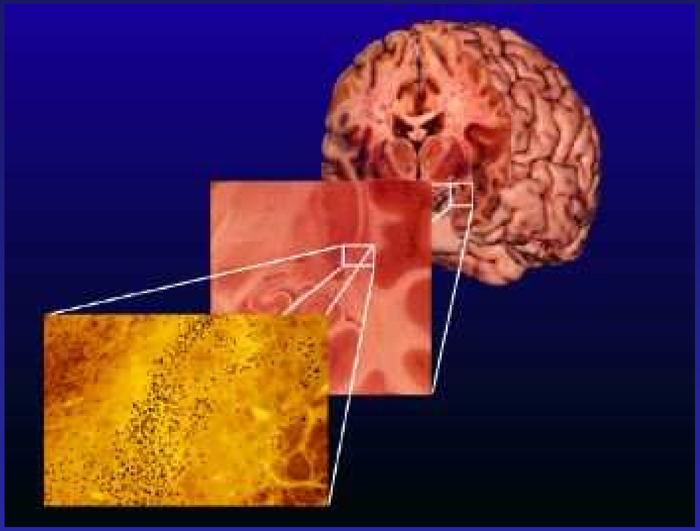

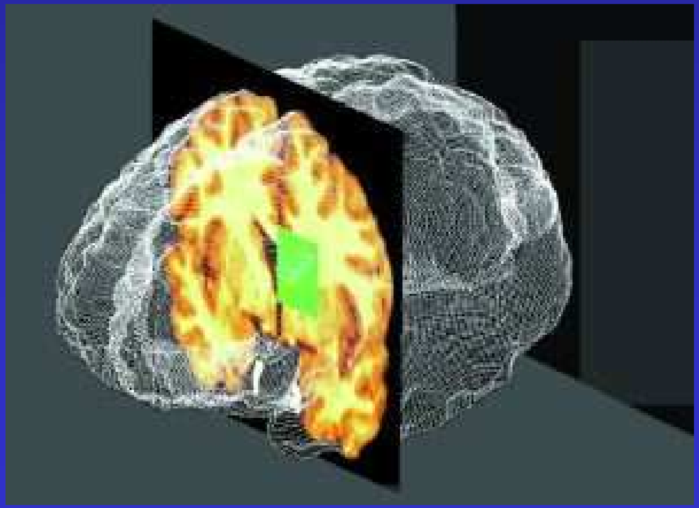

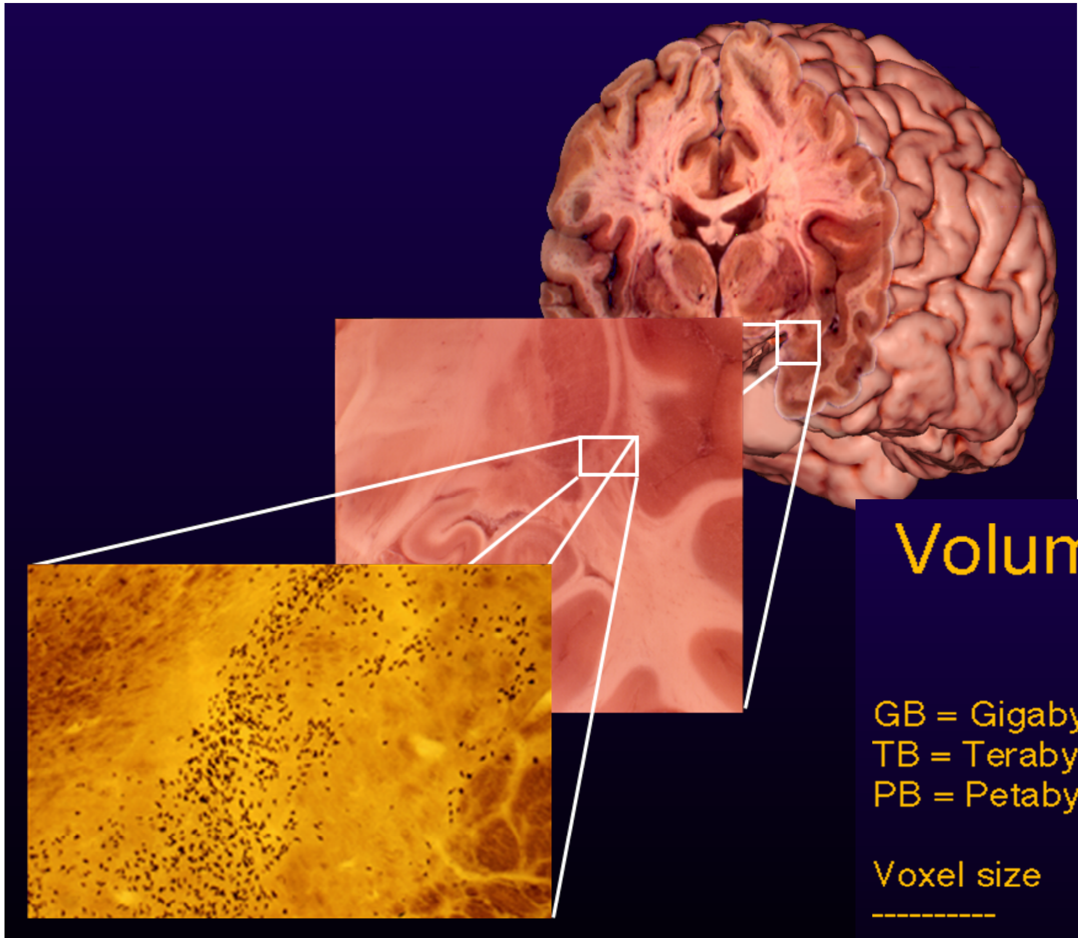

EACH BRAIN REPRESENTS A LOT OF DATA

Volume sizes by resolution - brain = 1500 cm^3

GB = Gigabyte = 10^9
TB = Terabyte = 10^{12}
PB = Petabyte = 10^{15}

Voxel size	B&W (1 B/p)	High res (2 B/p)	Color (3 B/p)
cm	1.5 KB	3 KB	4.5 KB
mm	1.5 MB	3 MB	4.5 MB
10 μm	1.5 TB	3 TB	4.5 TB
μm	1.5 PB	3 PB	4.5 PB

AND COMPARISONS MUST BE MADE BETWEEN MANY

We need to get to one micron to know location of every cell. We're just now starting to get to 10 microns – the TeraGrid will help get us there and further

Multi-disciplinary Simulations: Aviation Safety

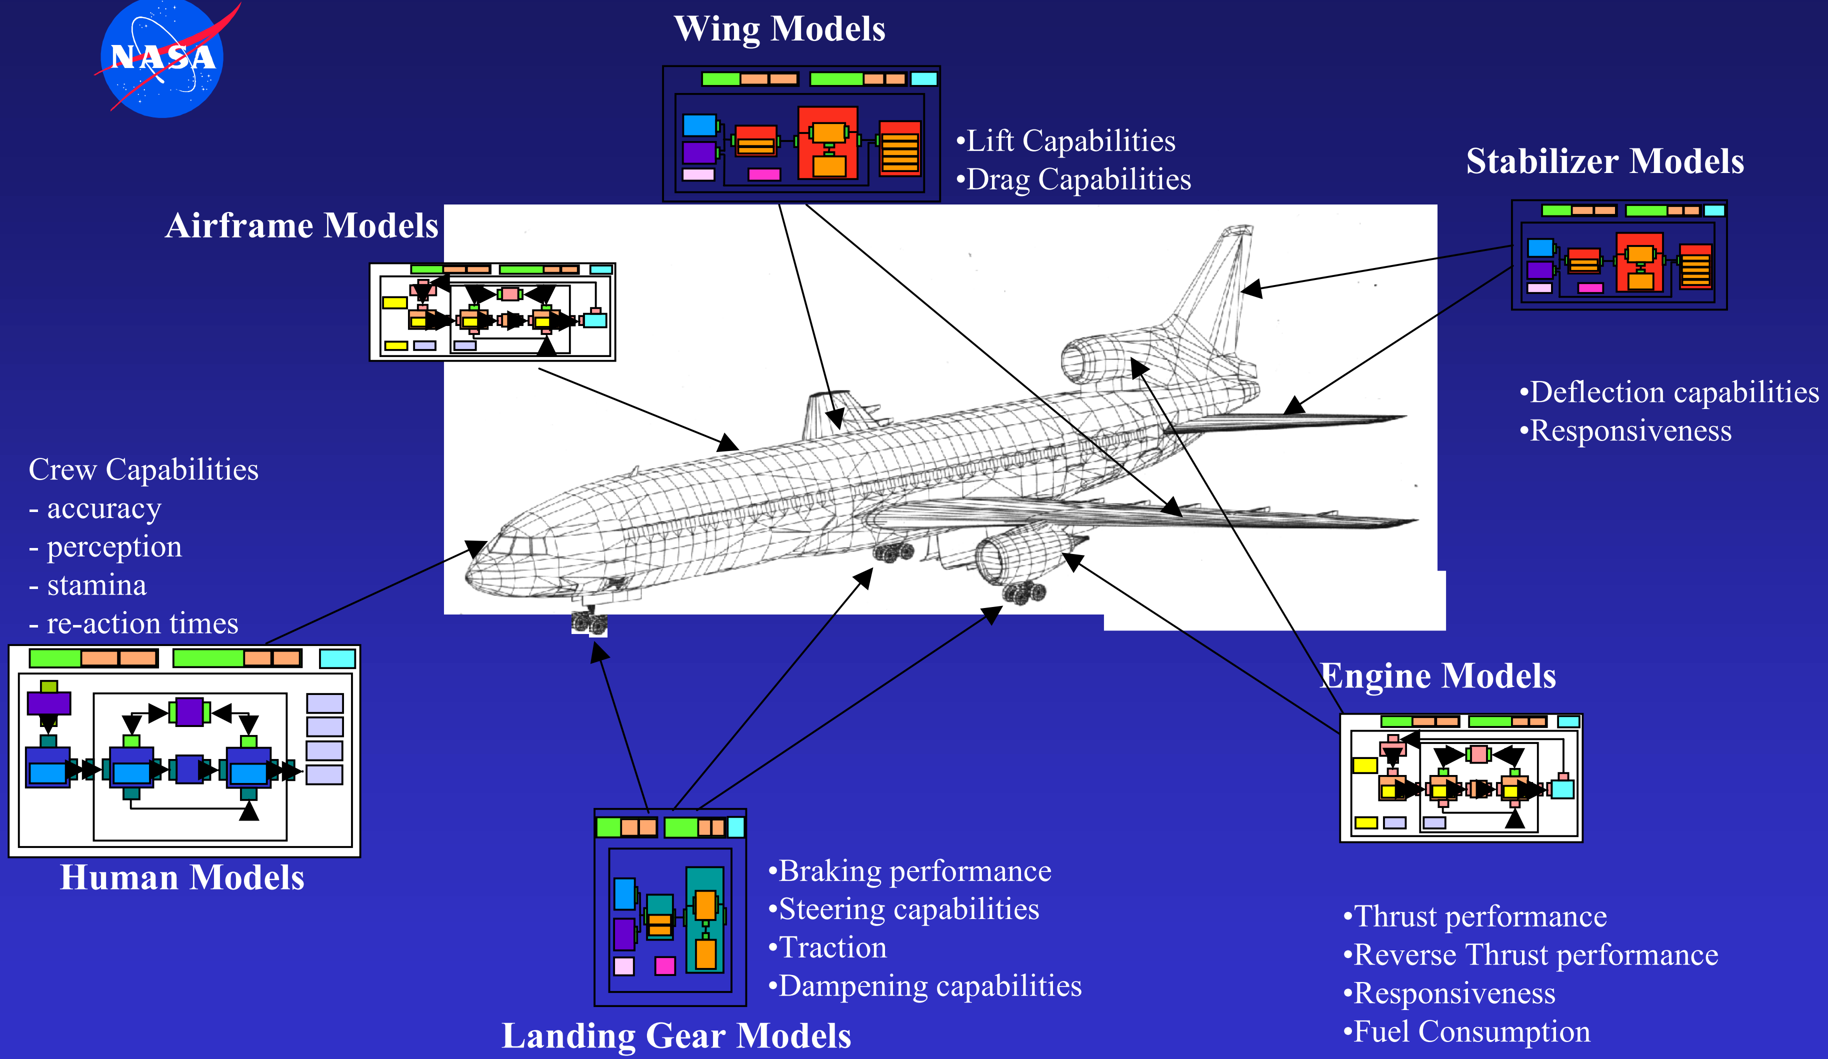

Whole system simulations are produced by coupling all of the sub-system simulations

Network for Earthquake Engineering Simulation

- NEESgrid: national infrastructure to couple earthquake engineers with experimental facilities, databases, computers, & each other
- On-demand access to experiments, data streams, computing, archives, collaboration

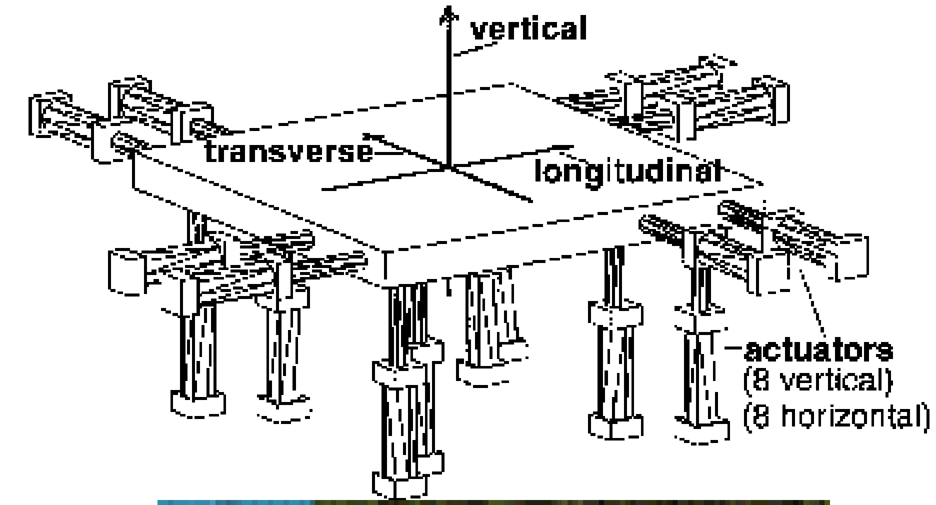

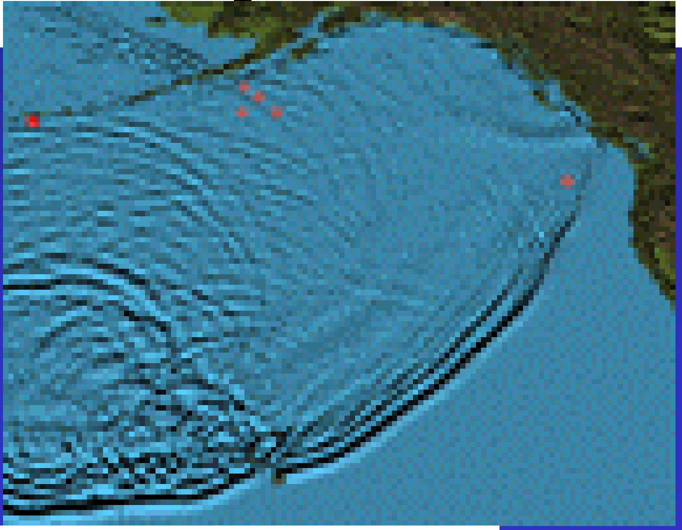

NEESgrid: Argonne, Michigan, NCSA, UIUC, USC

Digital Radiology

(Hollebeek, U. Pennsylvania)

mammograms
X-rays
MRI
cat scans
endoscopies, ...

- Hospital Digital Data
  - Very large data sources - great clinical value to digital storage and manipulation and significant cost savings
  - 7 Terabytes per hospital per year
  - dominated by digital images
- Why we chose Mammography
  - clinical need for film recall and computer analysis
  - large volume (4,000 GB/year) (57% of total)
  - storage and records standards exist
  - great clinical value to this application

Managing Large Scale Data

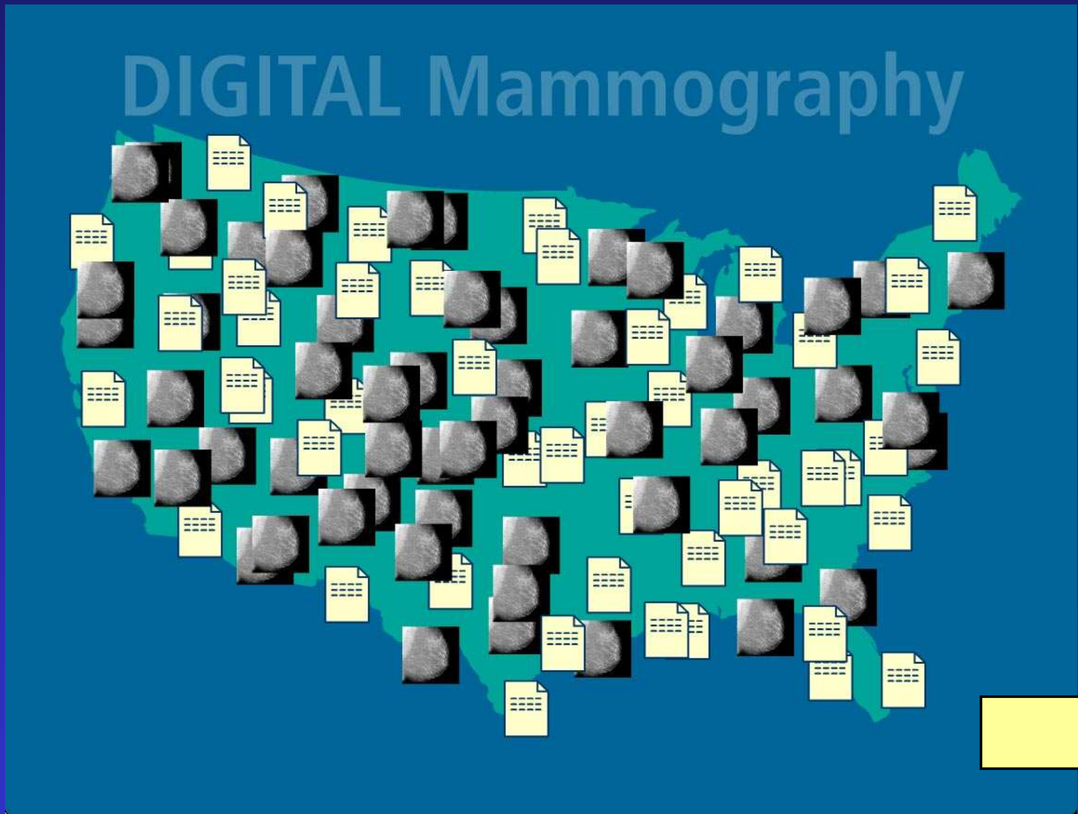

Highly Distributed Source

Hierarchical Storage and Indexing

Proposed Hierarchical Layout

Regional and Area Archives (A)

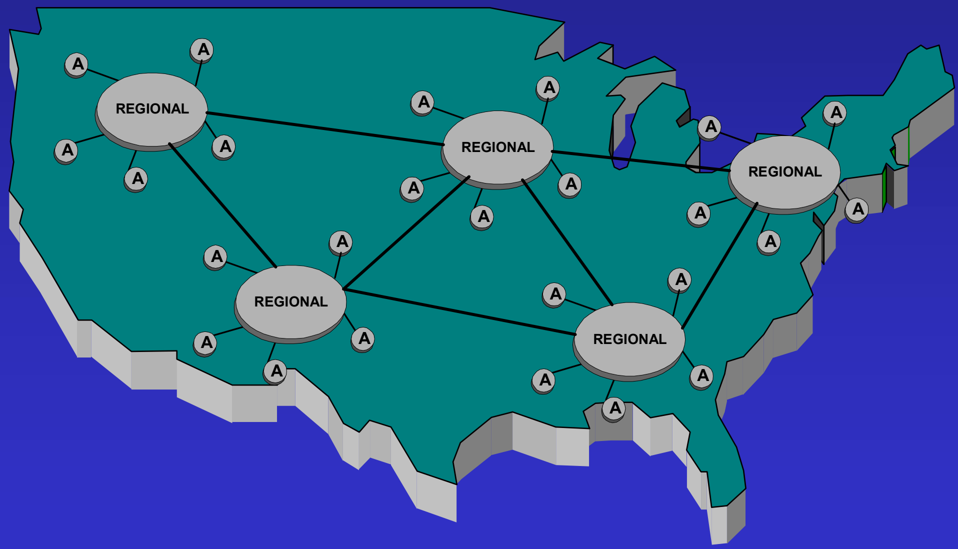

Scale of the Problem

Recent FDA approval and cost and other advantages of digital devices will encourage digital radiology conversion

- 2000 Hospitals x 7 TB per year x 2
- 28 PetaBytes per year
  - (1 Petabyte = 1 Million Gigabytes)
- **Pilot Problem scale in NDMA**
  - **4 x 7 x 2 = 56 Terabytes / year**

Data Integration and Mining: *(credit Sara Graves)* From Global Information to Local Knowledge

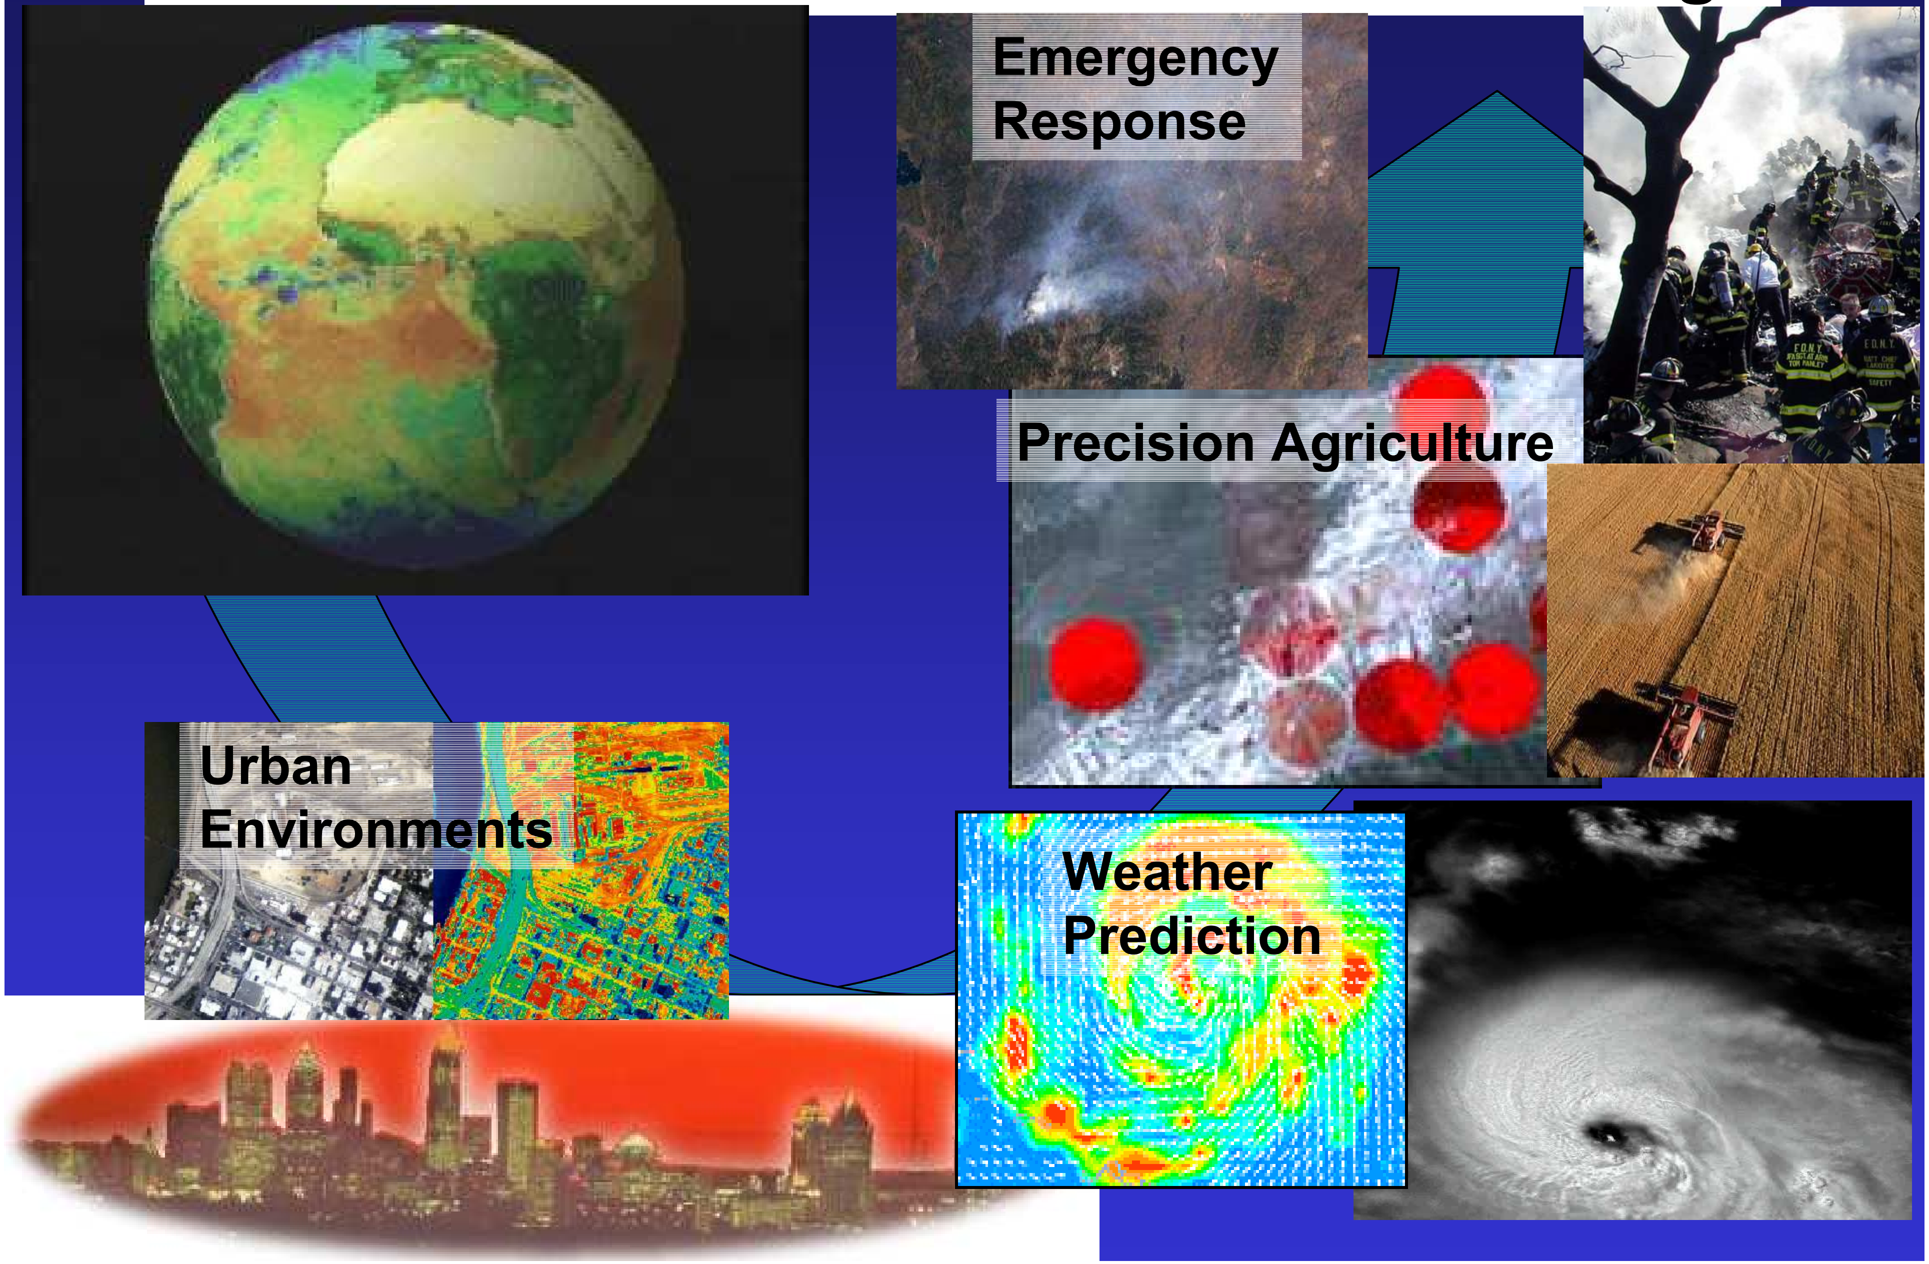

Biomedical Information Science and Technology Initiative report recommendations

- NIH should establish between five and twenty National Programs of Excellence in Biomedical Computing devoted to all facets of this emerging discipline, from the basic research to the tools to do the work
  - those National Programs will play a major role in educating biomedical-computation researchers
- To make the growing body of biological data available in a form suitable for study and use, NIH should establish a program on the principles and practice of information storage, curation, analysis, and retrieval

Biomedical Information Science and Technology Initiative report recommendations

- The NIH should provide additional resources and incentives for basic research to provide adequate support for those who are inventing, refining, and applying the tools of biomedical computing
- The NIH should foster a scalable national computer infrastructure
  - To assure that biomedical researchers can gain access to the computing resources they need beyond their desktops, NIH should provide funds to increase computing capacity, both local and remote

All kinds of biologists may want grids

- Field biologists need reliable, cheap, dependable wireless
- Field biologists need satellites for wireless communication and data collection and sharing
- Need remotely controlled sensor grids for field studies

NSF Blue Ribbon Panel

- Broad themes from BRP Testimony
  - Information Infrastructure is a fundamental **tool for science.**
  - **Grid Computing** is essential.
  - The U.S. must build a **persistent, evolutionary, grid-based National Information Infrastructure** to stay competitive and maintain its leadership role.
  - Dealing with immense amounts of **data is the critical problem** of the next decade.

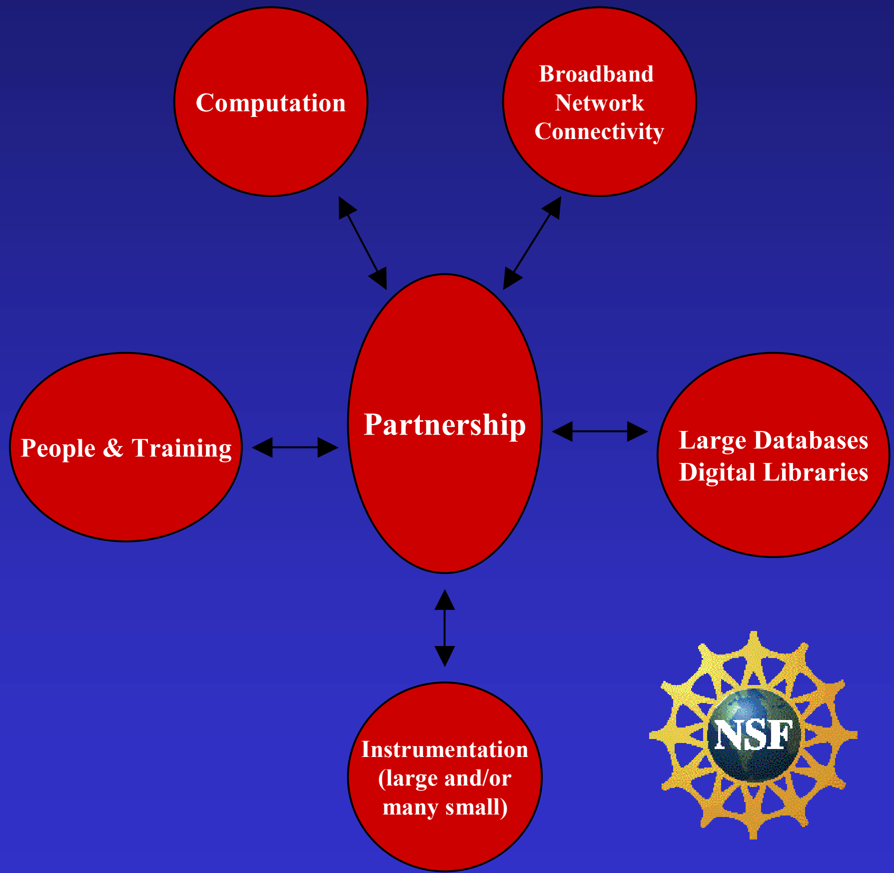

The Scientific Computing Environment of the year 2010

Thesis: Grids will be the prevalent computing environment for science

- The killer applications are there
  - biomedical
    - oops, bad choice of words for these apps
  - environmental
  - industrial
- Governments worldwide see major potential benefits, economic and societal
  - cf. last few slides
- The technologies -- software and hardware -- will continue to mature and improve

And they will be Petascale Grids

- Grids with resources measured in Peta units
  - Peta(fl)ops computers
  - Petabyte disk caches
  - Multi-Petabyte data archives
  - Terabit/s to Petabit/s network connectivity

…and then a miracle happens
(some harp arpeggios, please)

- I will wave a magic wand and assume that government funding agencies worldwide have MOUs that allow sharing of resources
  - with reasonable policies as to who has access and to level of resources allowed to use
- Research projects give timely access to their data
- GGF has successfully fostered grid standards and most institutions have adopted them
- Grid middleware and related tools are included in system software for most computers

The 2010 Petascale computing environment

- Governments and agencies explicitly fund a global grid to serve science goals
  - and resources to put on the grid
- National research network backbones run at 10 terabit/s
  - hundreds of sites connected at 1 Tb/s
  - high-end computing systems at 100s of universities and research labs
- Terabit class networks connect Europe, North America, and Asia/Pacific
- 1000s of data resources are on the global grid
  - and data collections are first class objects/resources

The 2010 Petascale computing environment

- Research groups will have private lambdas, easily (automatically?) dedicated to specific tasks, thus providing predictable BW
- There will be curated scientific data archives, supported indefinitely by governments who wisely realize that the data they paid so much to obtain can have many uses and a long useful lifetime
- The contents of all issues of scientific journals more than 6 months old will be freely available online

Grids and Petaflops

- Aggregate power of systems on grids will certainly be many petaflops
  - for some applications, these will be useful petaflops
- But some applications need single systems with petaflops power
  - some general purpose, some special purpose (GRAPE VIII??!!)
- The grid will make those petaflops machines easier to use
  - as "just another resource" on the grid, petaflops systems will have access to suitable ancillary resources
  - networks will be fast enough to move data between system and remote resources
- Both types of petaflops systems needed by the scientific community

A Unifying Infrastructure

- Easier to integrate specialized resources into computing/information environment
  - peta- and hexa-flops systems
  - signal processors
  - advanced visualization equipment
  - sensors
  - exotic technology (molecular, quantum, etc.) computers

A Unifying Infrastructure

- A long time ago I said that high-speed networks would provide the “links” that will truly tie theory, experiment, and computation as research methodologies
- I was wrong -- it’s the *Grid* that is doing that
- This aspect may in the end be the most important contribution, for the scientific community
- Models, vast data collections, and computations are now all available in concert or for comparison

The Petascale Grid will facilitate better science

- Validation of simulations will be easier
  - data from experiments will "all" be readily available
  - use of alternate simulation programs for comparison will be easier
- Collaboration with remote colleagues will be more pleasant
  - sorry, current videoconferencing not good enough
- Consultation of experts will be easier
- Medical treatments will be based on much better statistical information, fundamental knowledge, and human expertise

Additional potential benefits

- Democratization (ugh) of scientific research
  - (closer to equal) access to
    - first-rate research infrastructure (including experimental)
    - data
    - publications
- (more) science-based policy making
  - especially for environmental issues?
- Users will have ~1 computing environment to learn to use
  - and complexity of use linearly related to complexity of task
- Better support of truly multidisciplinary research
  - not just CS and physics, also math, chem, bio, engineers

In other words

- With the petascale grid, the vision of the applications that I mentioned will be realized!

Summary

- Grids will make “more real” the emergence of computation as the third methodology for scientific research, alongside experiment and theory
- They will expand the pool of people who can do forefront research
- They will leverage investments in research infrastructure
- Grids open new opportunities for research

This slide intentionally left blank