Agents in an Exponential World

Ian Foster

Computation Institute

Argonne National Lab & University of Chicago

http://www.ci.uchicago.edu

Context

"People tend to overestimate the short-term impact of change, and underestimate the long-term impact"

— Roy Amara

"The future is here, it is just not evenly distributed."

— William Gibson

Fundamental Dynamics
(Pace of change of the digital infrastructure)

Digital power =
computing ⊗ communication ⊗ storage ⊗ content

computing	communication	storage	content
Moore's law	fiber law	disk law	community law
doubles every 18 months ⊗	doubles every 9 months ⊗	doubles every 12 months ⊗	2^n where n is # people

(Thanks to John Seely Brown)

Exponential Growth

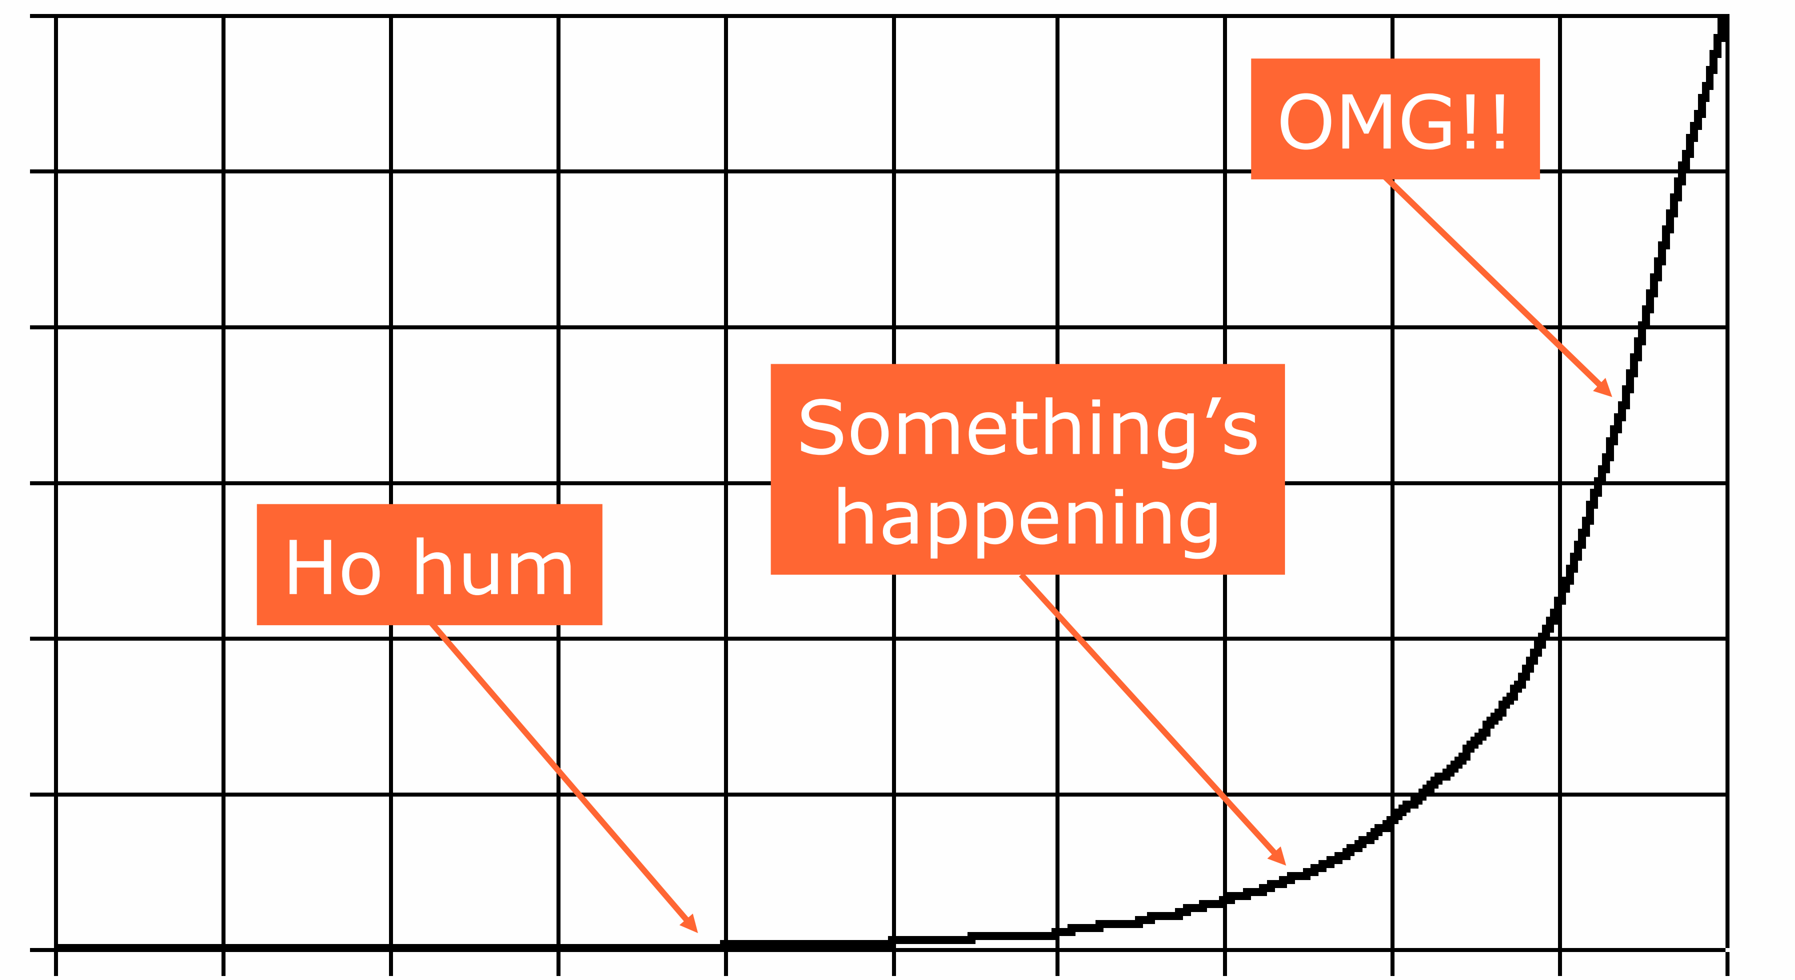

Relevant Trends

- Computing
  - High performance via massive parallelism
- Communications
  - Doping and optical switching
- Storage
  - Enormous increases in available data
- Community
  - Social networks in many forms

Projected Performance Development

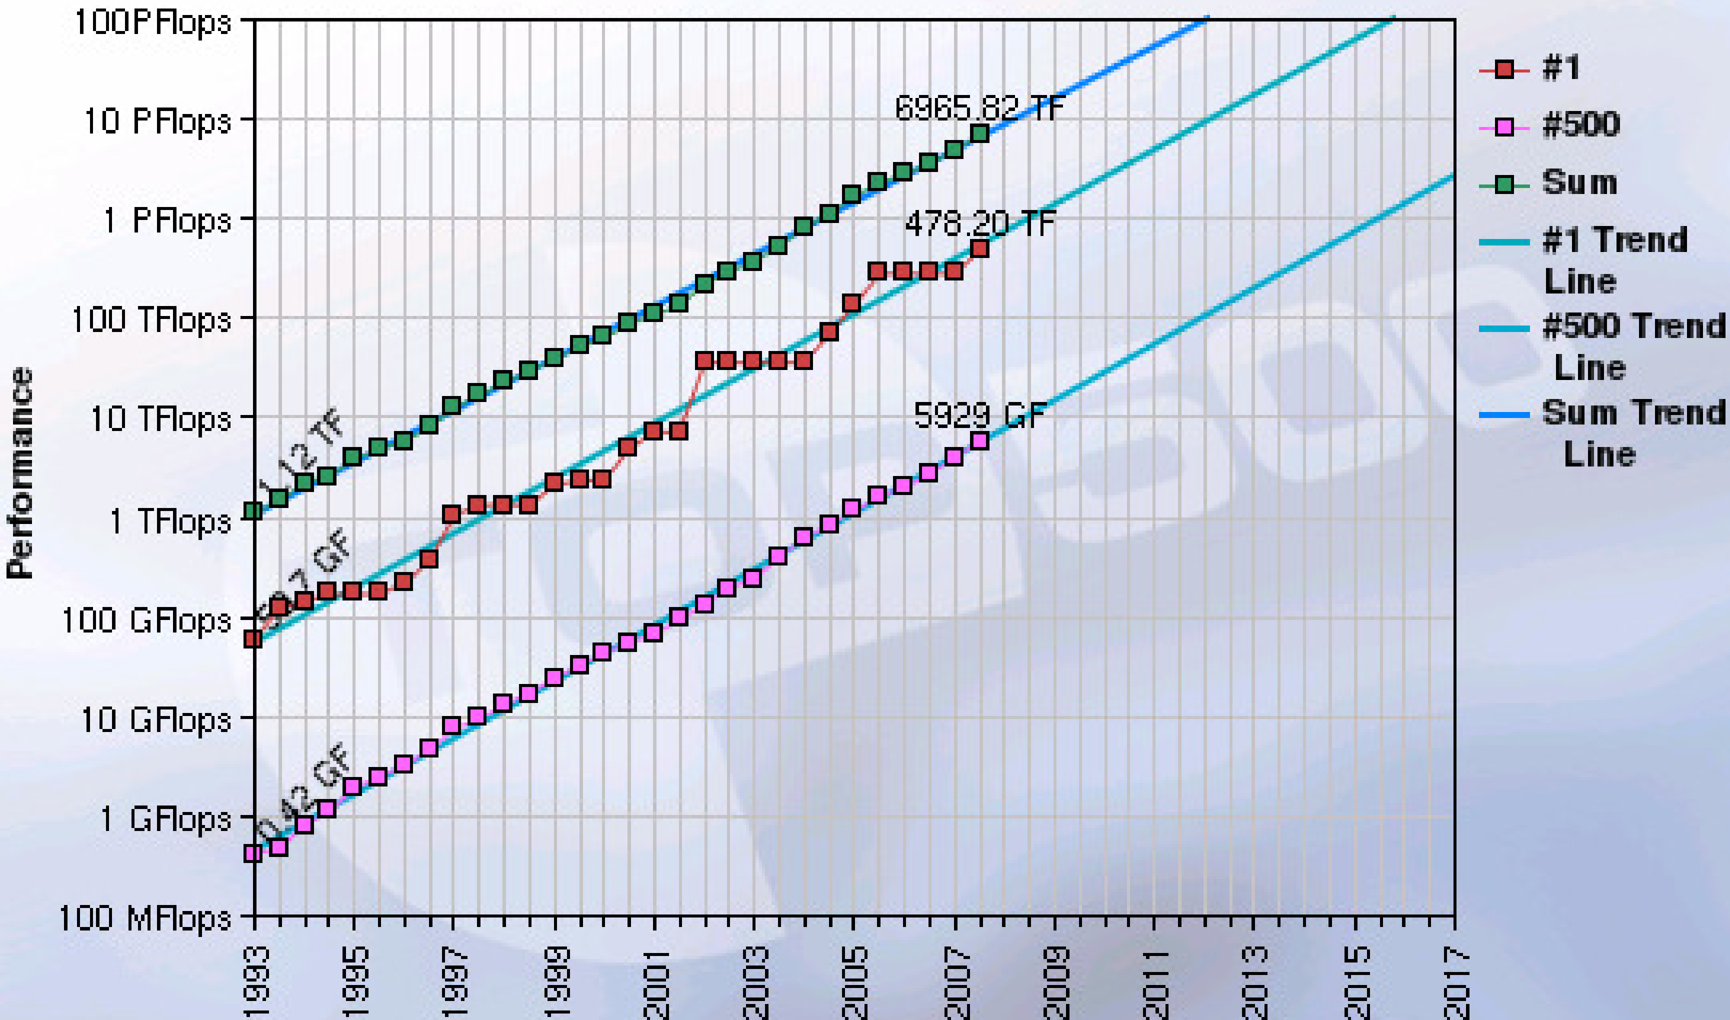

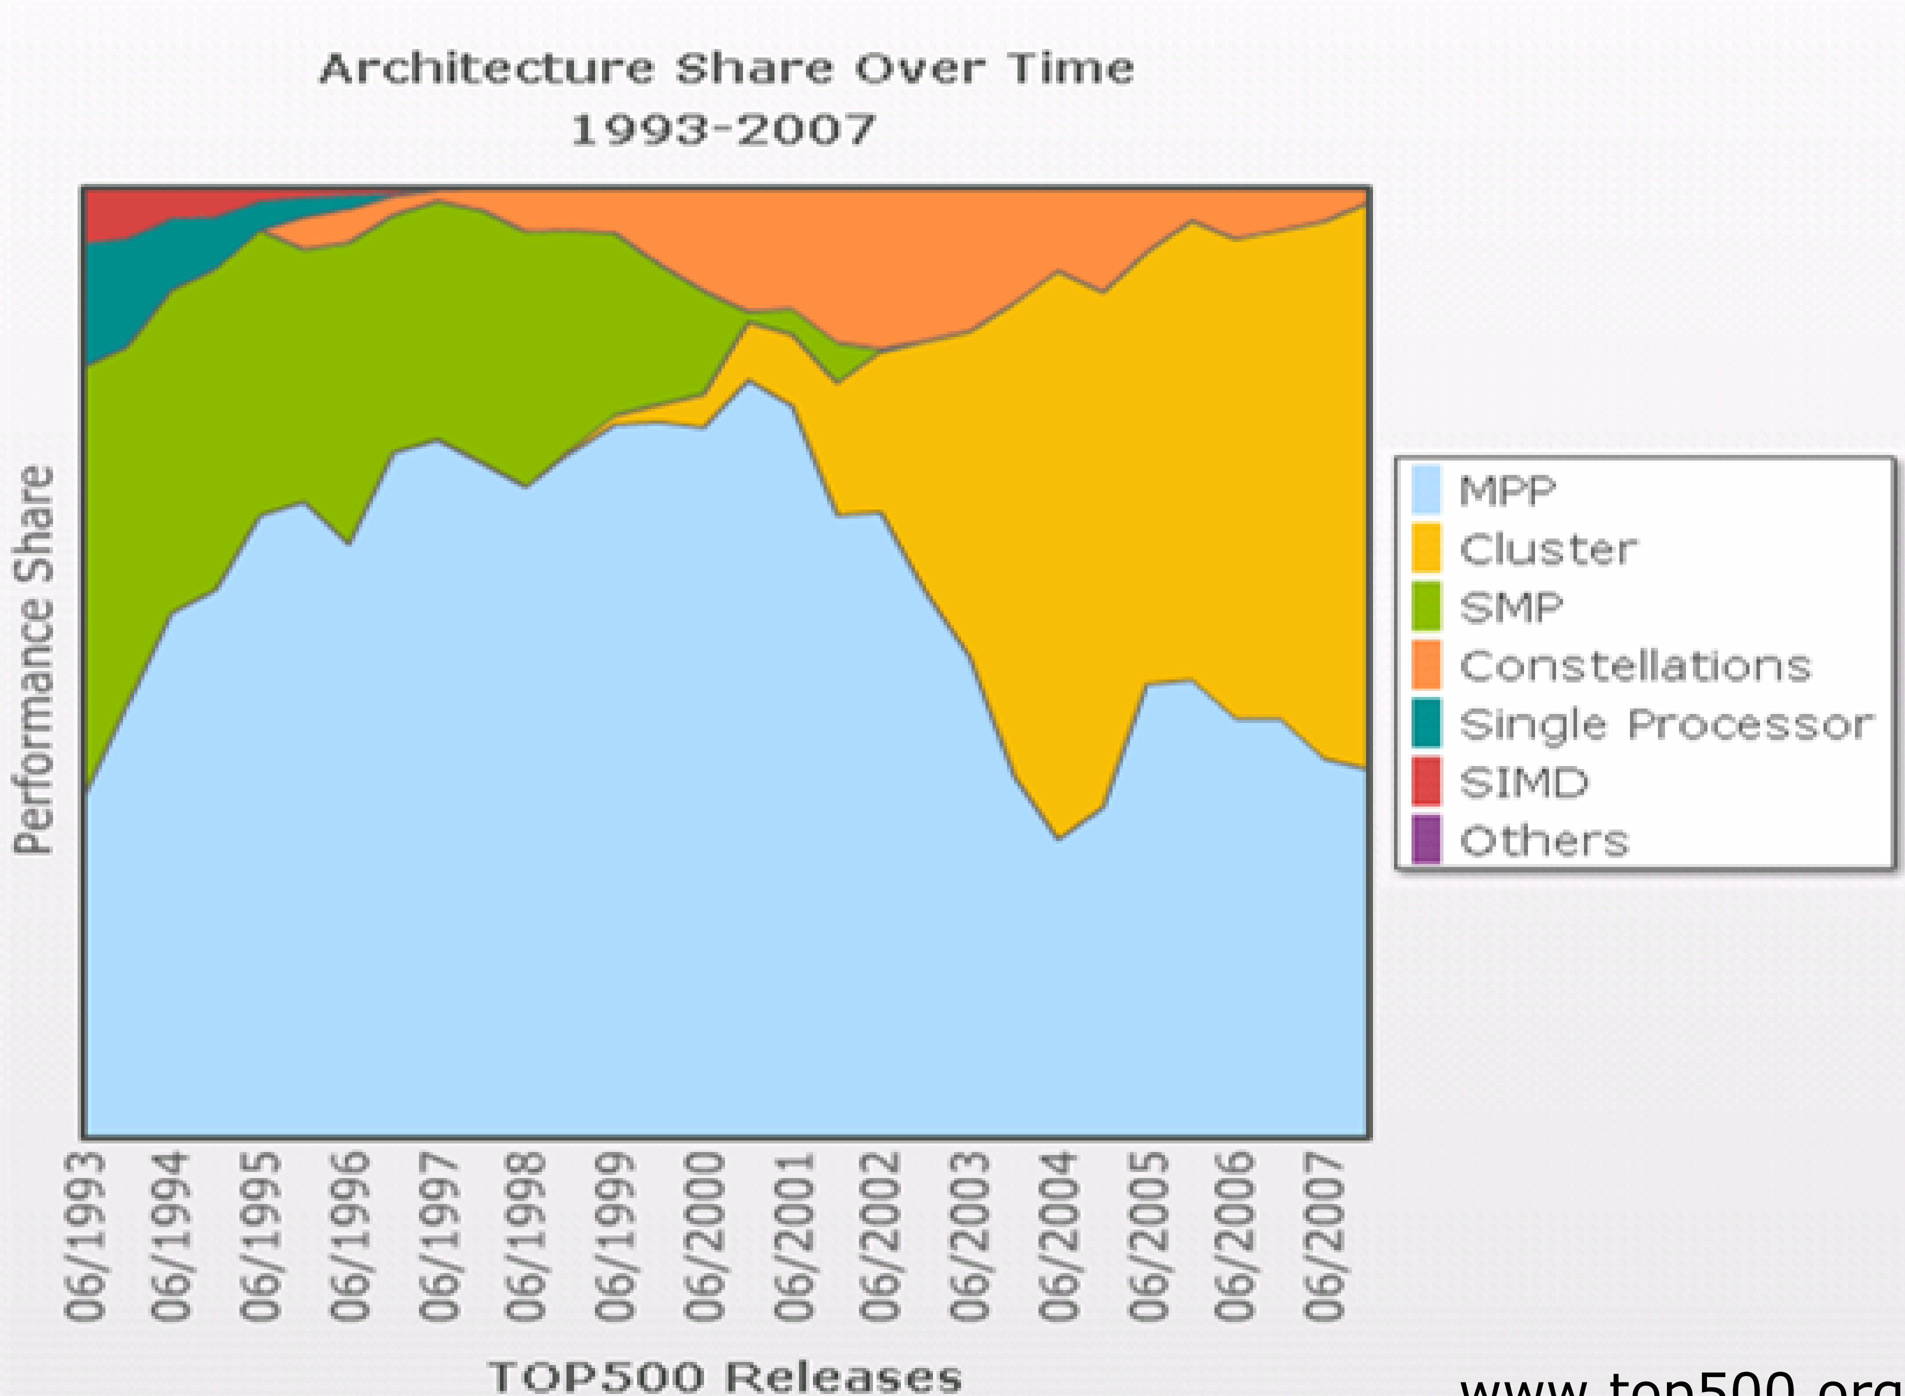

www.top500.org

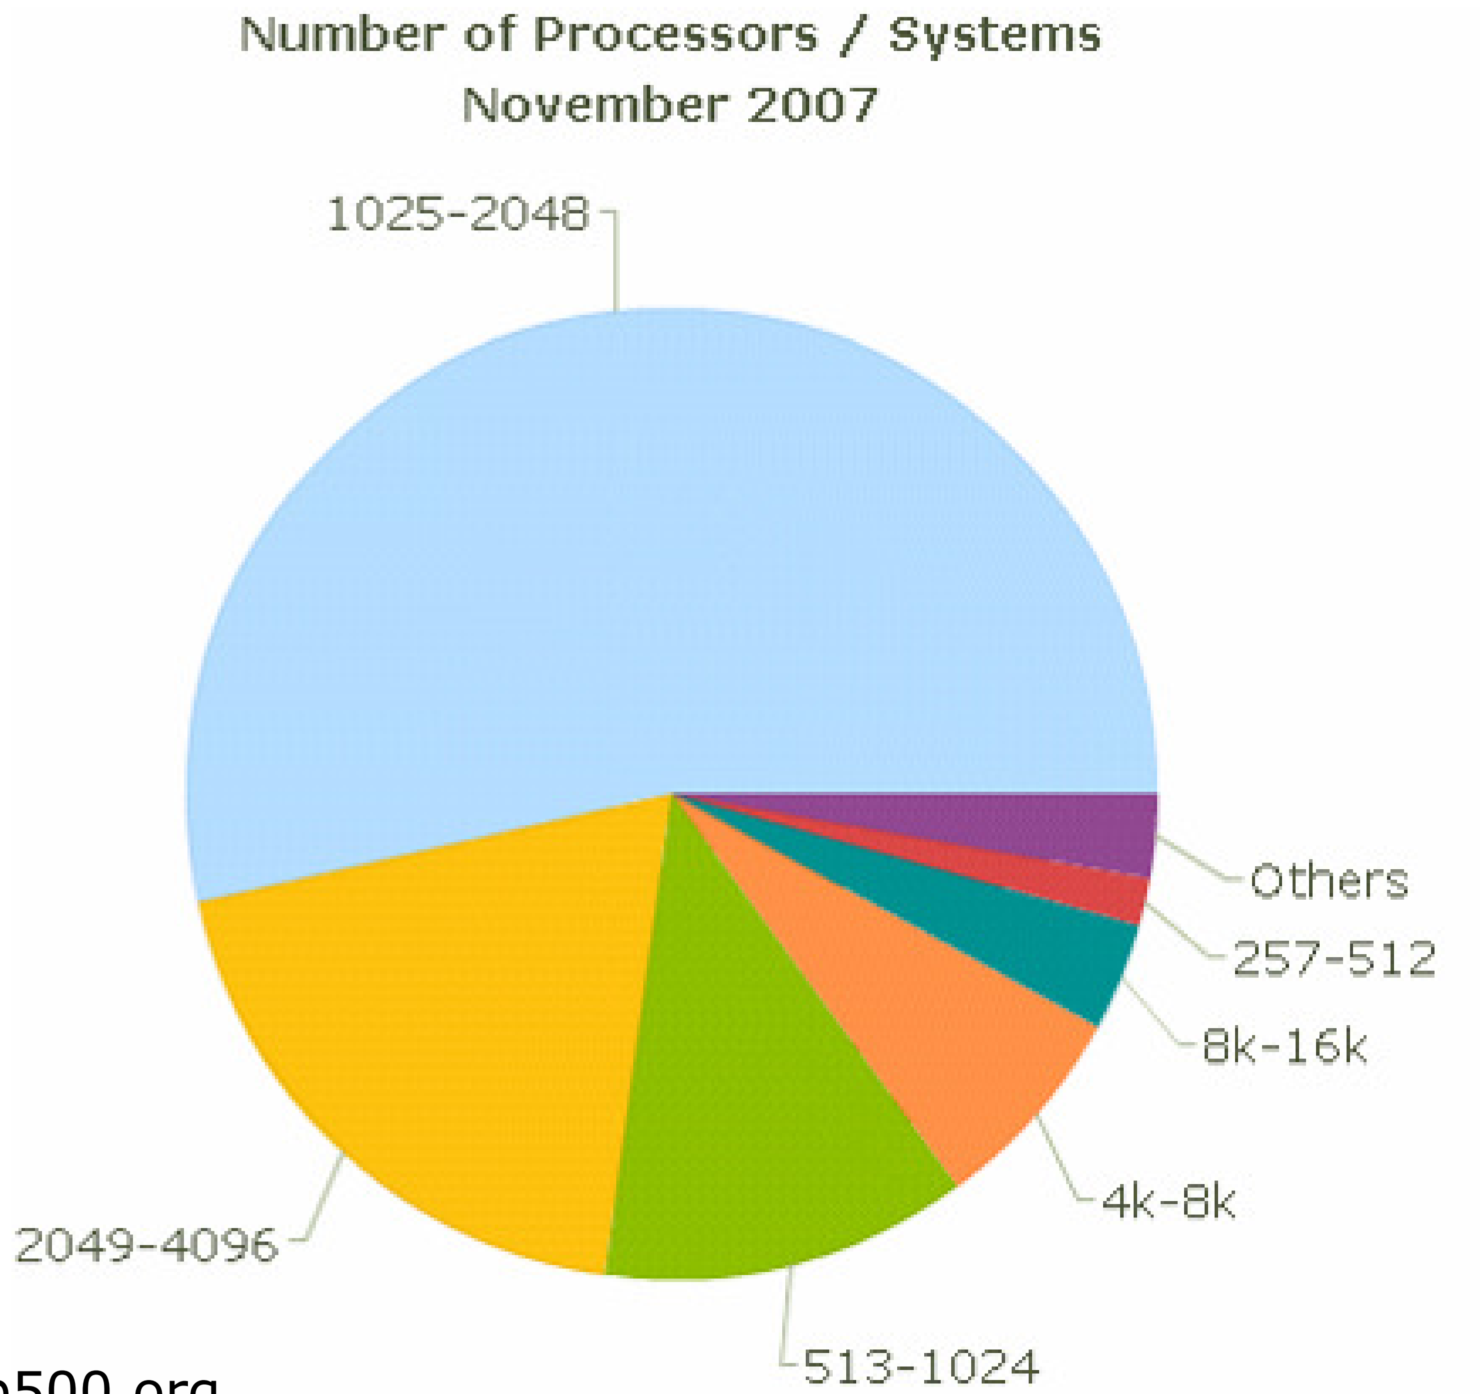

www.top500.org

1974: IBM 370/195

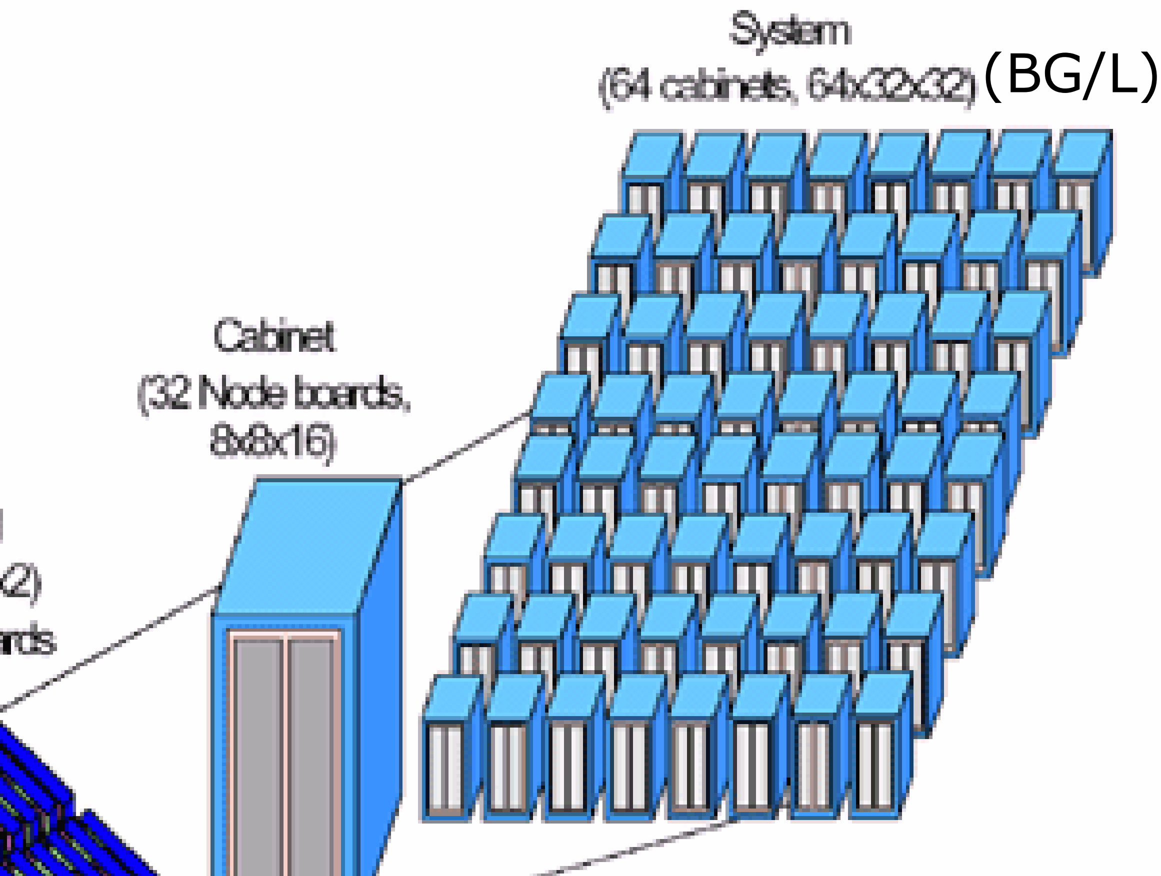
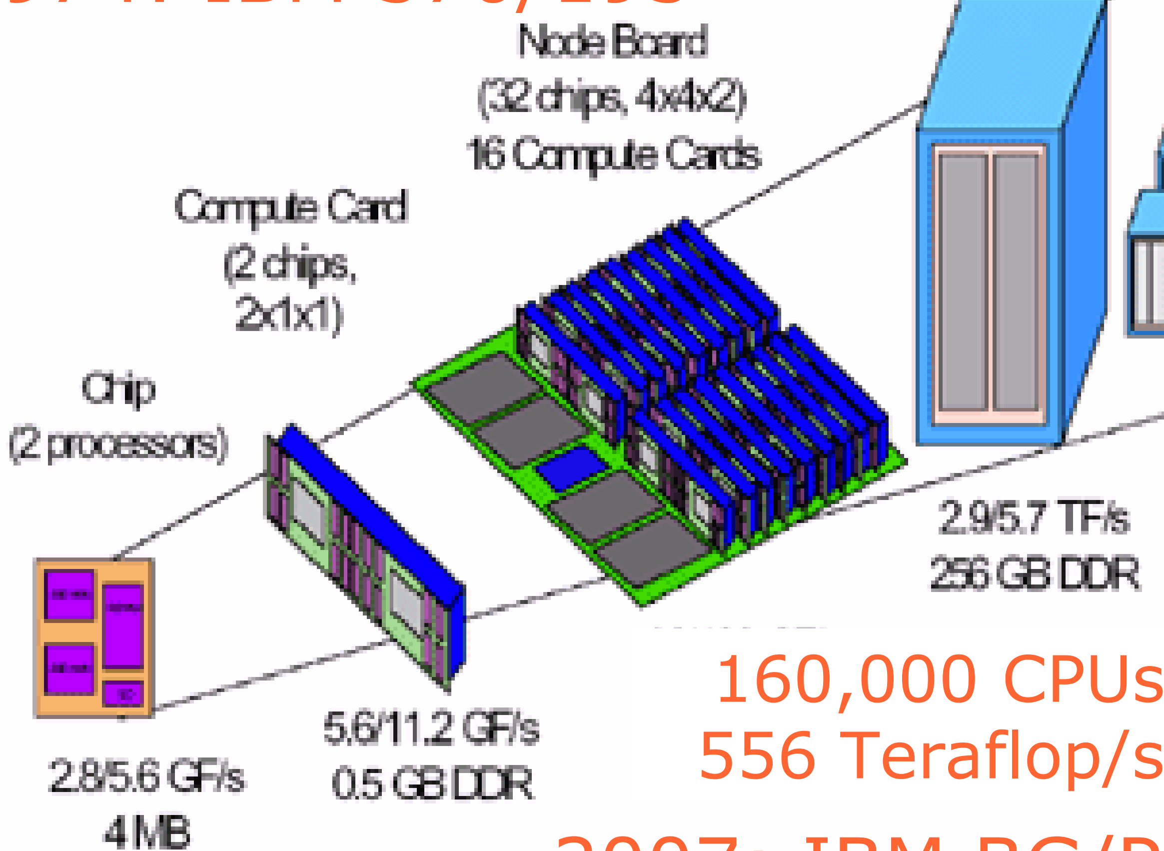

160,000 CPUs
556 Teraflop/s

2007: IBM BG/P

0.5 TeraFlop/s, 1500 Watts

SiCortex SC648, powered by 8 Trek bicycles @ ~260 Watts each

Microprocessor Trends

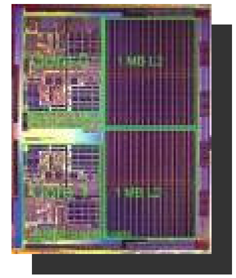
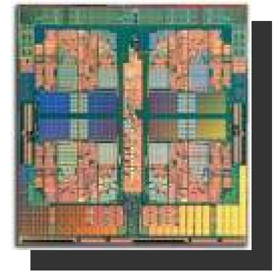

- AMD
  - Dual core (April 2005)
  - Quad core (October 2007)
- Intel
  - Dual core (July 2005)
  - Quad core (December 2006)

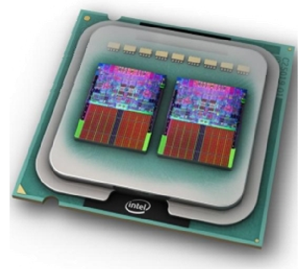

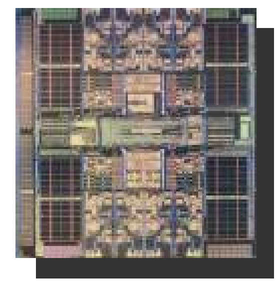

- SUN
  - Niagara: 8 cores * 4 threads/core (November 2005)
  - Niagara2: 8 cores * 8 threads/core (August 2007)
- IBM POWER6
  - 2 cores * 4 threads/core (May 2007)
- Tilera 64 cores

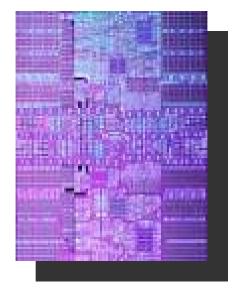

(Slide credit: Dan Reed)

Marching Towards ManyCore

- Intel's 80 core prototype
  - 2-D mesh interconnect
  - 62 W power
- Tilera 64 core system
  - 8x8 grid of cores
  - 5 MB coherent cache
  - 4 DDR2 controllers
  - 2 10 GbE interfaces
- IBM Cell
  - PowerPC and 8 cores

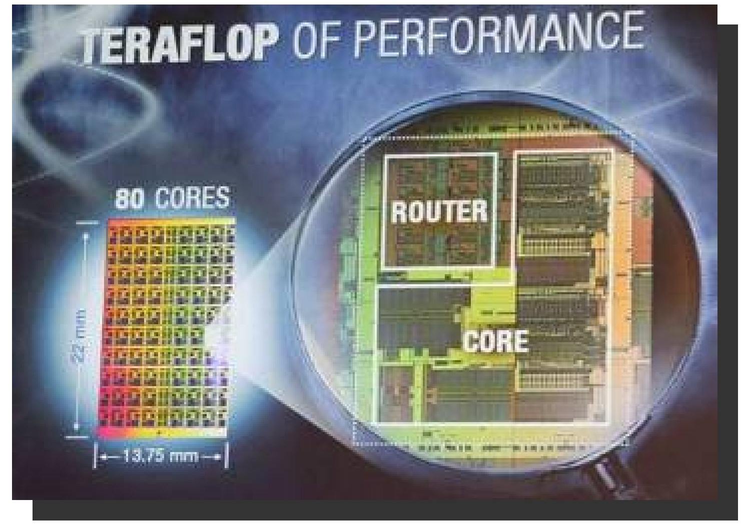

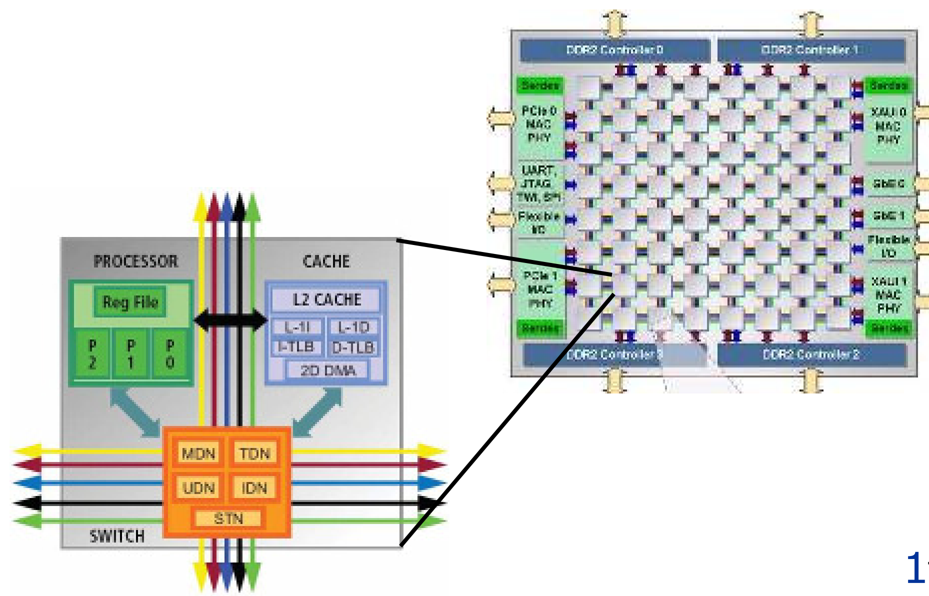

(Slide credit: Dan Reed)

Gamers Want Manycore

"I wish to have no connection with ships that do not sail fast, for I intend to go in harm's way."

John Paul Jones

(Slide credit: Dan Reed)

Criminals control 4 Petaflops supercomputer

Filed under HPC by Andy | 2 comments

09.25.2007

Zdnet reports that the botnet created by the Storm Worm trojan is a more powerful distributed supercomputer than any listed on the Top500. The botnet created by Storm Worm trojan is estimated to have captured between 1 and 10 million CPUs. Taking a conservative 2GHz processor, this adds up to at least 4 Petaflops of computing power in the hands of the criminals behind this operation.

Folding@home project nabs Guiness World Record title, thanks to the PS3

Posted Oct 31, 2007 at 09:02AM by Sally B.
Listed in: PlayStation 3
Tags: Sony, Stanford University, Folding@Home, Vijay Pande

Thanks to the ginormous help afforded by Sony PlayStation 3, Stanford University's Folding@home project was eventually recognized by the Guinness World Records as the World's Most Powerful Distributed Computing network. This big achievement is thanks to the overwhelming participation of PS3 owners all around the world (QJ included).

Folding@home (also known as F@h) managed to reach the one petaflop mark on September 16 last month, while continued widespread participation of PS3 users enabled the PS3 to surpass one petaflop on September 23, not counting the input from other computers and devices.

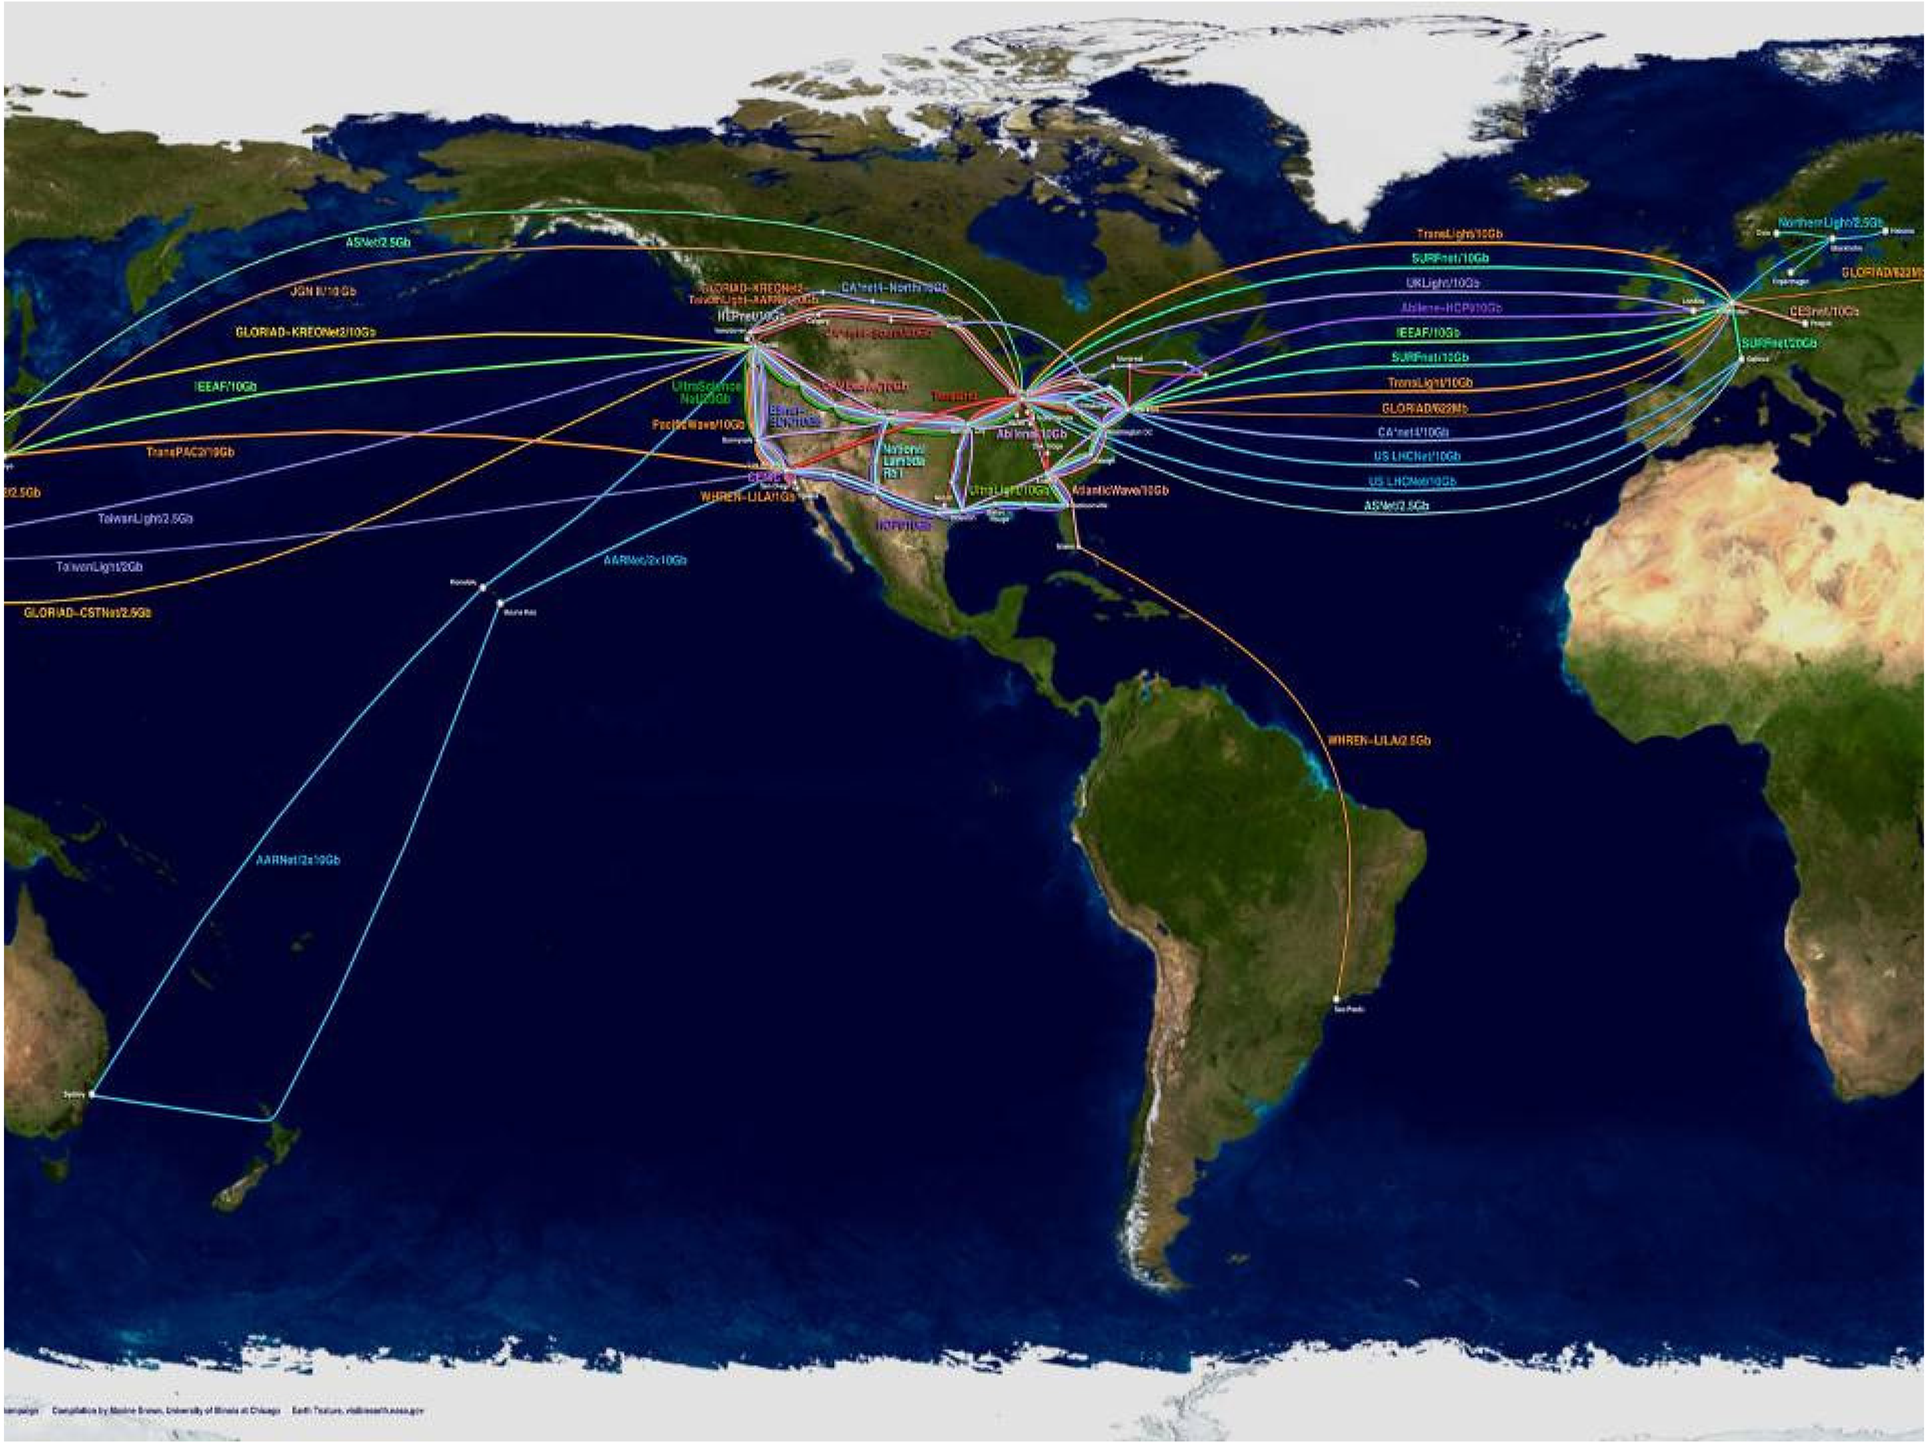

ASNet/2.5Gb
JGN II/10Gb
GLORIAD–KREONet2/10Gb
IEEAF/10Gb
TransPAC2/10Gb
TaiwanLight/2.5Gb
TaiwanLight/2Gb
GLORIAD–CSTNet/2.5Gb
CA*net4–North/10Gb
UltraScience Net/20Gb
PacificWave/10Gb
National LambdaRail
WHREN–LILA/10Gb
AARNet/2x10Gb
HOPI/10Gb
Abilene/10Gb
UltraLight/10Gb
AtlanticWave/10Gb
TransLight/10Gb
SURFnet/10Gb
UKLight/10Gb
Abilene–HOPI/10Gb
IEEAF/10Gb
SURFnet/10Gb
TransLight/10Gb
GLORIAD/622Mb
CA*net4/10Gb
US LHCNet/10Gb
US LHCNet/10Gb
ASNet/2.5Gb
NorthernLight/2.5Gb
CESnet/10Gb
SURFnet/20Gb
WHREN–LILA/2.5Gb

Optical Networking Breakthrough!

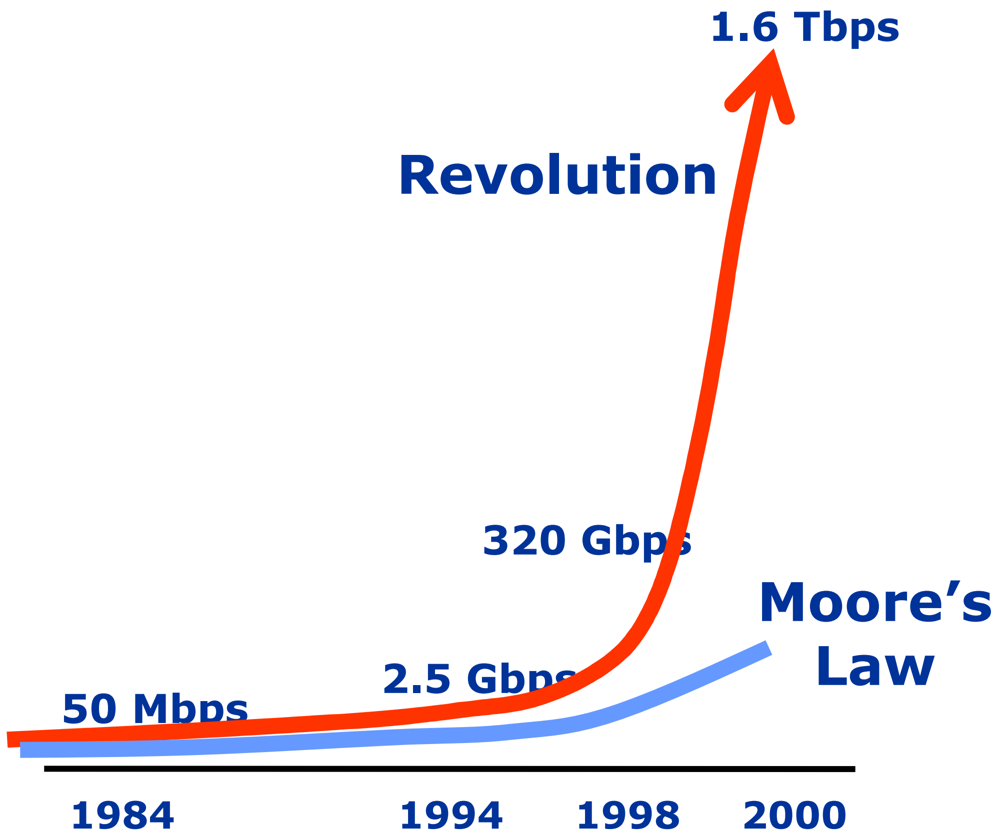

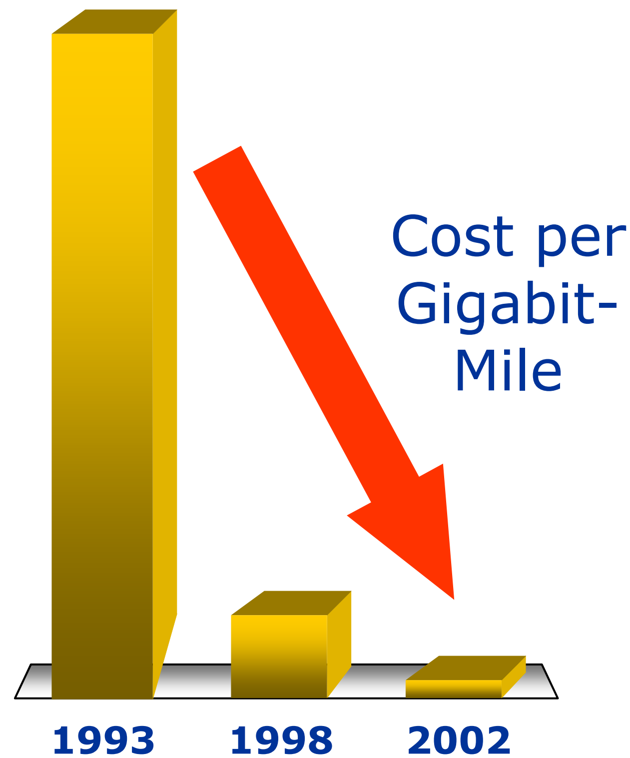

Capacity increase and new economics

Source: Nortel

Optical Switches

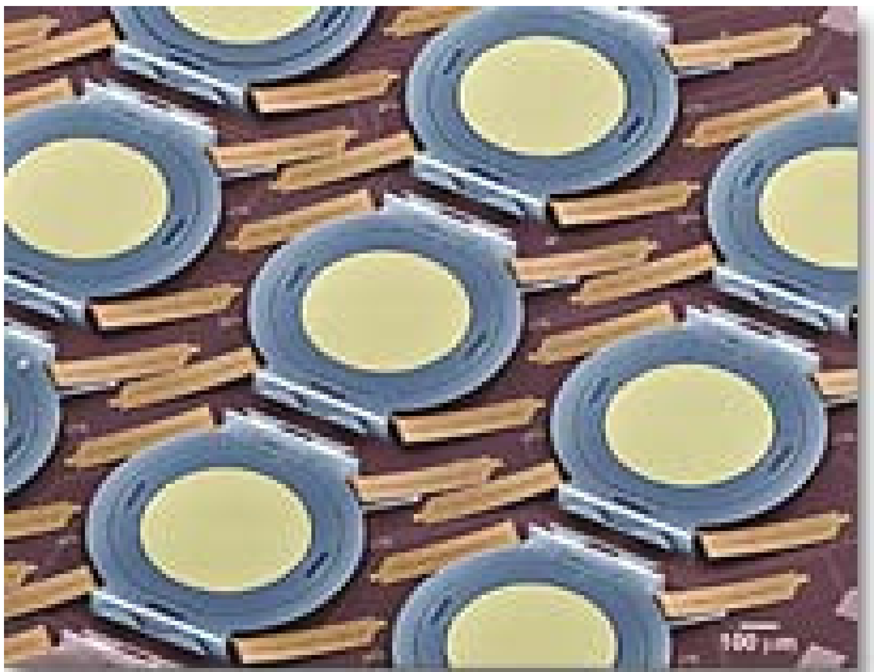

Lucent

Redirecting Light with Tiny Mirrors

Researchers have developed optical switches that use tiny electro-mechanical mirrors to steer light. Such devices would allow fiber-optic networks to handle higher volumes of data and voice traffic than would be possible using electronic components. Here is a MEMS design used by Lucent Technologies.

WHAT IT DOES

Called an optical cross connect, it allows light from any one of 256 optical input fibers to be routed to any one of 256 output fibers.

HOW IT WORKS

1. Light coming in through an optical fiber is focused at a MEMS mirror. Tipped at a slight angle, the mirror redirects the light to another MEMS mirror and out a different fiber.
2. By changing the angle of the MEMS mirrors, the light is sent in a different direction and out a different fiber.

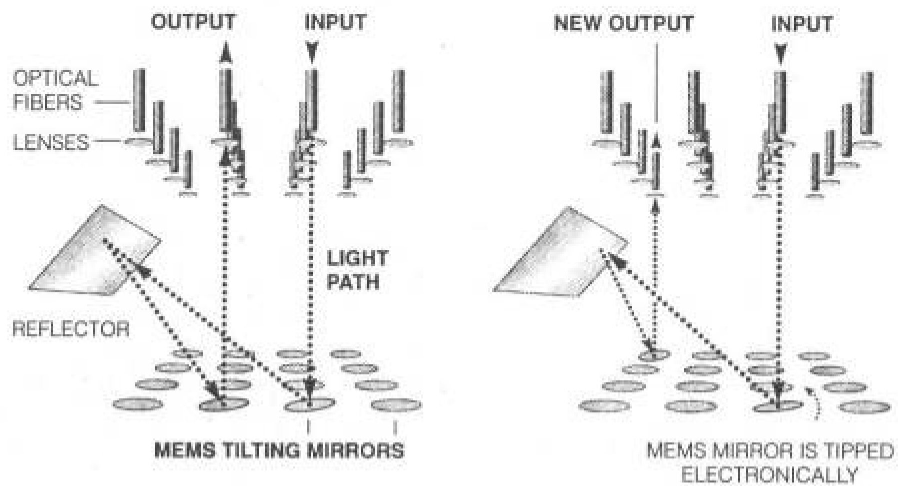

ADVANTAGES

Currently, switches must convert light into an electronic digital signal, redirect it, then convert it back into light, a process that is slower and vulnerable to signal loss. In addition, optical switches use less power and can be expanded without replacing existing equipment.

Source: Lucent Technologies' Bell Labs

Globus Downloads Last 24 Hours

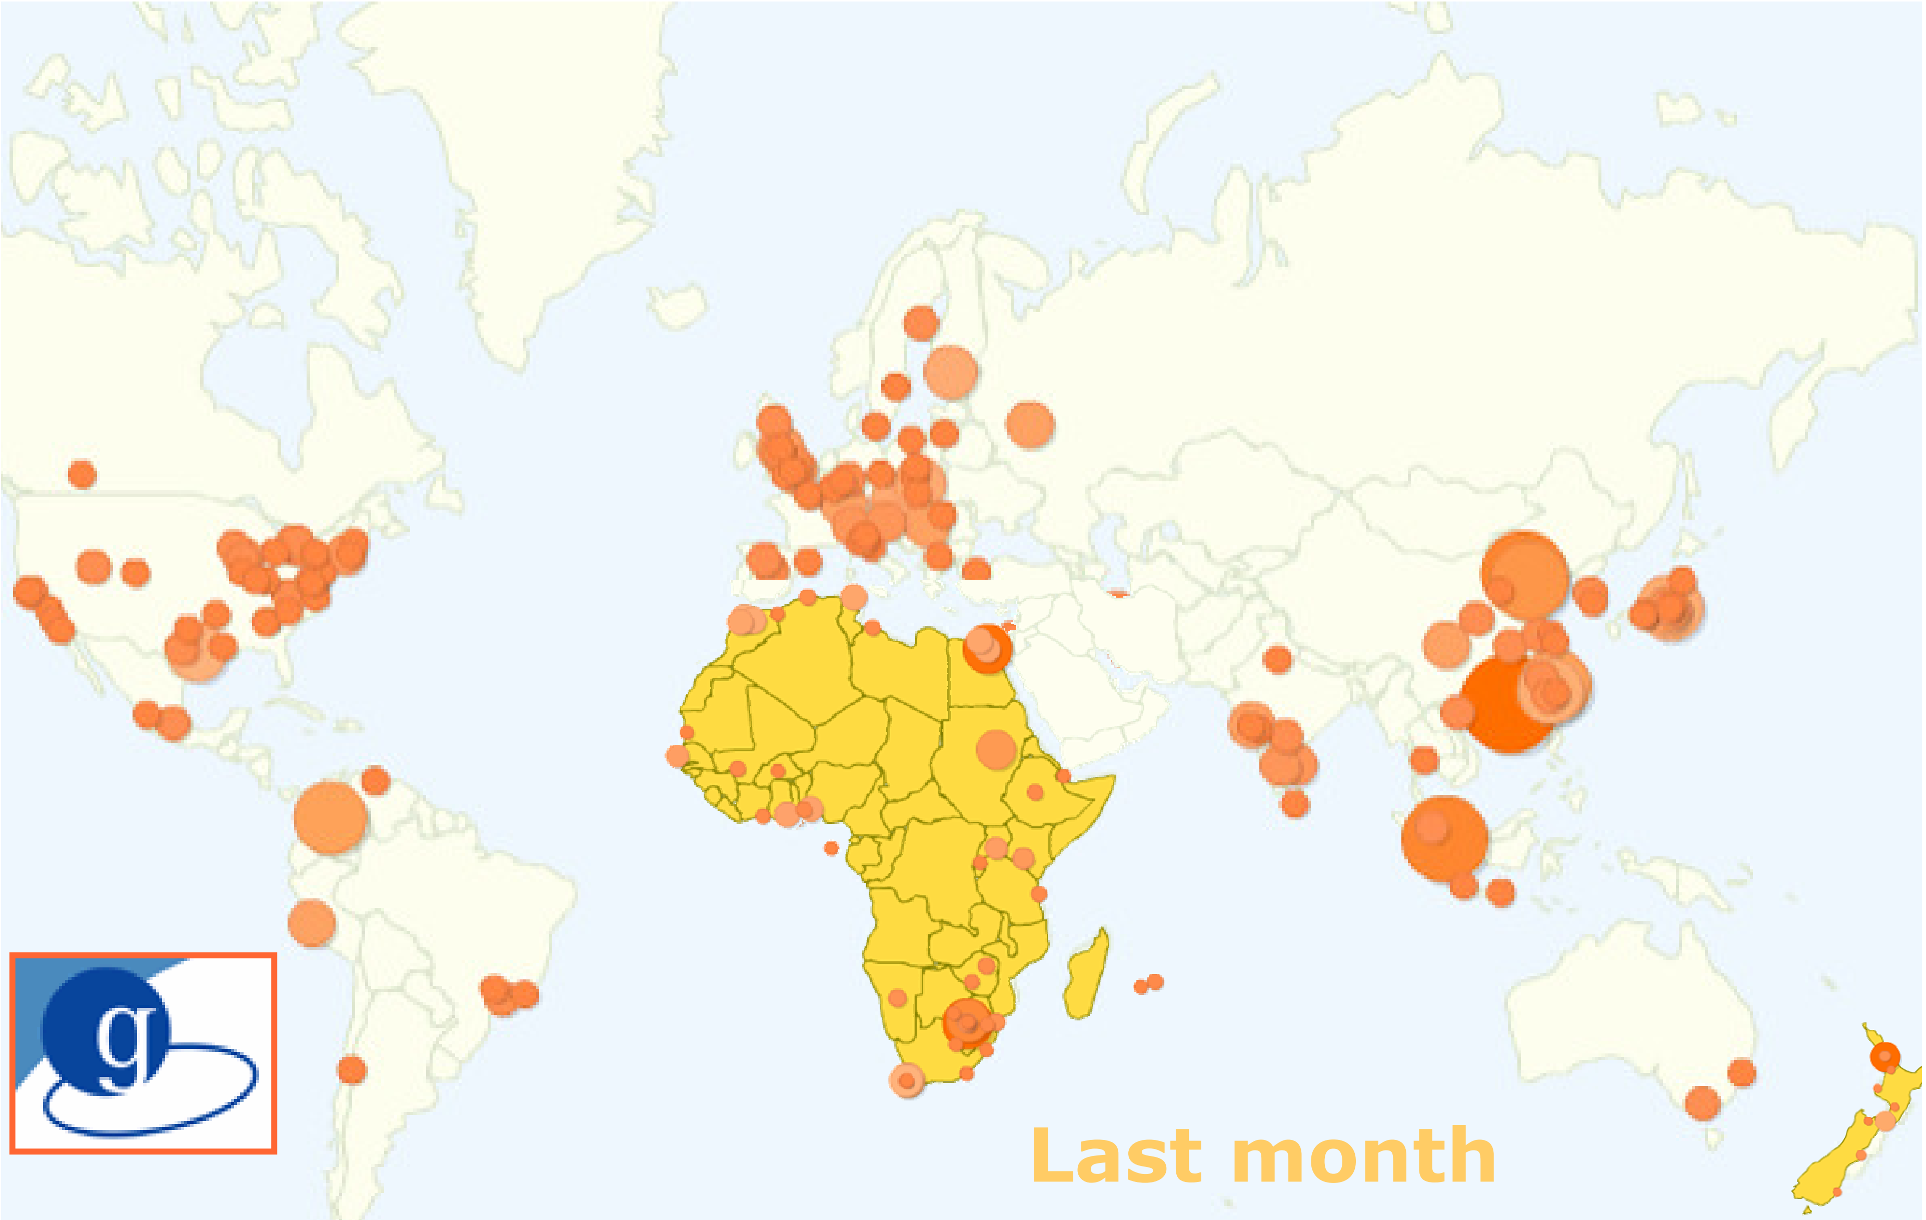

Phones: Wireless vs. Wired

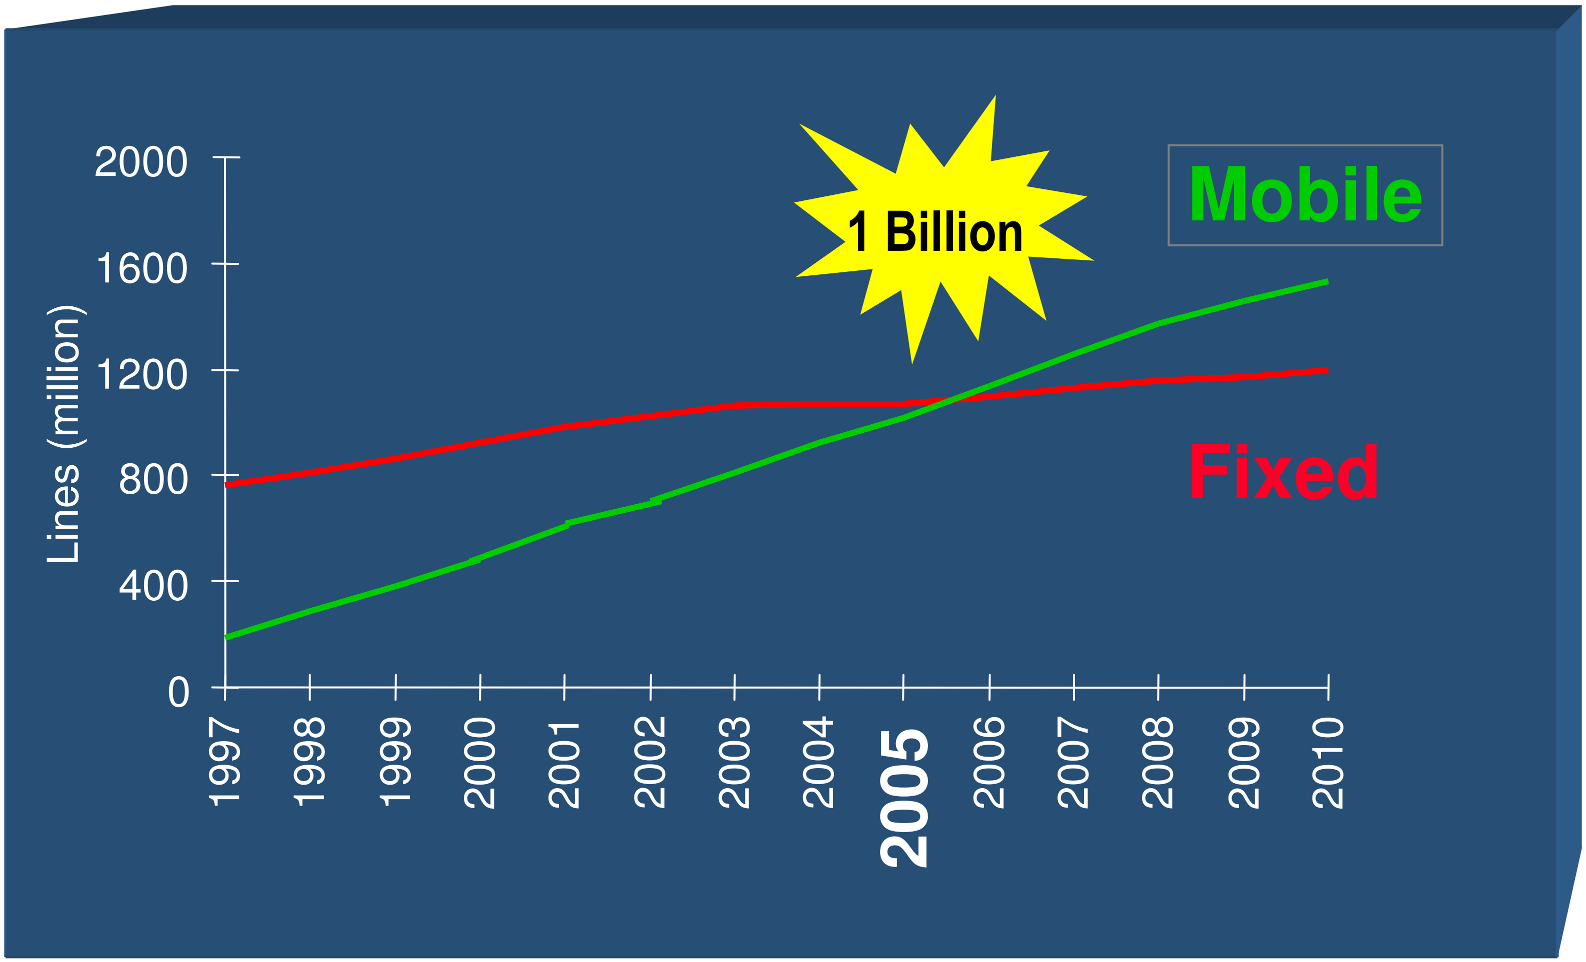

Source: Nortel

Earth to be Paradise; Distance to Lose Enchantment

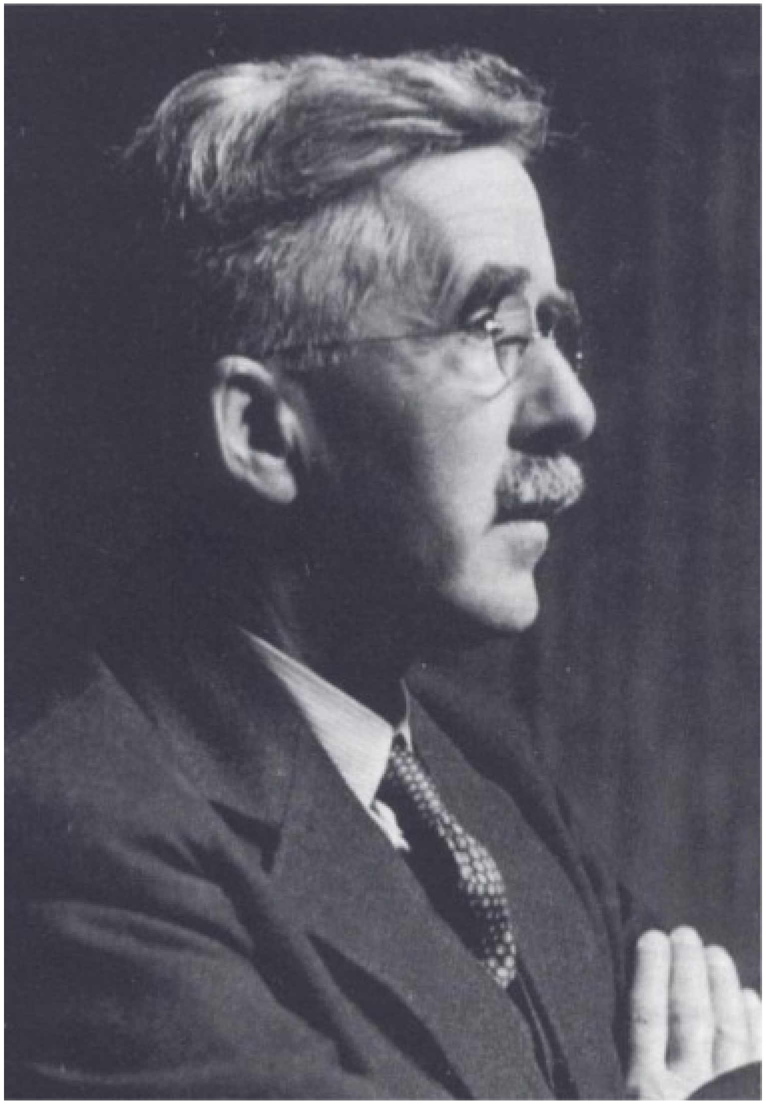

"If, as it is said to be not unlikely in the near future, the principle of sight is applied to the telephone as well as that of sound, earth will be in truth a paradise, and distance will lose its enchantment by being abolished altogether."

— Arthur Mee, 1898

Team Science

- Driven by complexity of problems to be addressed
- Enabled by quasi-ubiquitous Internet
- Exploiting diverse distributed expertise

Team Science meets Data Deluge

Information Big Bang

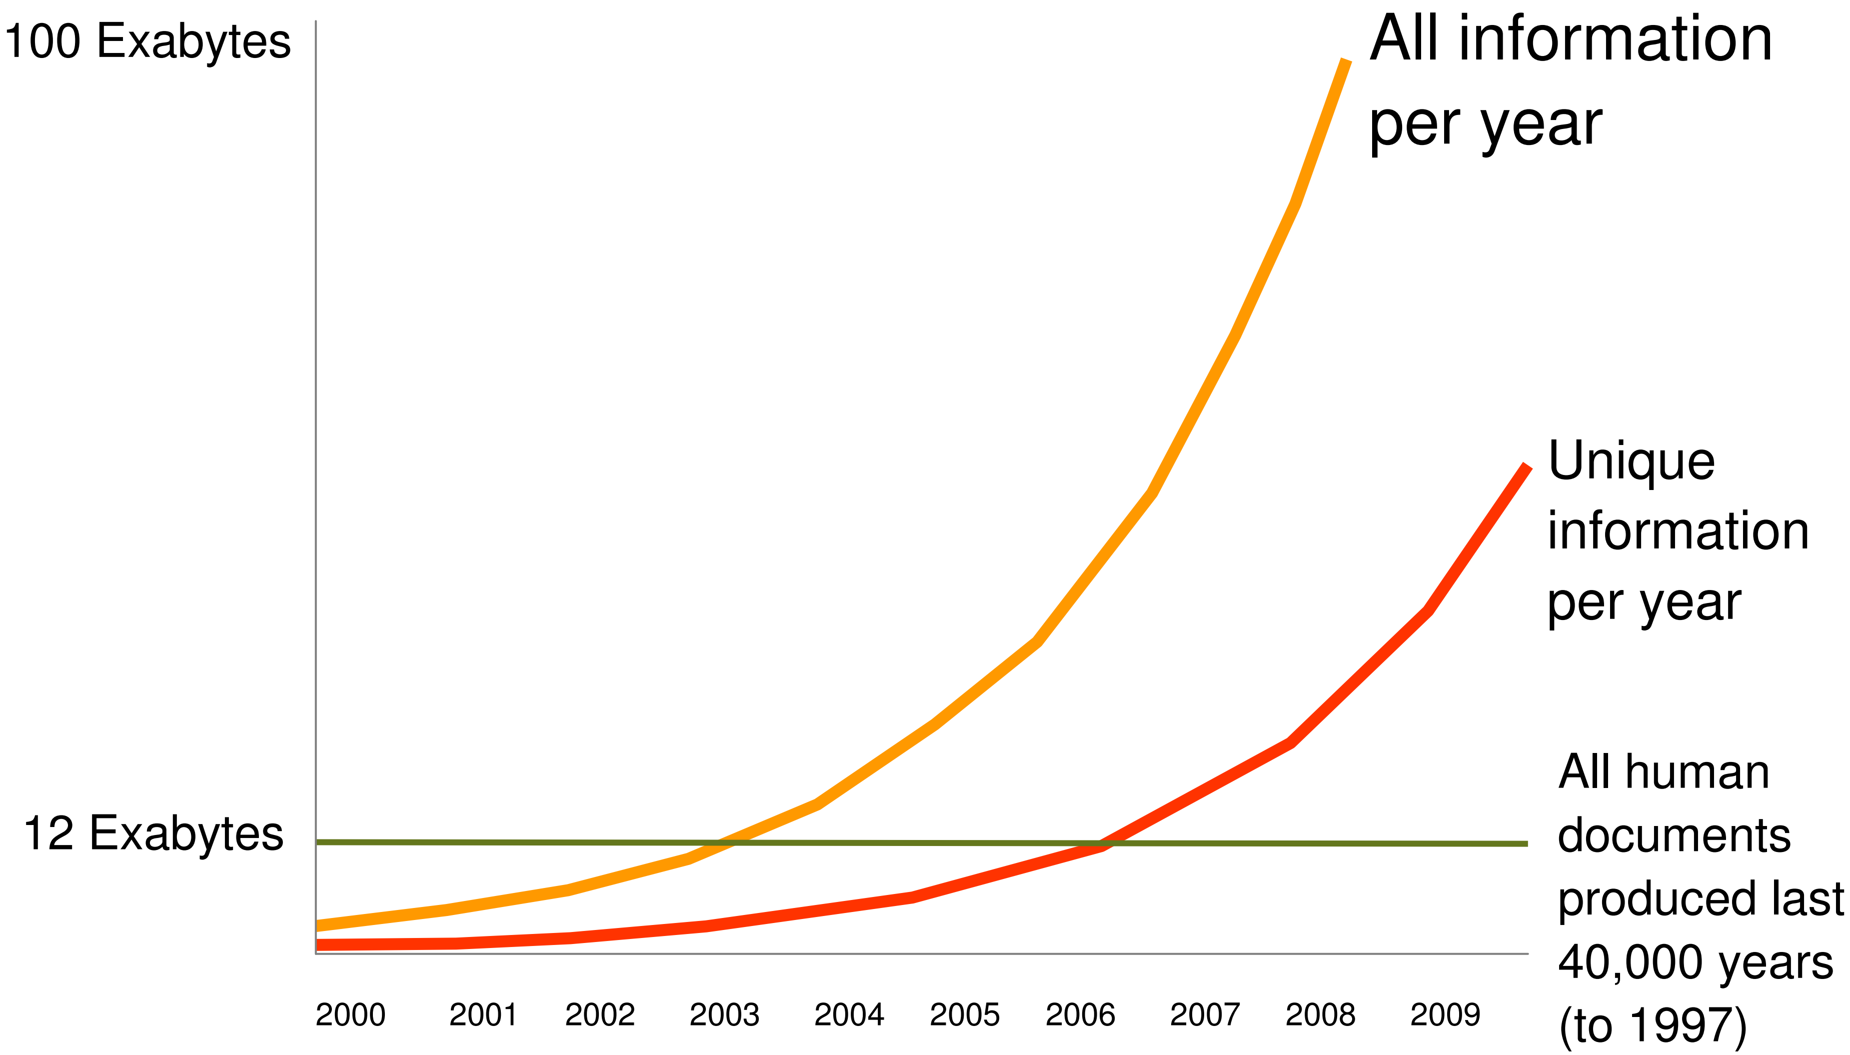

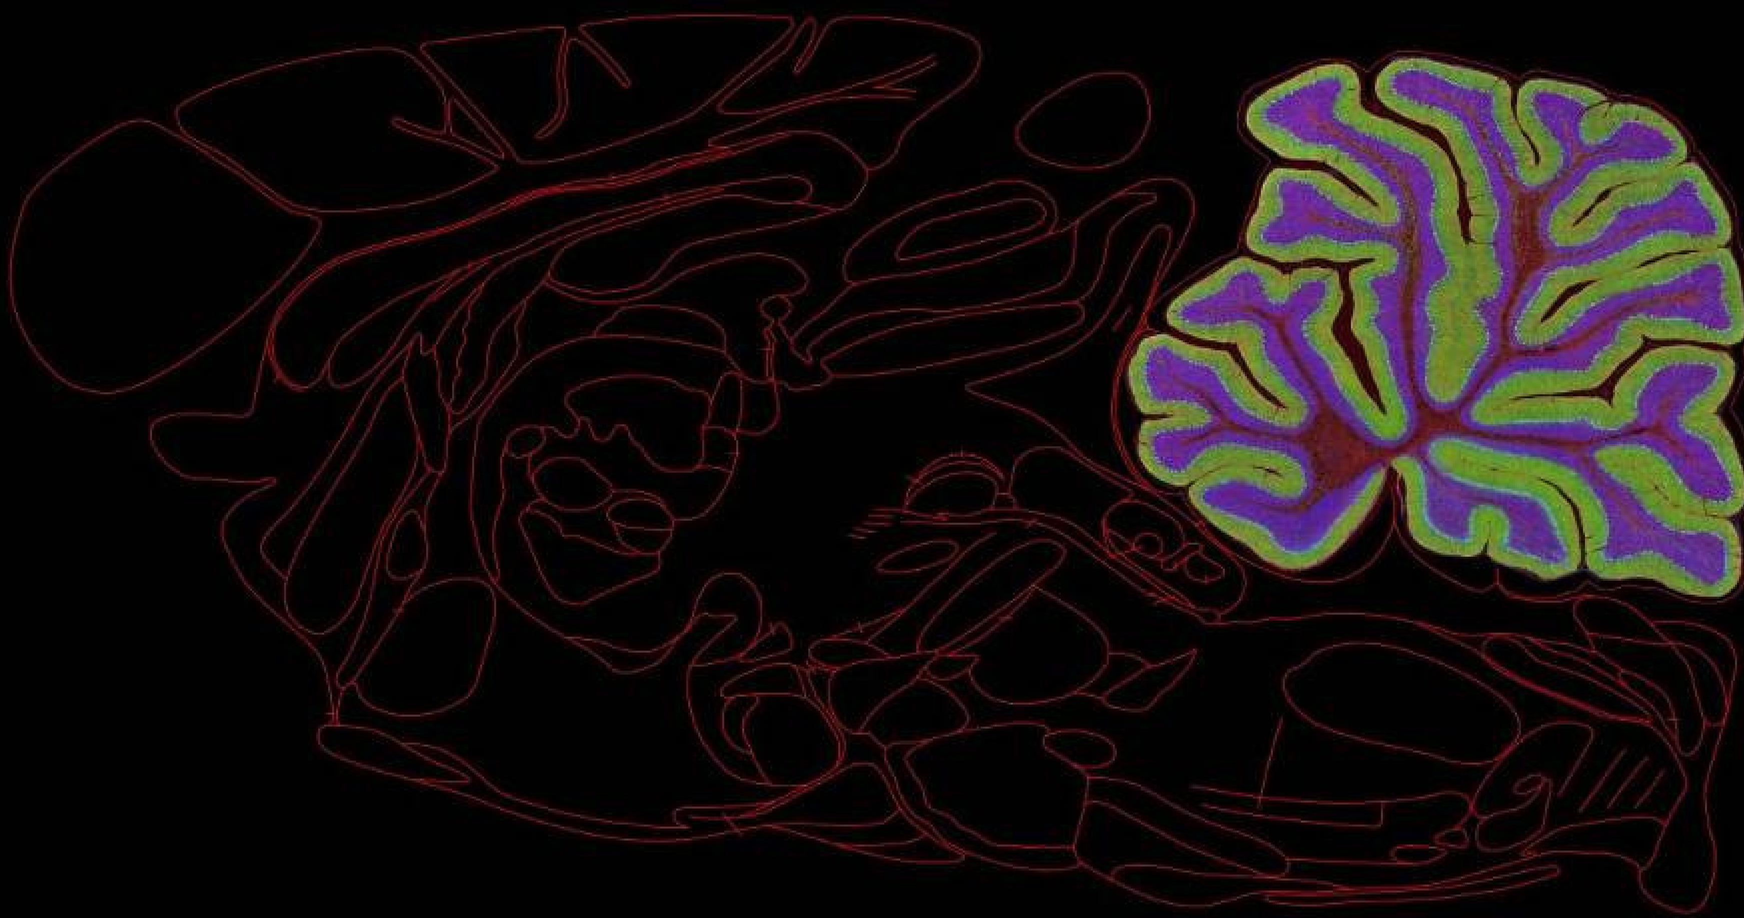

Images courtesy Mark Ellisman, UCSD

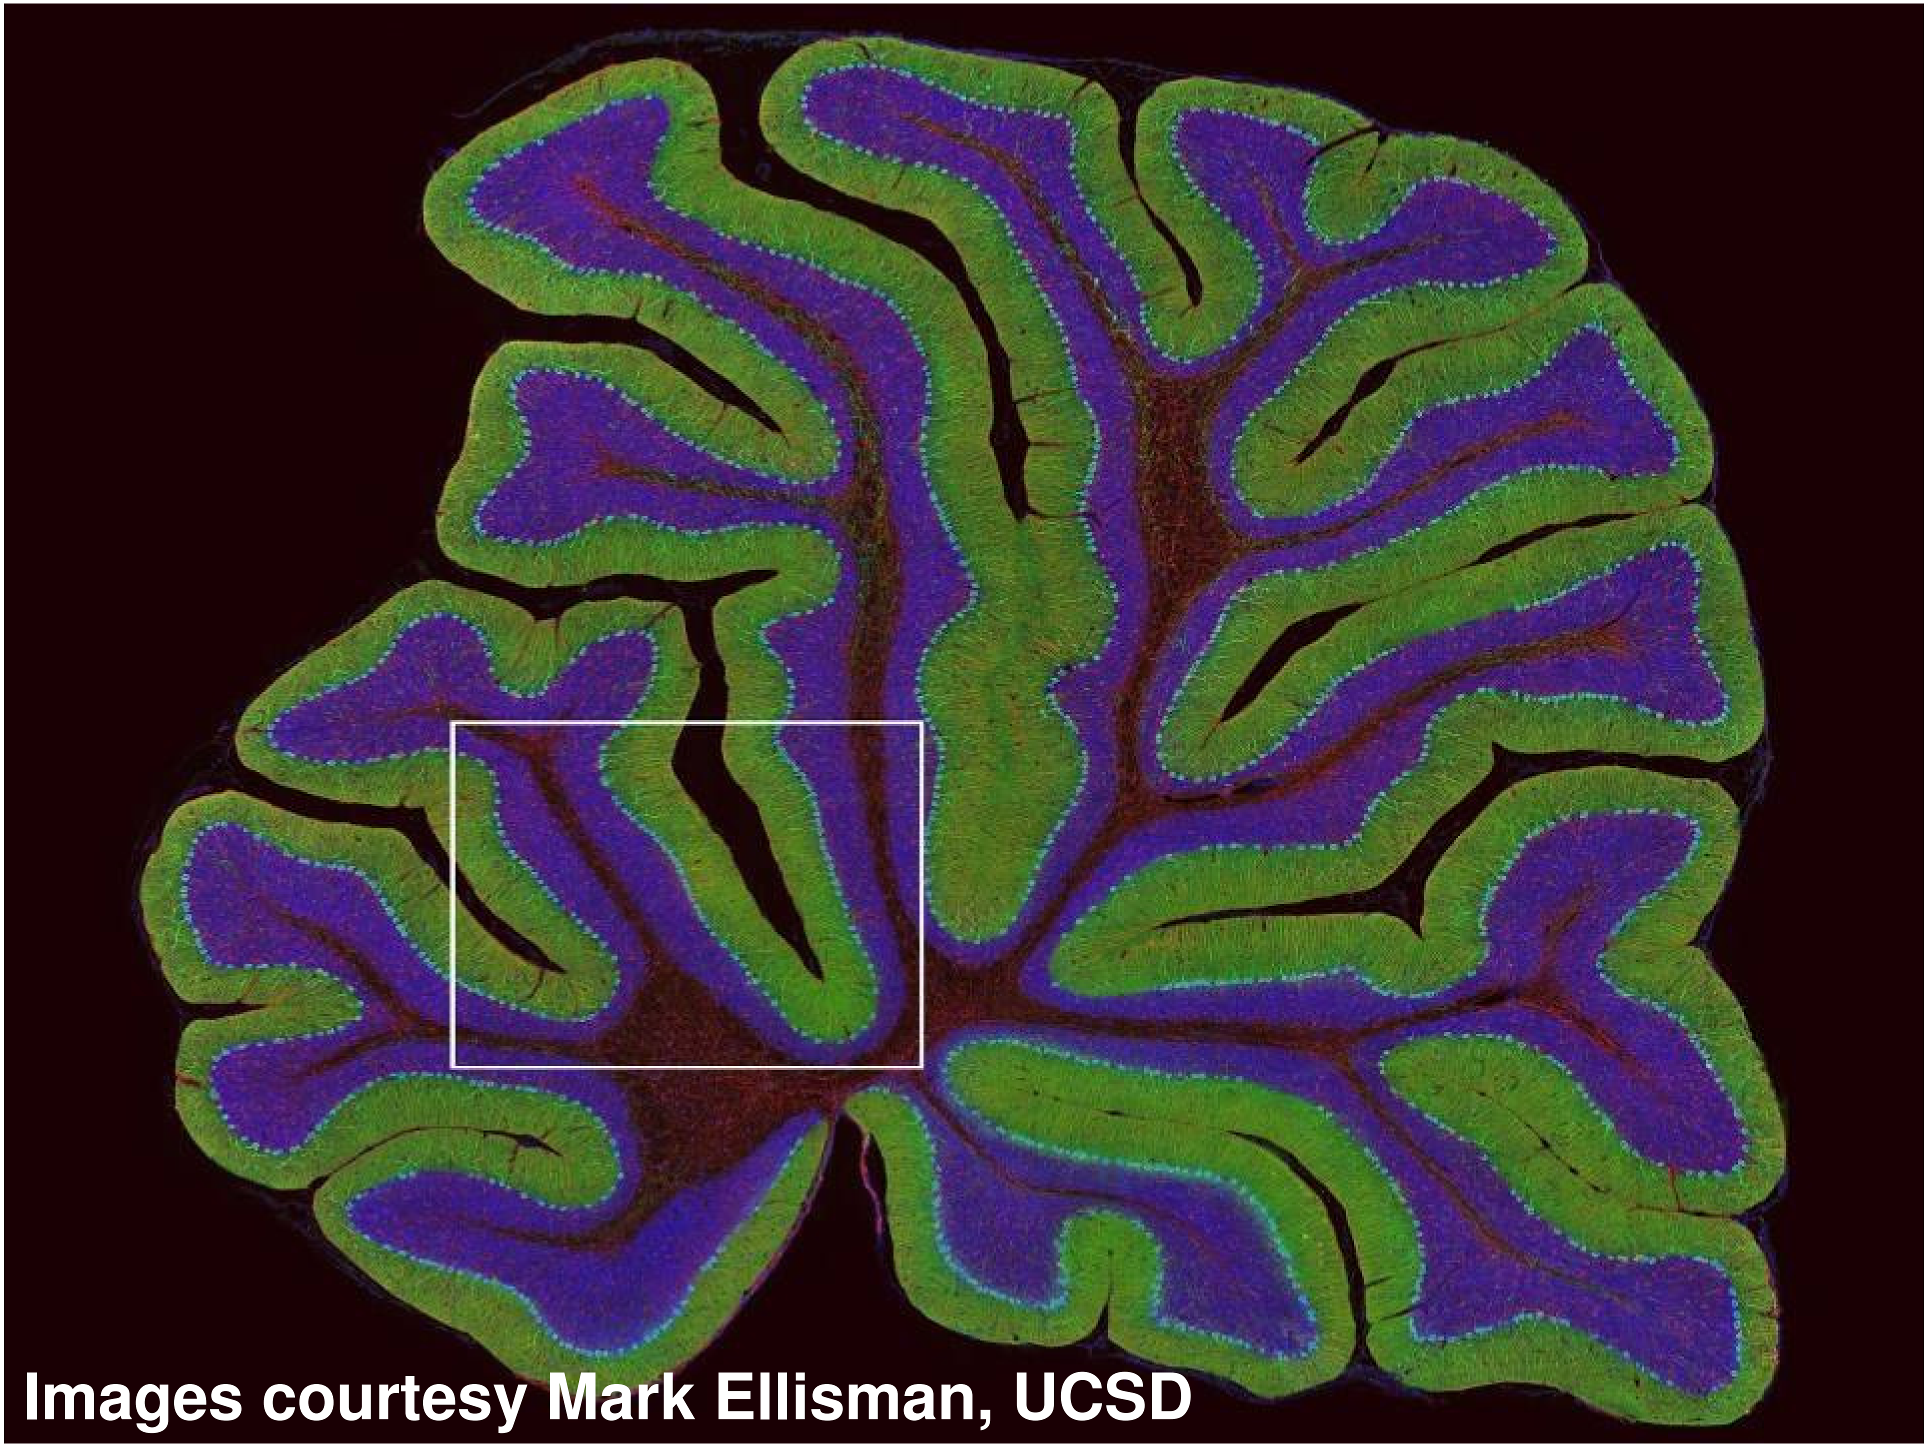
Images courtesy Mark Ellisman, UCSD

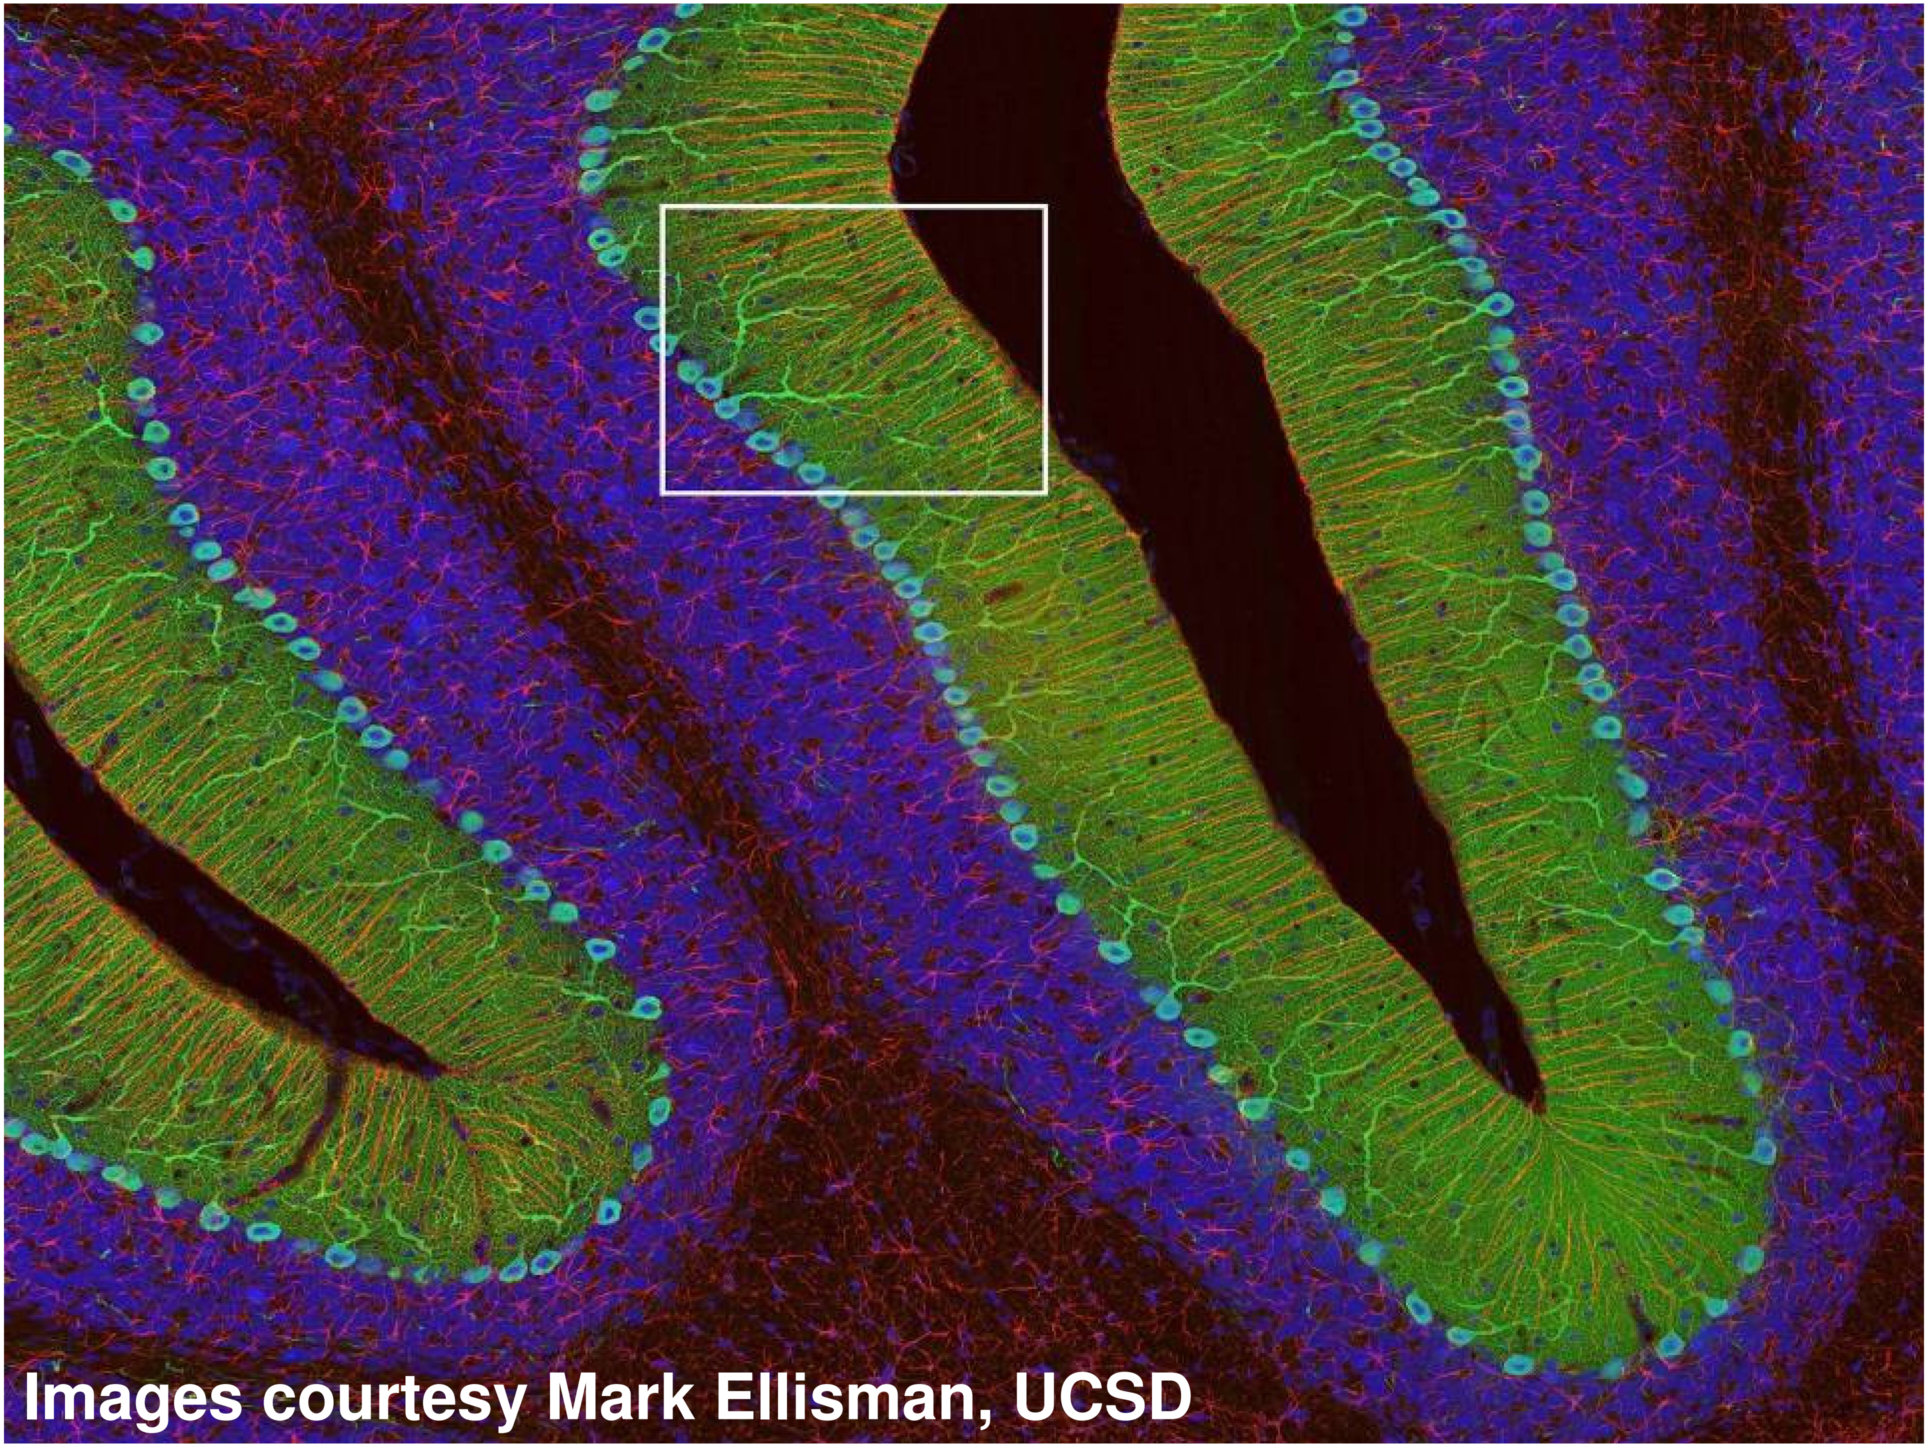
Images courtesy Mark Ellisman, UCSD


Images courtesy Mark Ellisman, UCSD

Post-World Cup celebration, Rome, July 9-10

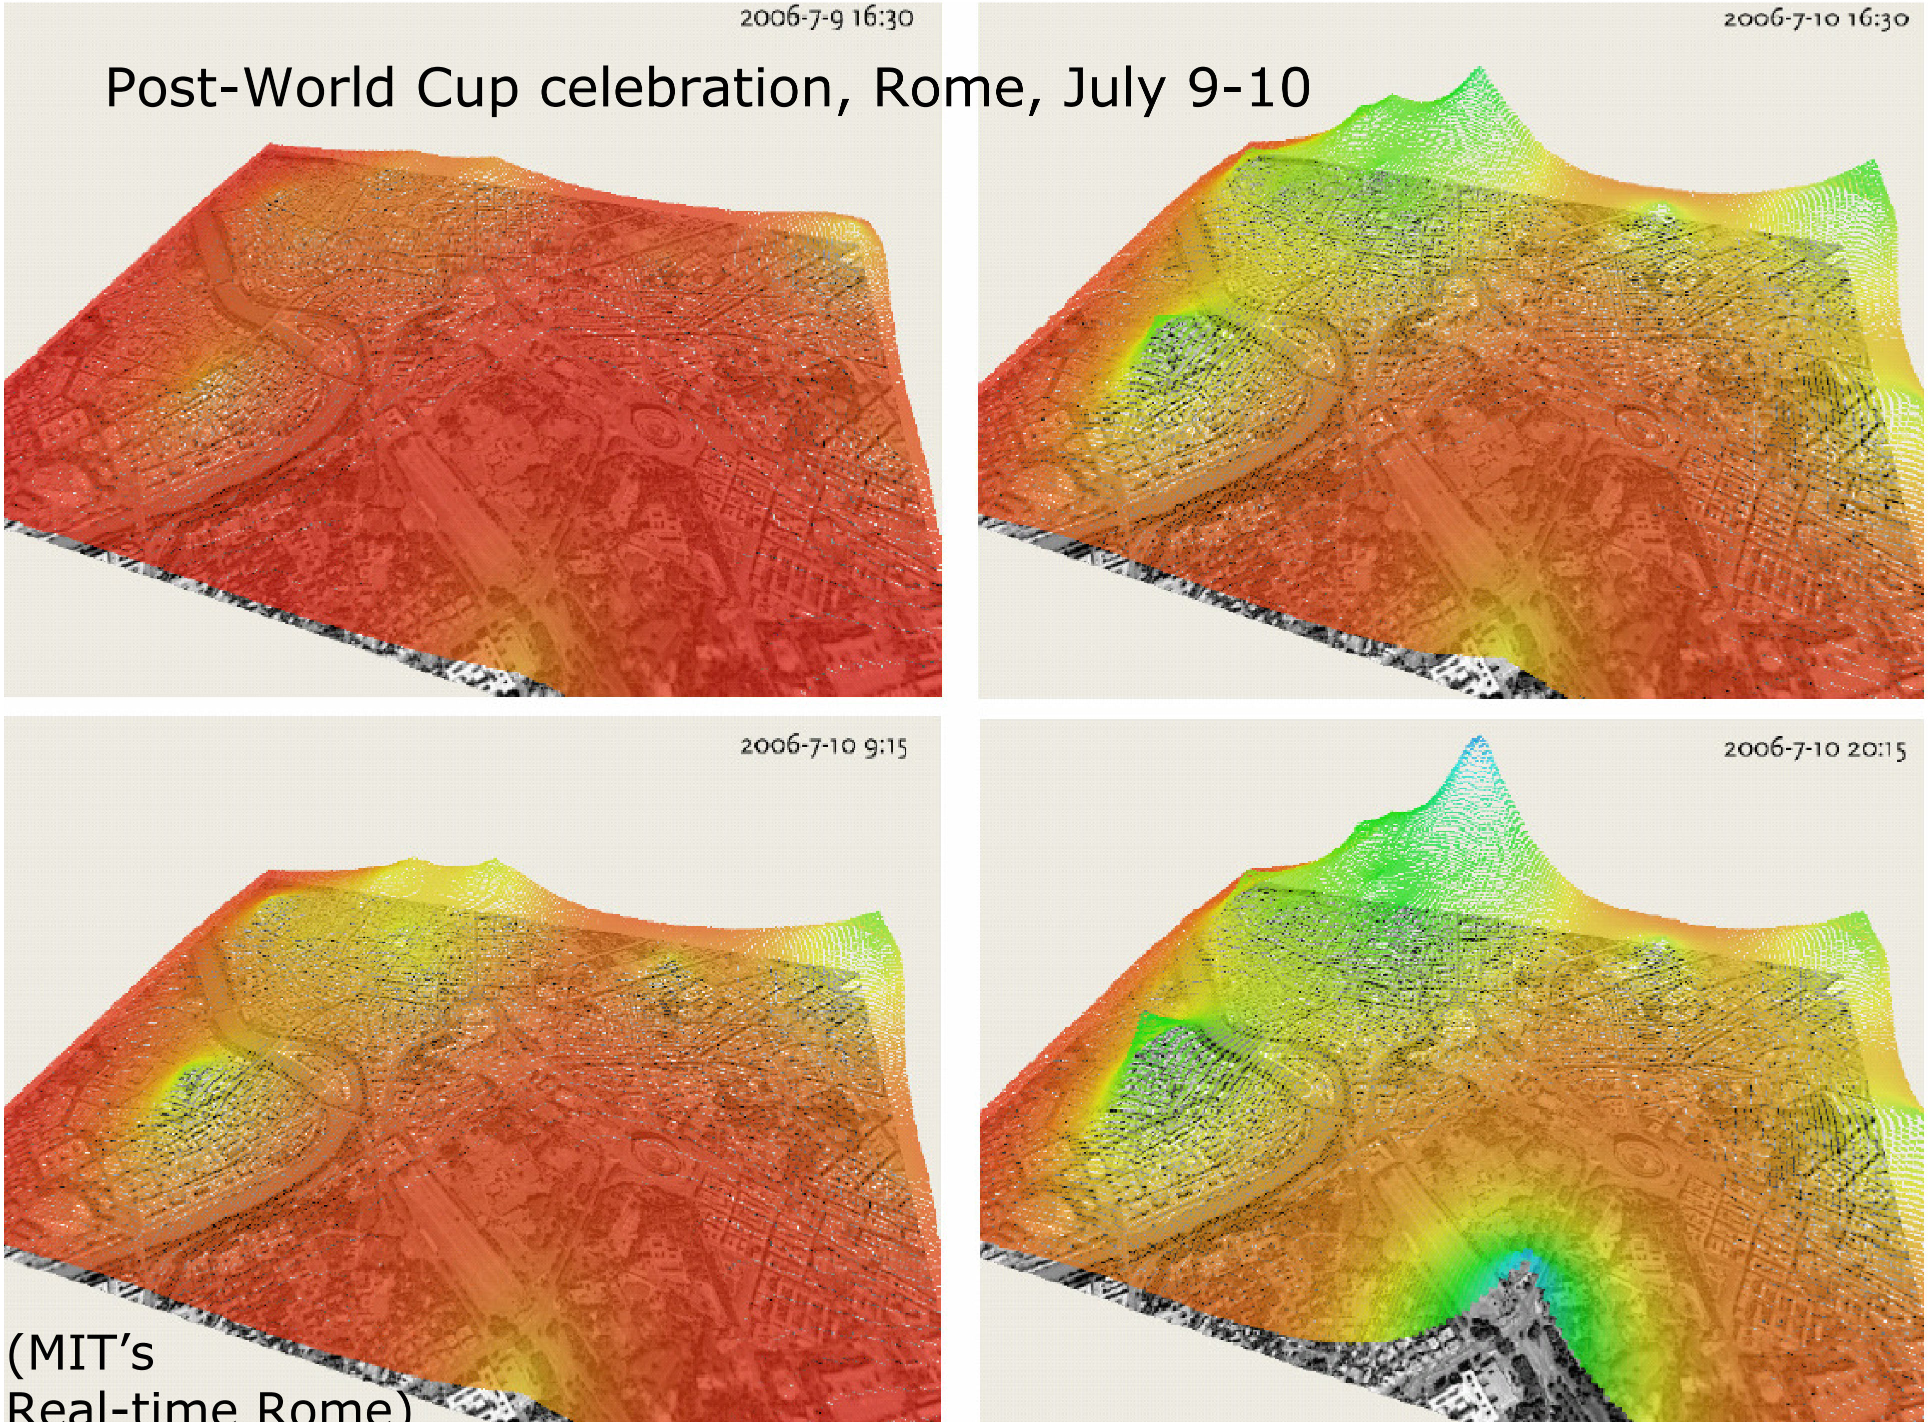

(MIT's
Real-time Rome)

Madonna Concert, cellphone activity

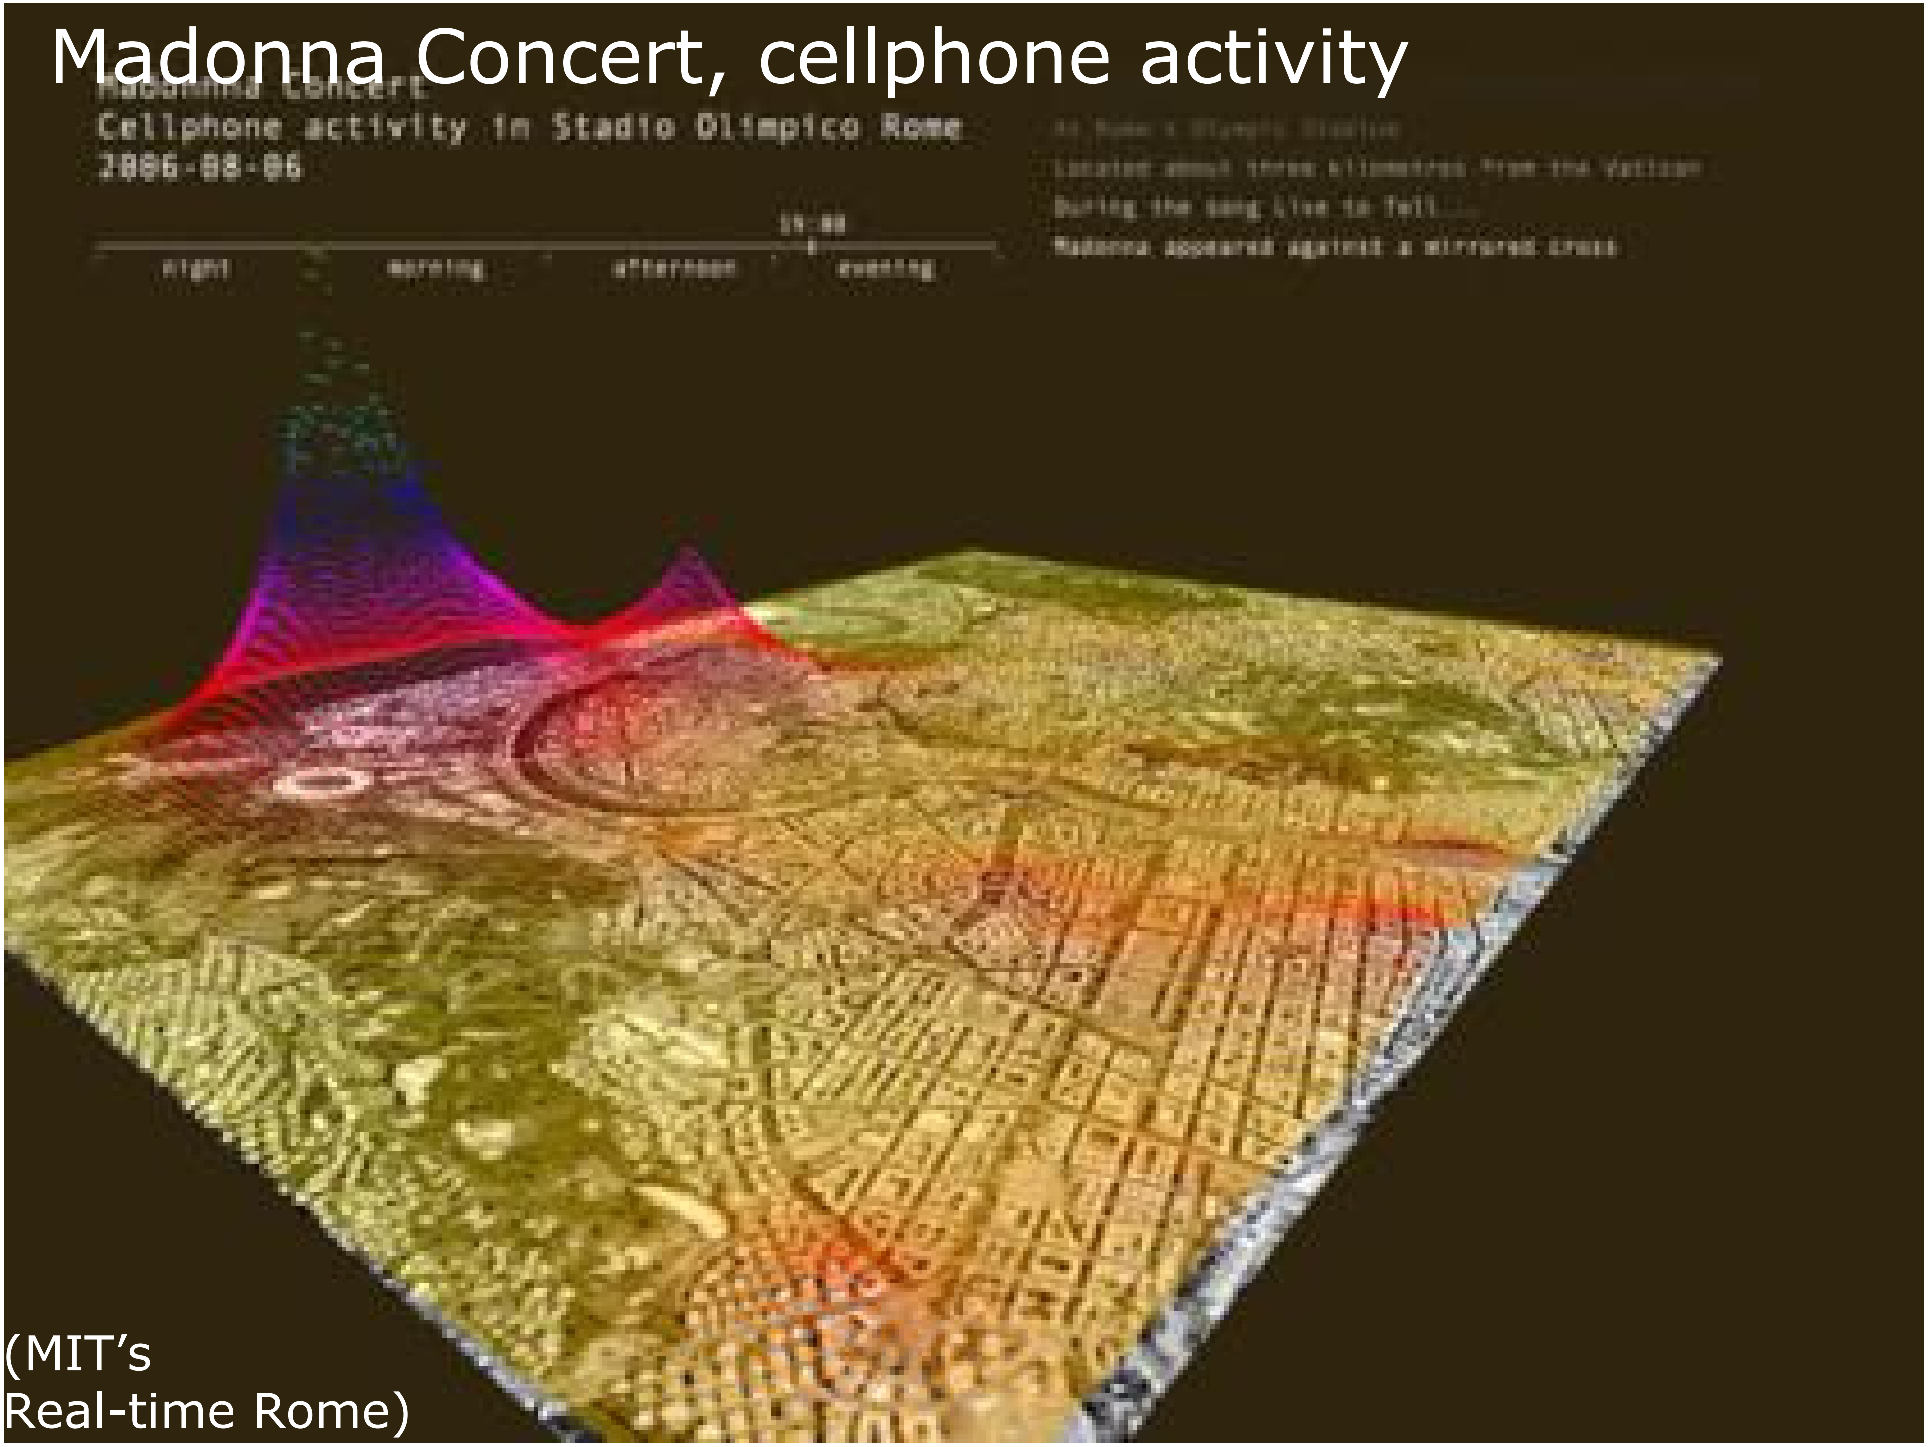

(MIT's
Real-time Rome)

Cell Phone Dynamics

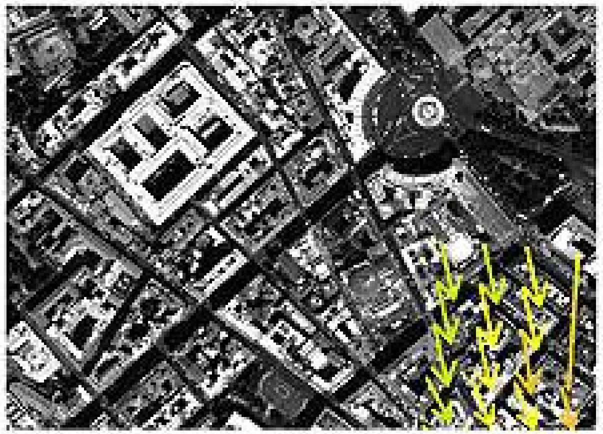
00 - 00 AM


04 - 00 AM


08 - 00 AM


12 - 00 PM


04 - 00 PM


20 - 00 PM

(Real-time Rome, MIT)

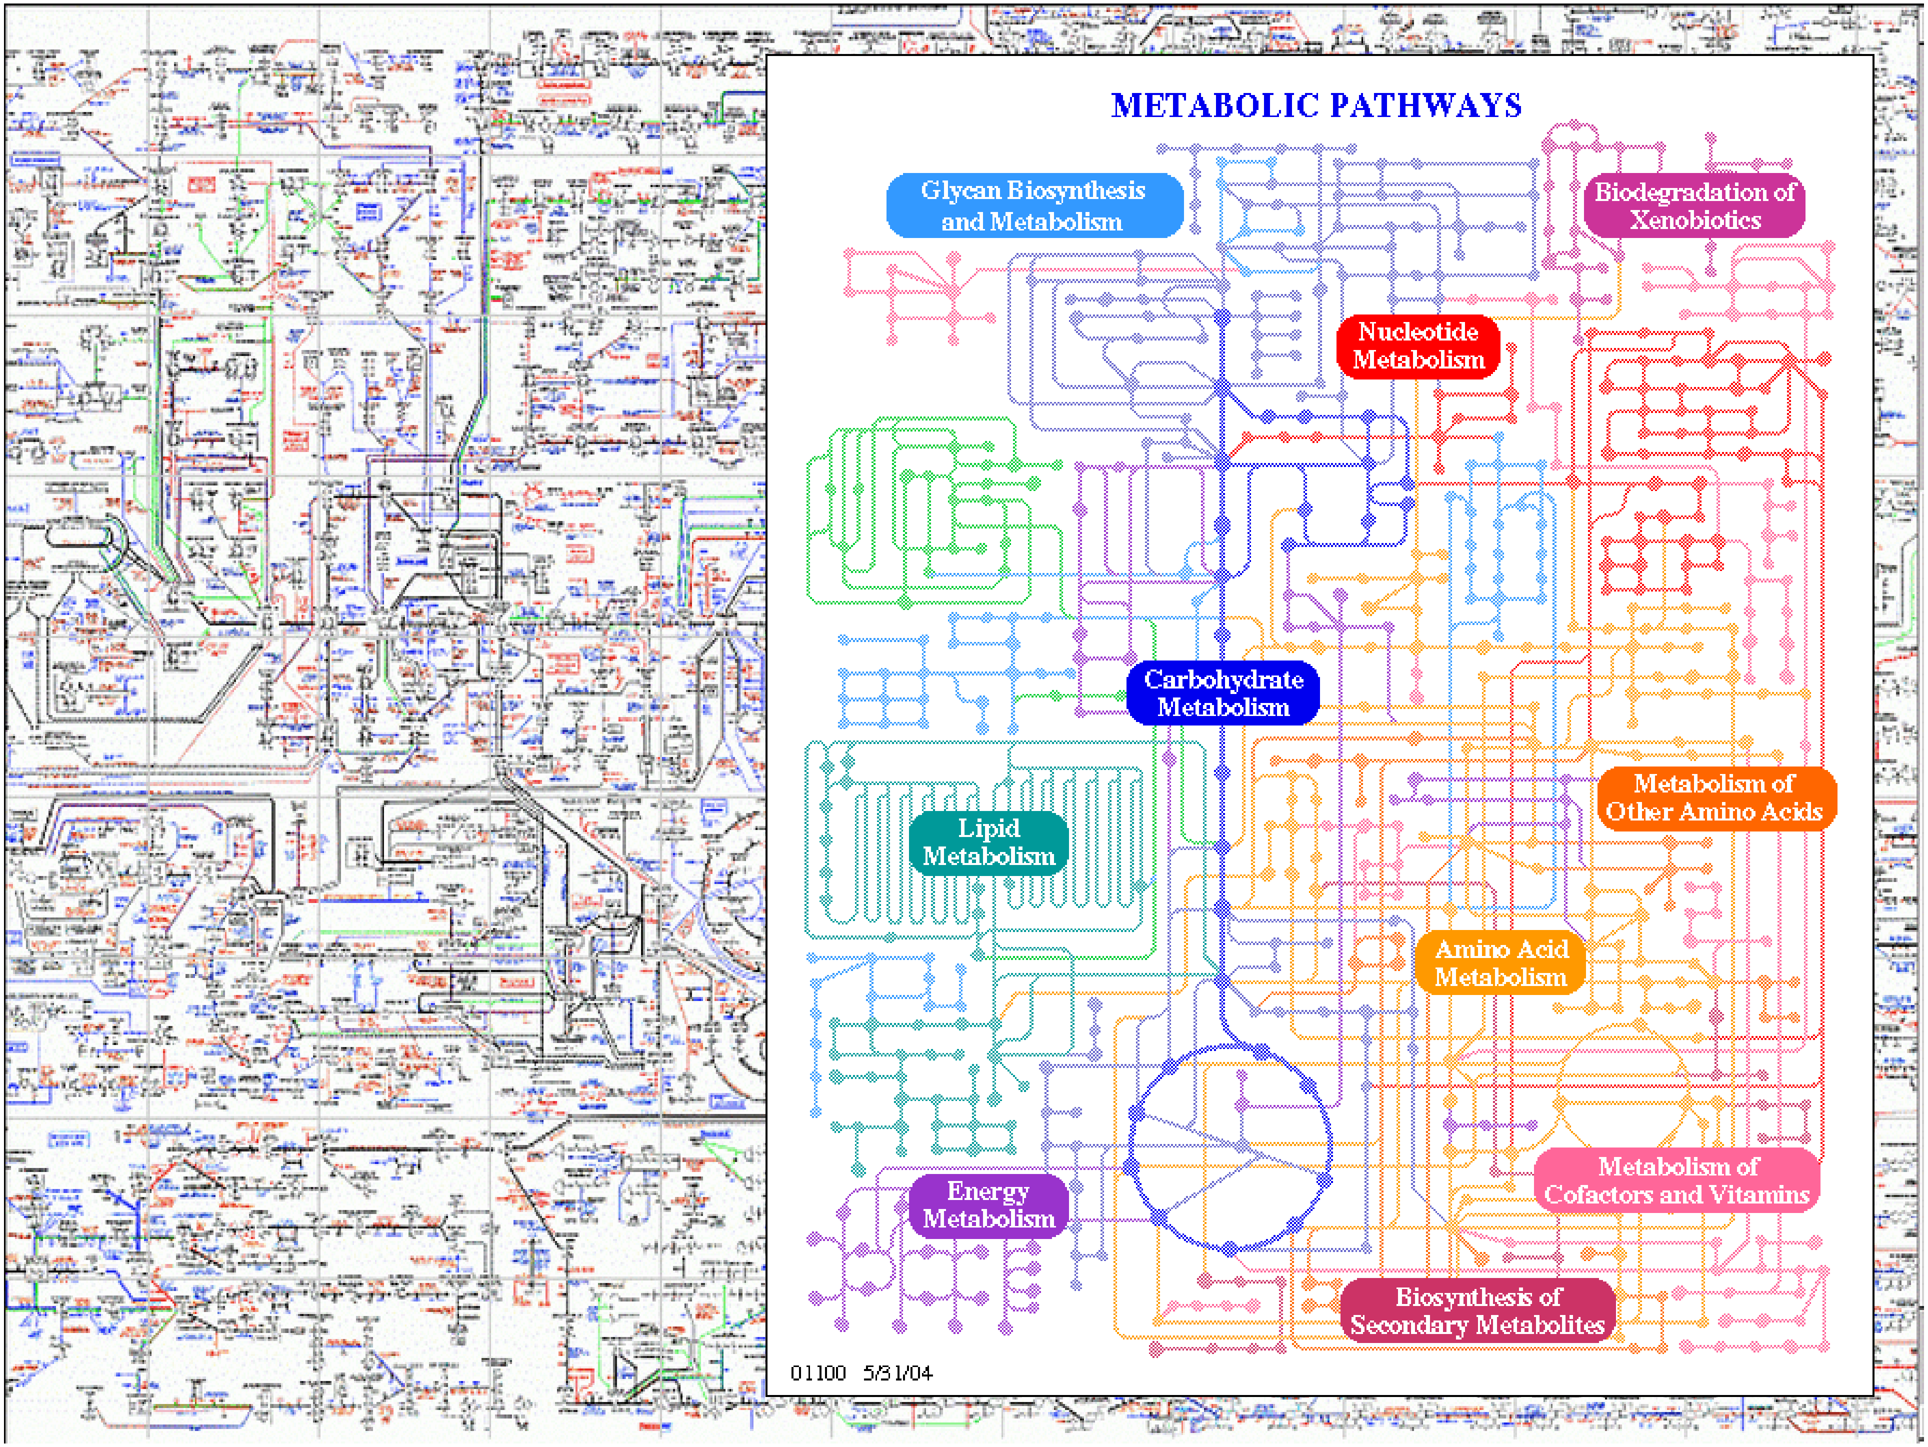
METABOLIC PATHWAYS
Glycan Biosynthesis and Metabolism
Biodegradation of Xenobiotics
Nucleotide Metabolism
Carbohydrate Metabolism
Lipid Metabolism
Metabolism of Other Amino Acids
Amino Acid Metabolism
Energy Metabolism
Metabolism of Cofactors and Vitamins
Biosynthesis of Secondary Metabolites
01100 5/31/04

New Data Repositories

Global Observation Database (View)

VCR-Style Control Panel

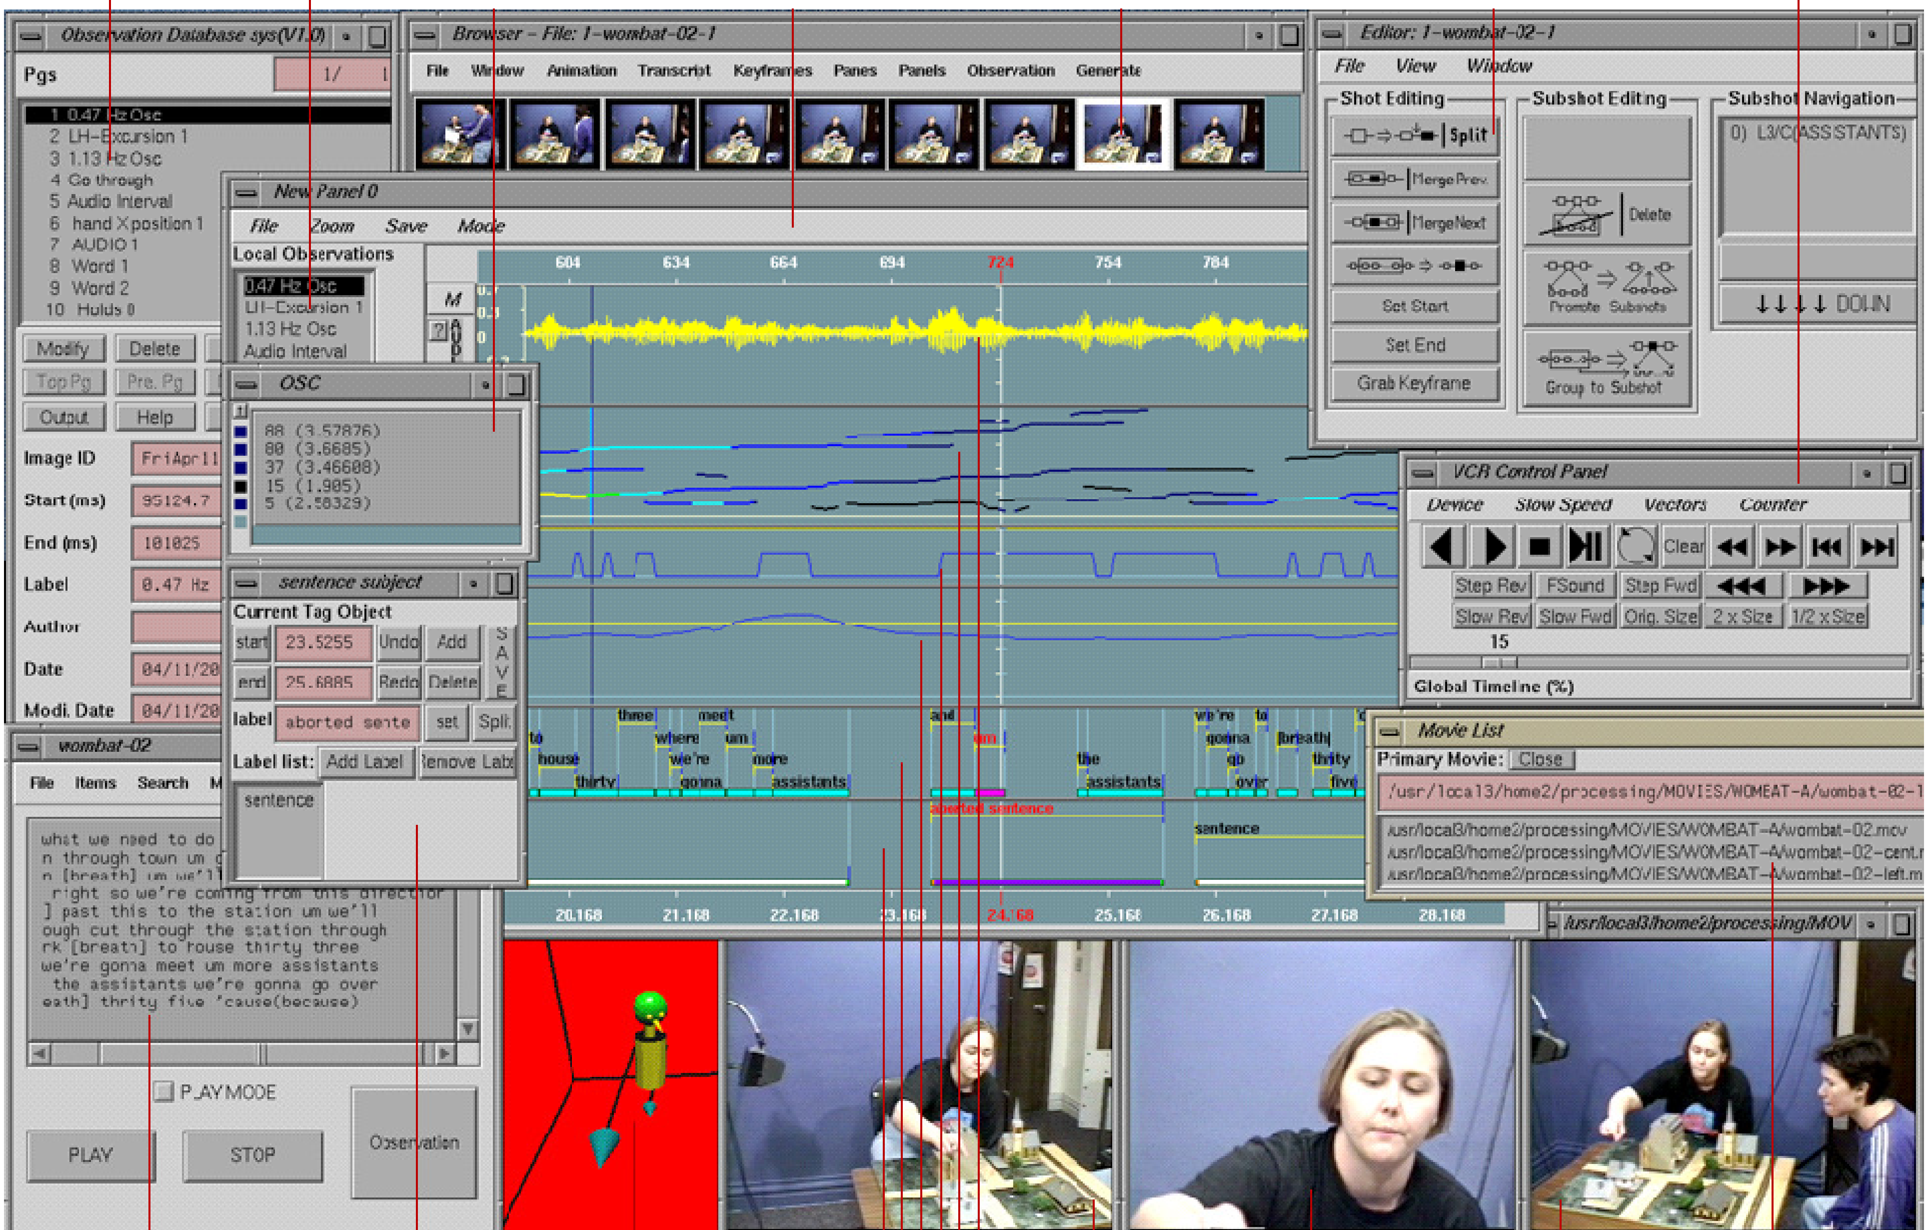

Animated Text Transcript (Paragraph Representation)

Tag Transcript Editor

Animated Avatar Representation

Animated Graph Panes

Video Displays

Video List

Bennett Berthenthal et al., www.sidgrid.org

Arabidopsis articles

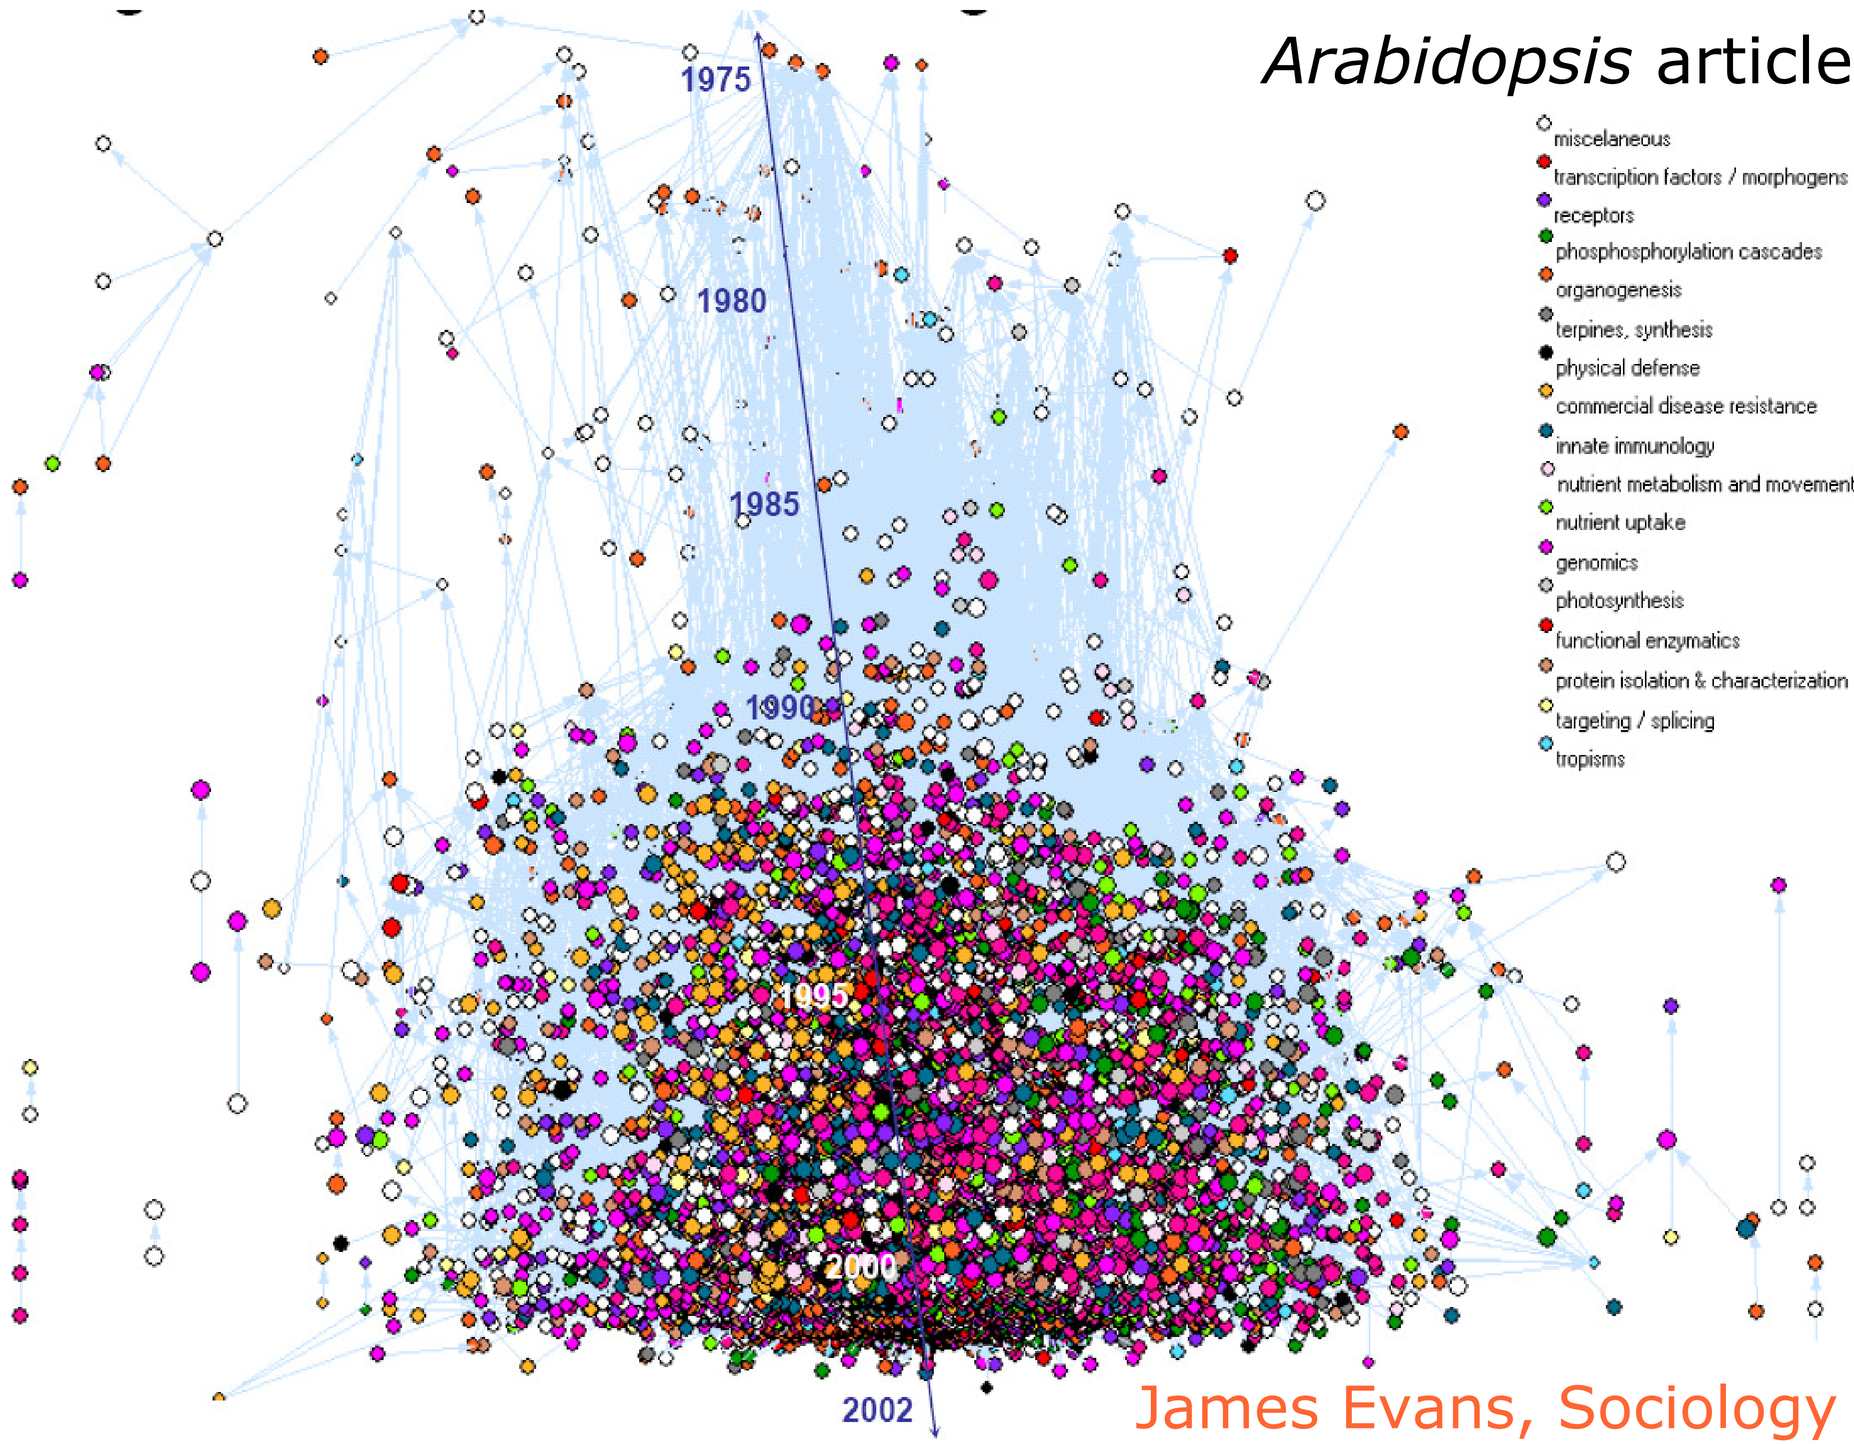

James Evans, Sociology

Virtual Environments as Social Science Laboratories

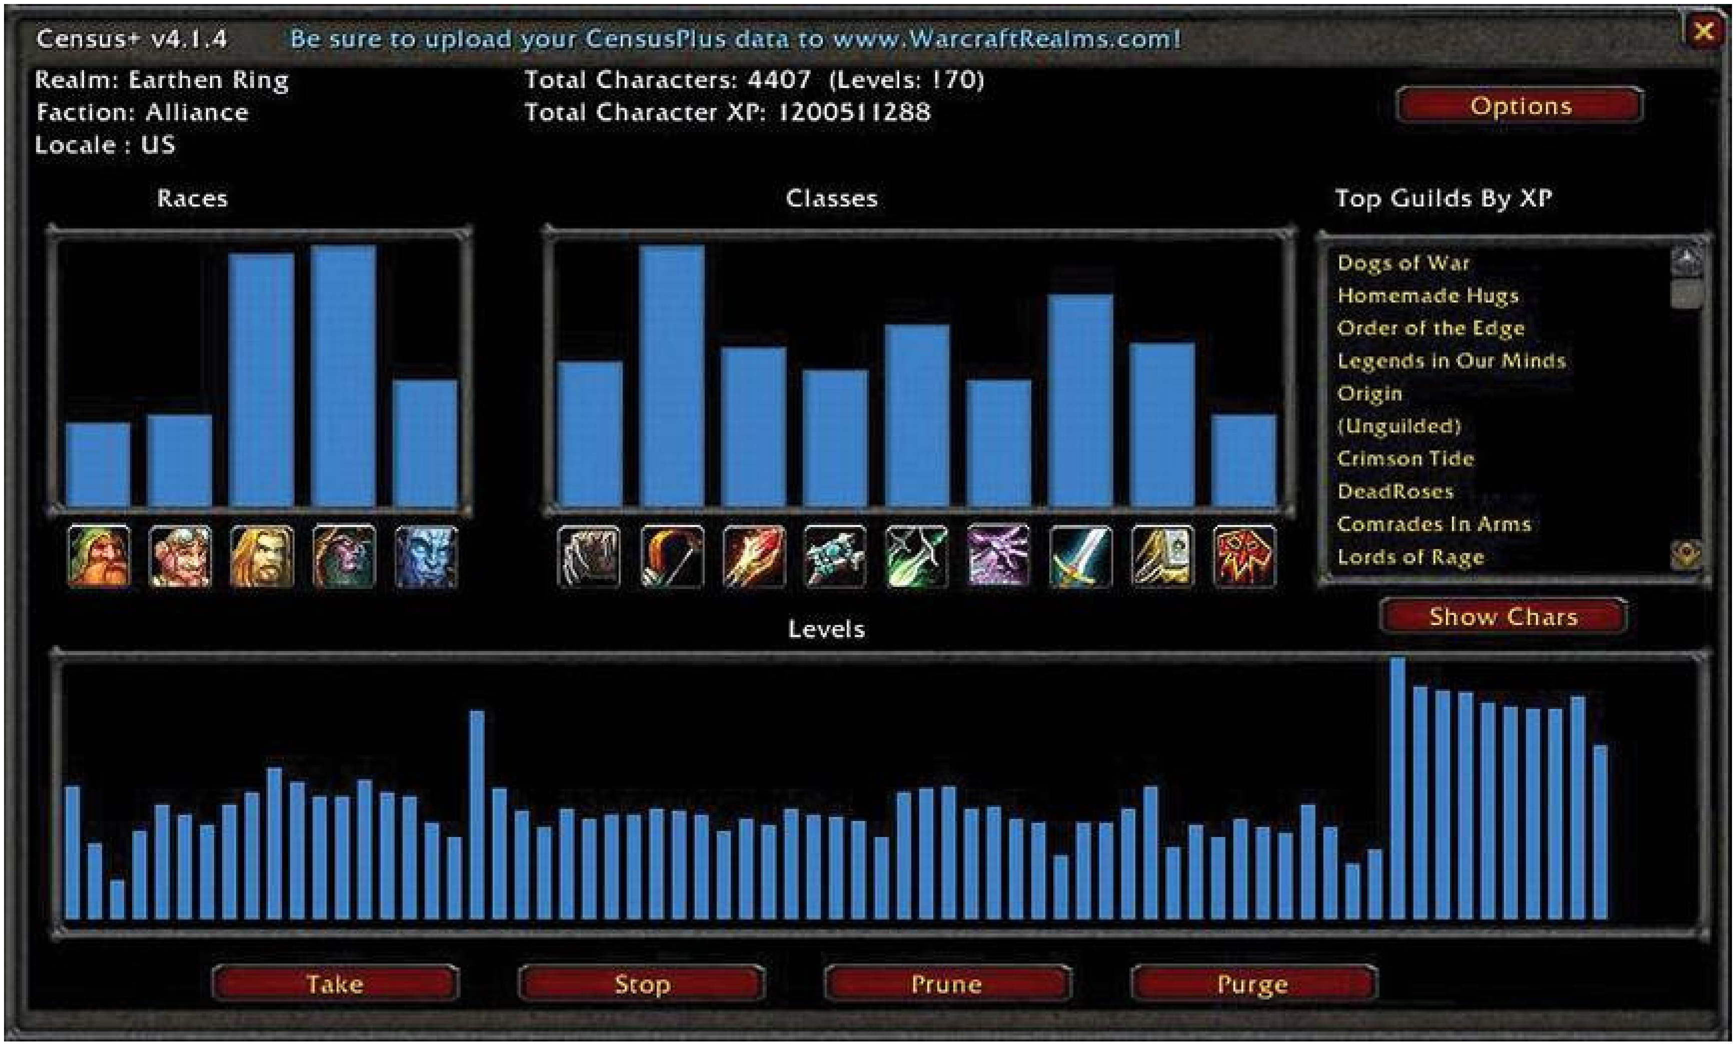

Fig. 3. Example of CensusPlus output from WoW
W. S. Bainbridge Science 317, 472-476 (2007)

First Generation Grids: **Batch Computing**

Focus on federation of many computers for massively (data-)parallel applications

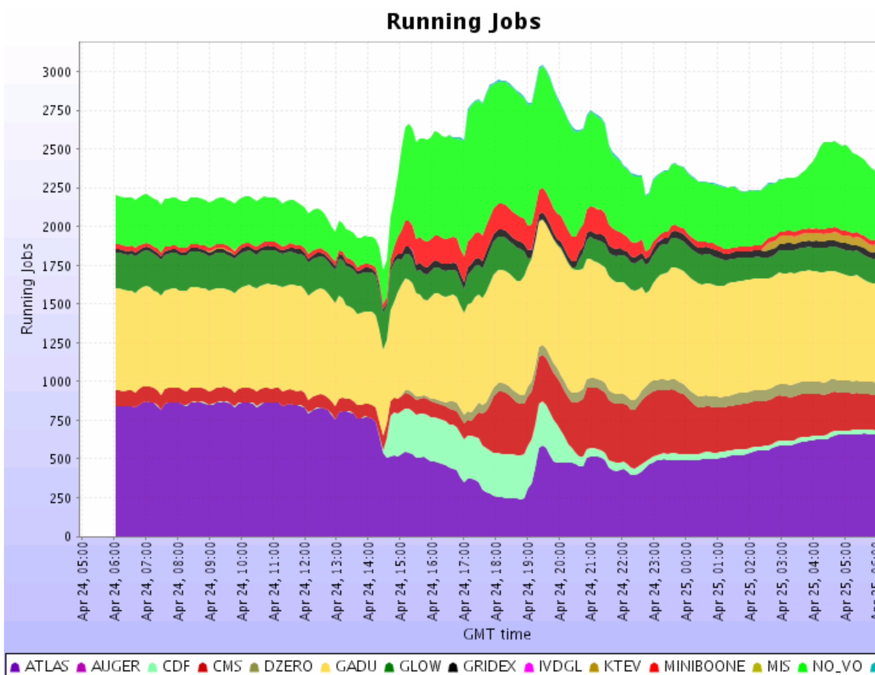

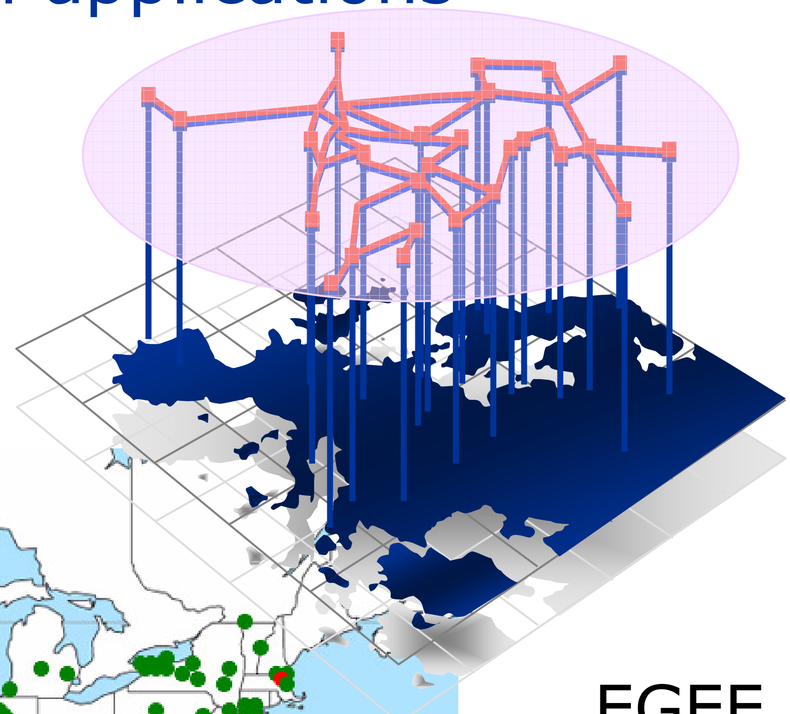

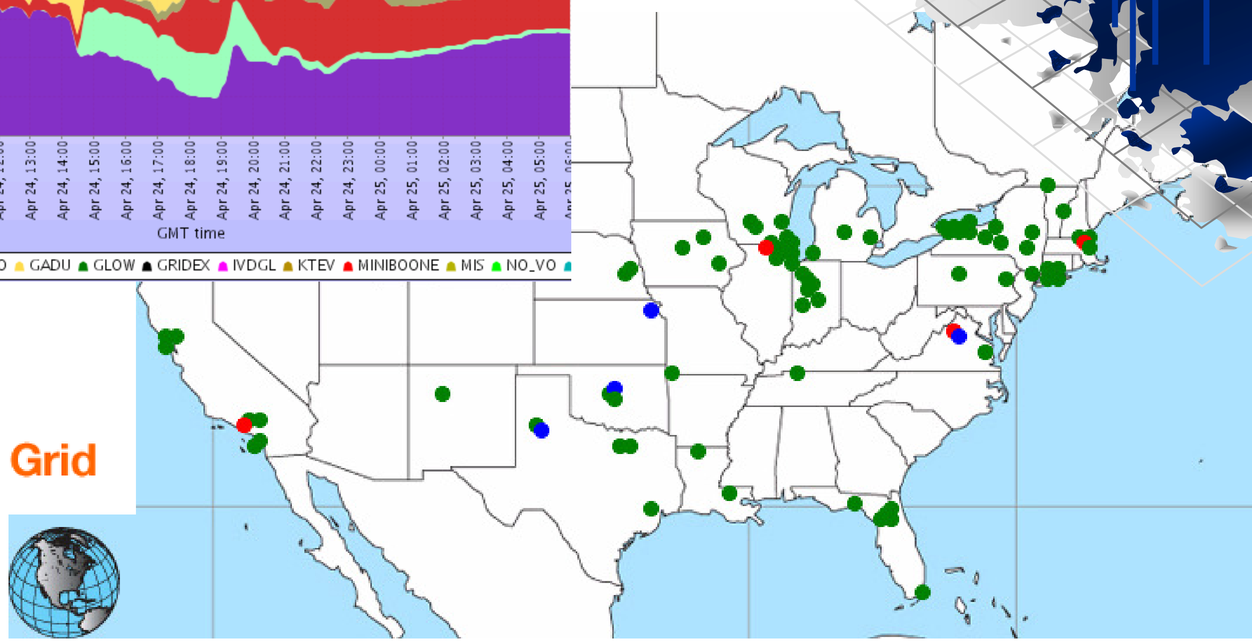

EGEE

Open Science Grid

Globus

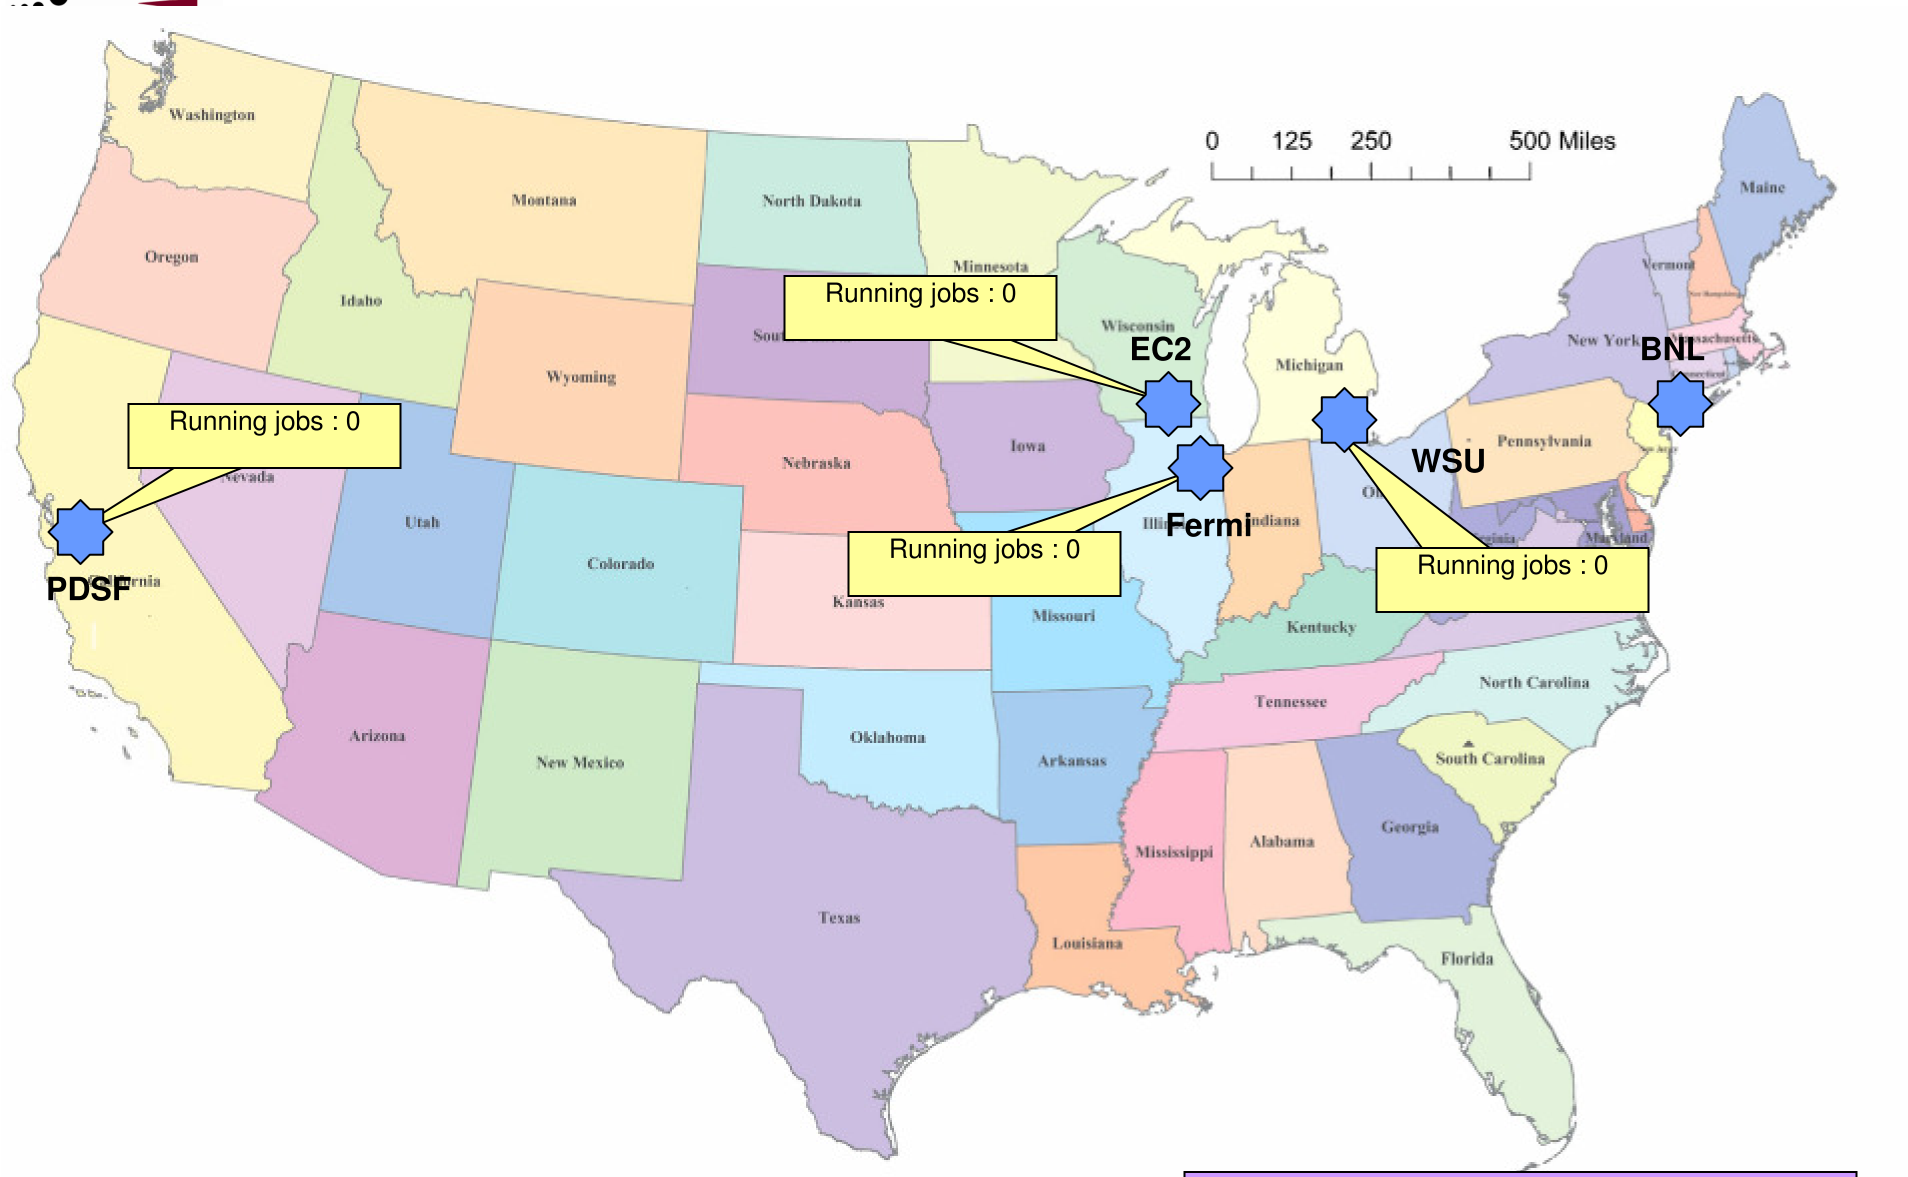

Job Completion :

File Recovery :

An Example Application: Functional MRI Data Analysis

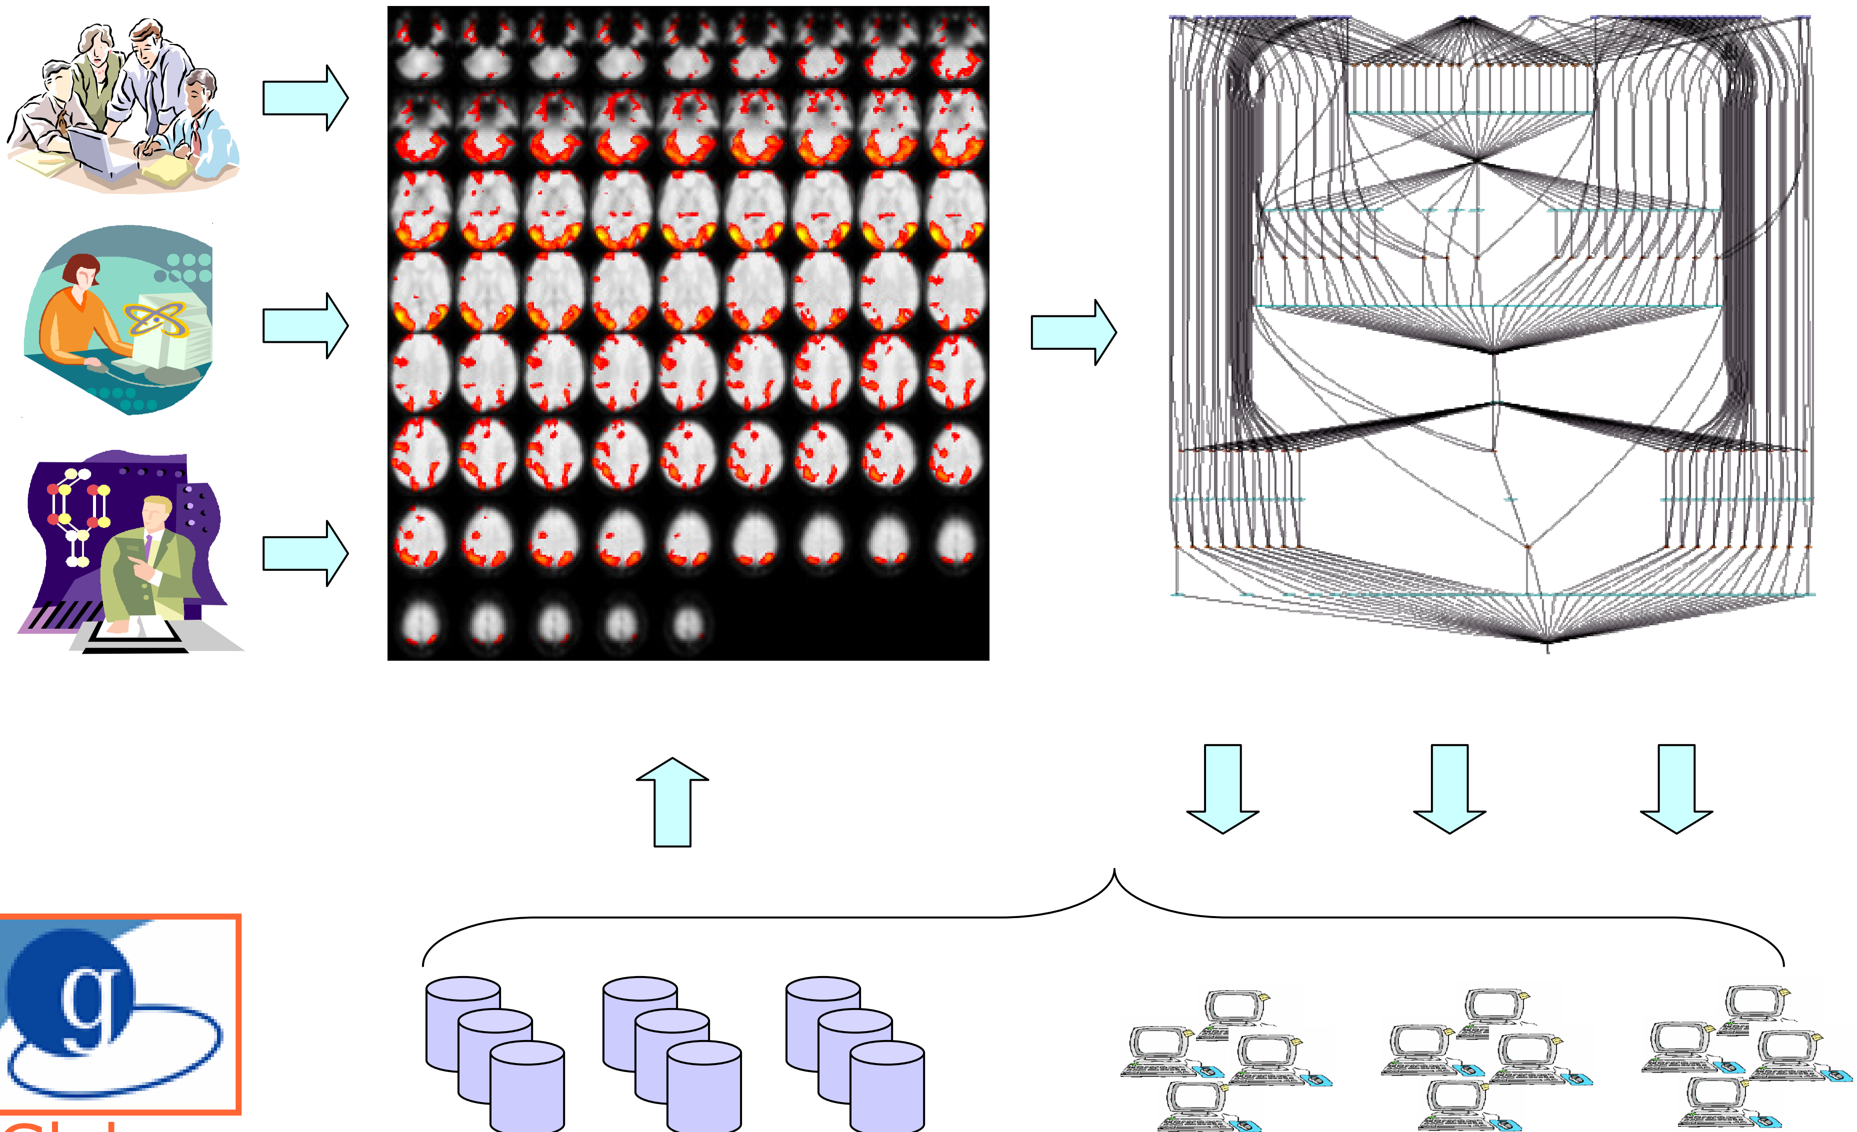

Globus

A Simple Analysis

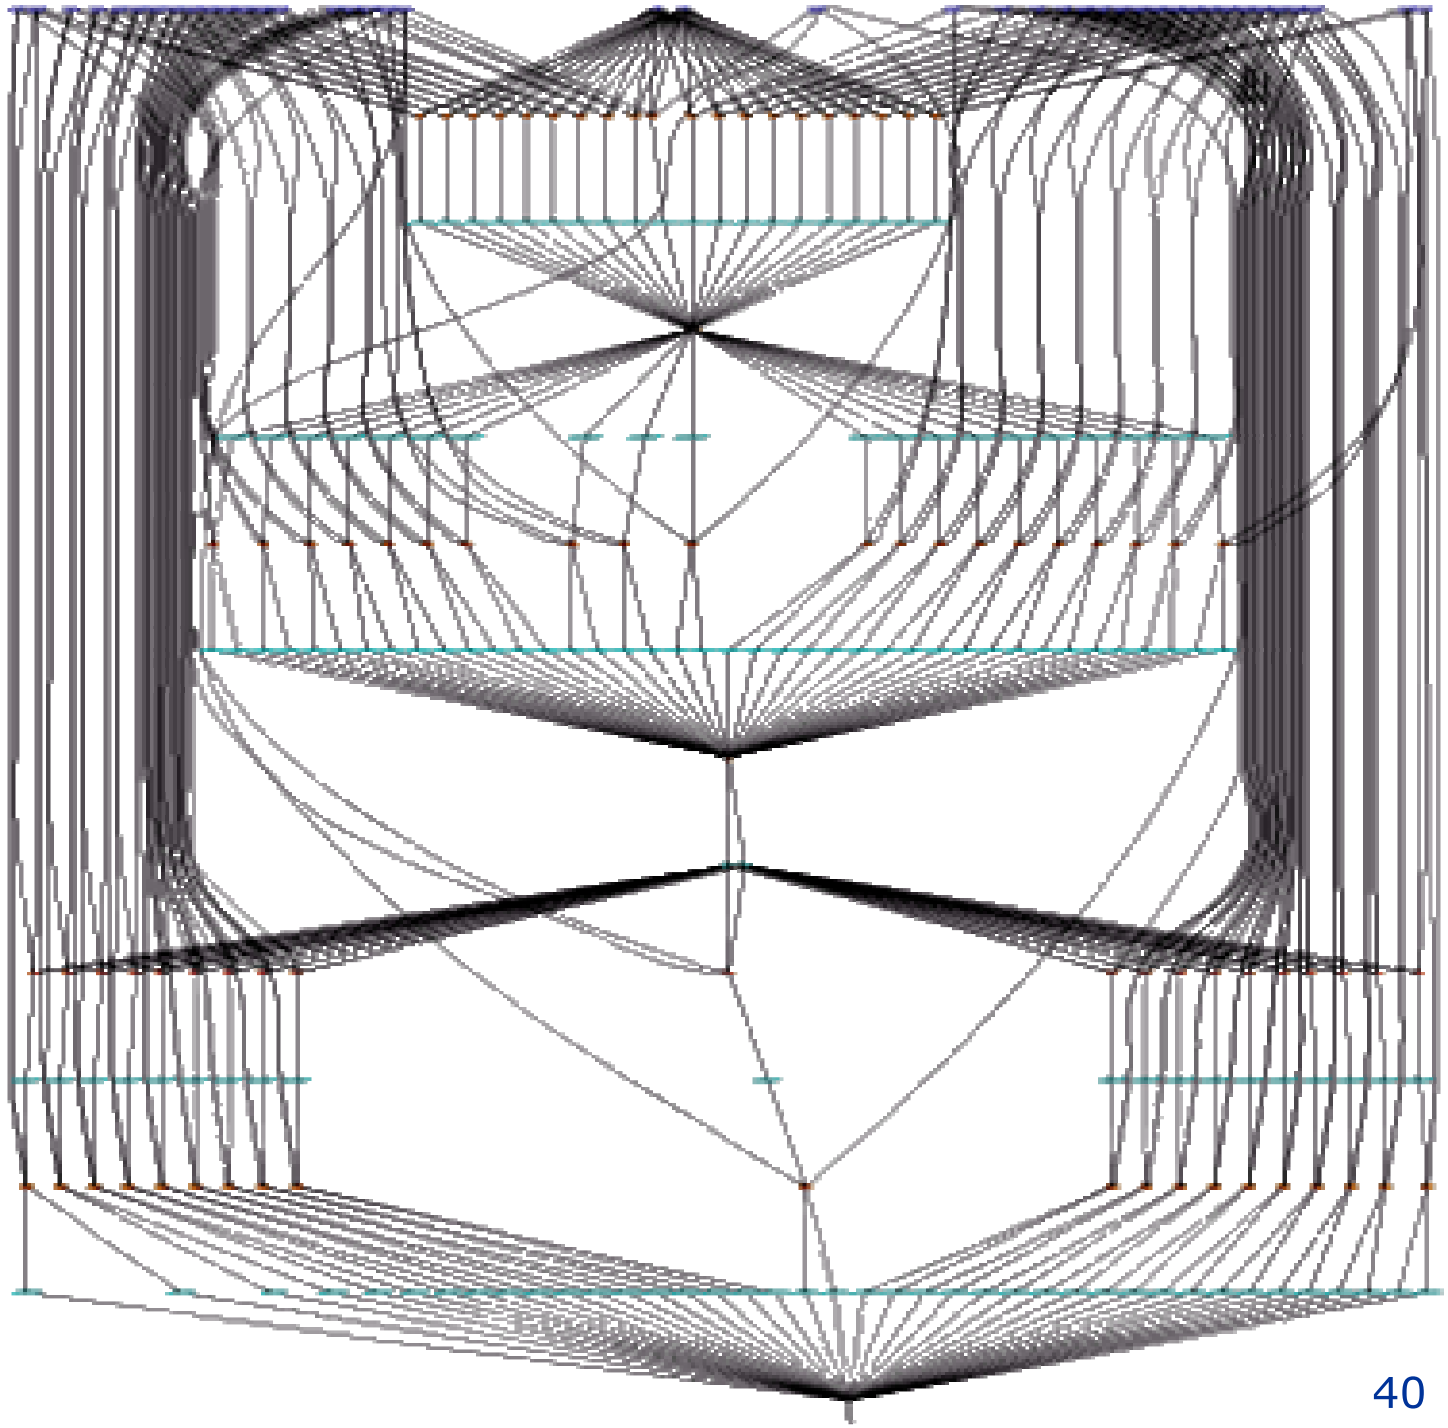

SwiftScript

```
(Run snr) functional ( Run r, NormAnat a,
                       Air shrink ) {
    Run yroRun = reorientRun( r , "y" );
    Run roRun = reorientRun( yroRun , "x" );
    Volume std = roRun[0];
    Run rndr = random_select( roRun, 0.1 );
    AirVector rndAirVec = align_linearRun( rndr, std, 12, 1000, 1000, "81 3 3" );
    Run reslicedRndr = resliceRun( rndr, rndAirVec, "o", "k" );
    Volume meanRand = softmean( reslicedRndr, "y", "null" );
    Air mnQAAir = alignlinear( a.nHires, meanRand, 6, 1000, 4, "81 3 3" );
    Warp boldNormWarp = combinewarp( shrink, a.aWarp, mnQAAir );
    Run nr = reslice_warp_run( boldNormWarp, roRun );
    Volume meanAll = strictmean( nr, "y", "null" )
    Volume boldMask = binarize( meanAll, "y" );
    snr = gsmoothRun( nr, boldMask, "6 6 6" );
}
```

```
(Run or) reorientRun (Run ir,
                      string direction) {
    foreach Volume iv, i in ir.v {
        or.v[i] = reorient(iv, direction);
    }
}
```

www.ci.uchicago.edu/swift

SF-Express
(1998, Caltech & ISI)

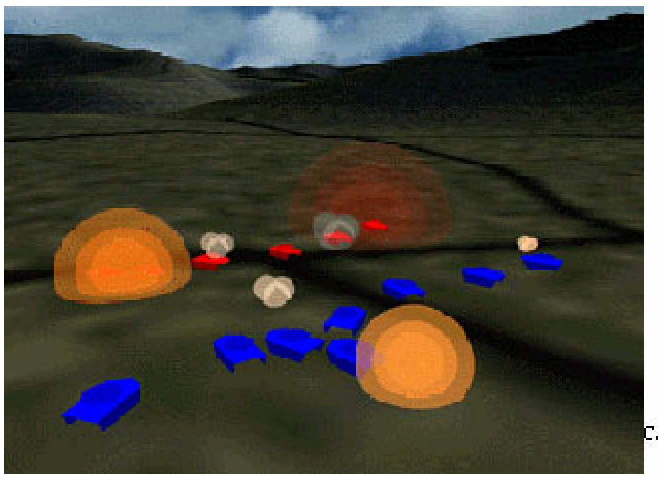

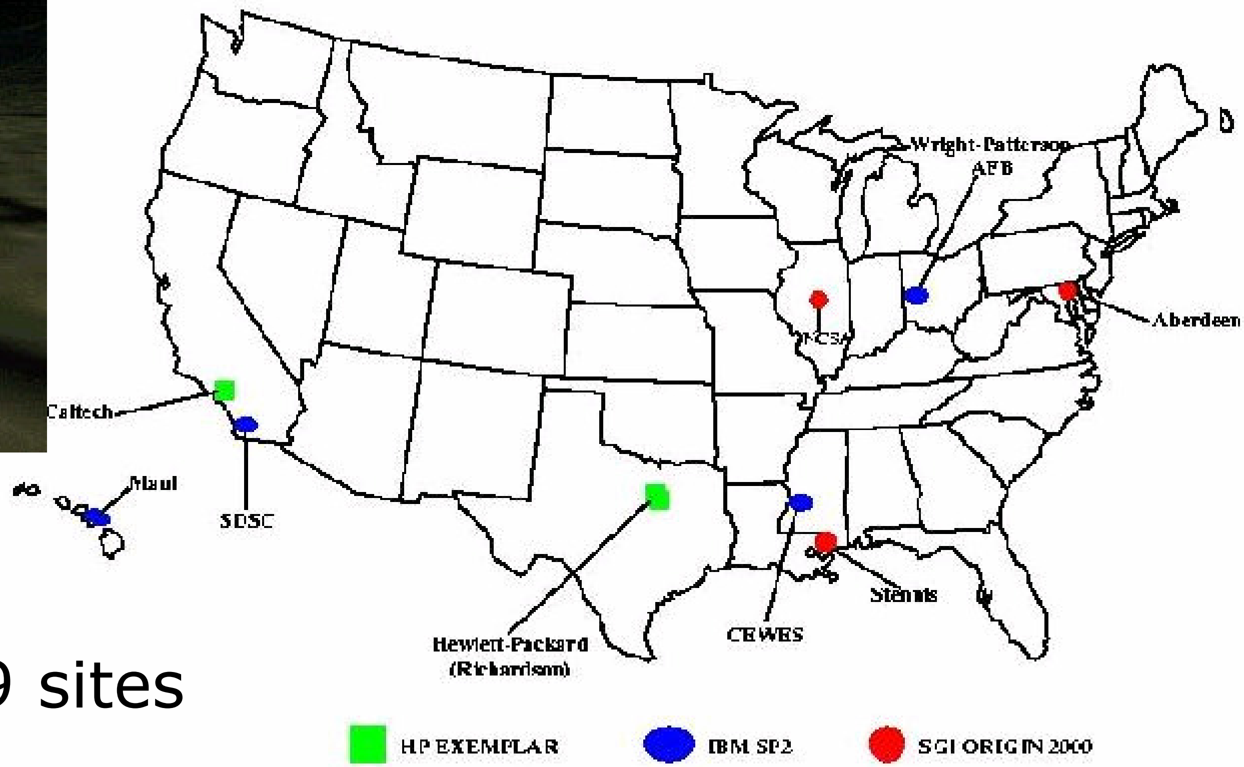

100,000 entities
11 computers @ 9 sites
1386 processors

Site	Platform	Total Processors	Vehicles
ASC	SP	130	10,818
ARL	SGI	60	4,333
ARL	SGI	60	3,347
Caltech	HP	240	21,951
CEWES	SP	232	17,049
HP	HP	128	8,599 *
MHPCC	SP	148	9,485
MHPCC	SP	100	6,796
NAVO	SGI	60	4,238
NCSA	SGI	128	6,693
SDSC	SP	100	6,989
	Totals	**1386**	**100,298**

Digital Human: Simulation of the Human Arterial Tree on the TeraGrid

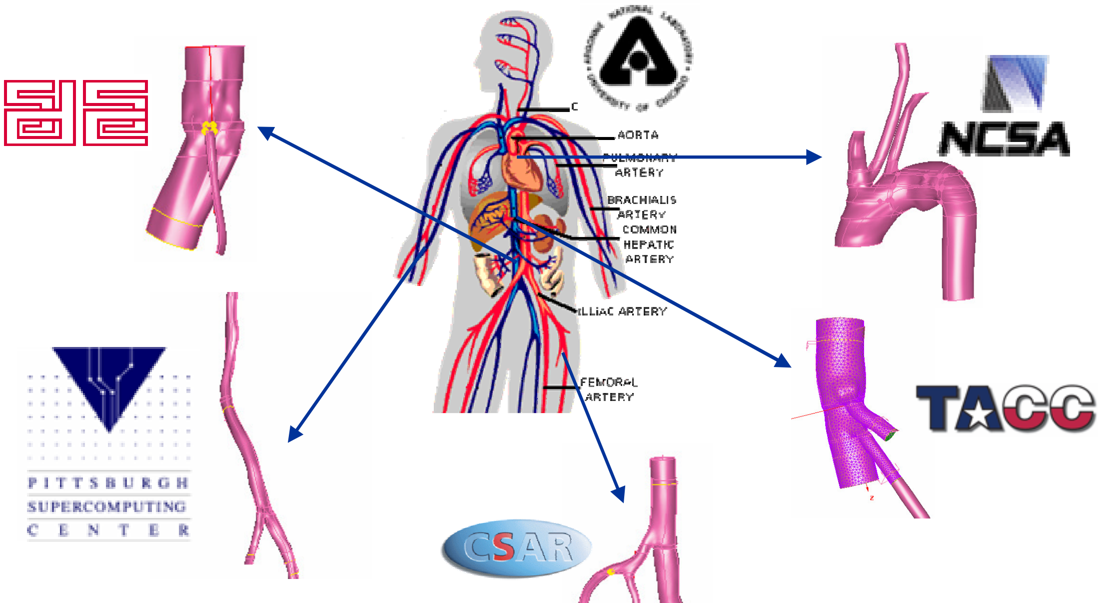

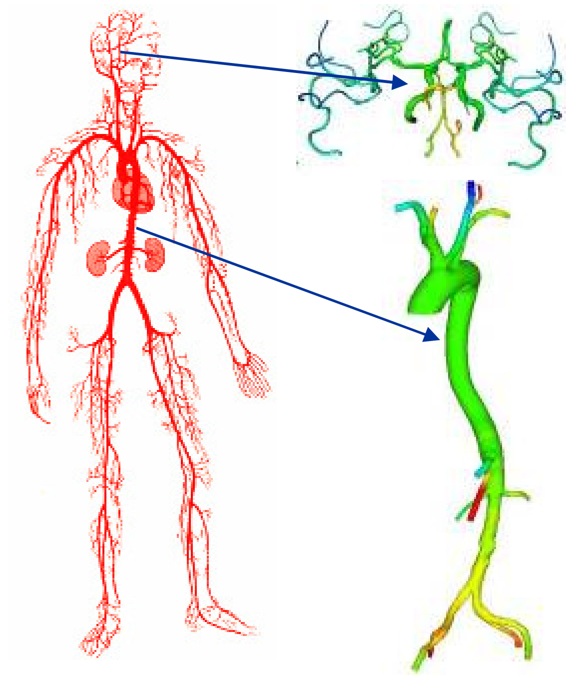

Team Members

Brown University: L. Grinberg[1], S. Dong[2], A. Yakhot, G.E. Karniadakis

Imperial College, London: S.J. Sherwin

Northern Illinois Univ.: N.T. Karonis, J. Insley, J. Binns, M. Papka

[1] L. Grinberg *et al.*, “Spectral/*hp* simulation of the human arterial tree on the TeraGrid”, USNCCM9.

[2] S. Dong *et al.*, “Simulating and visualizing the human arterial system on the TeraGrid”, *Future Generation Computer Systems*, Volume 22, Issue 8, October 2006, pp. 1011 - 1017

Globus

Supported by NSF (IMAG, CI-TEAM and DDDAS)

We Can
Access Computing on Demand

Public PUMA Knowledge Base

Information about proteins analyzed against ~2 million gene sequences

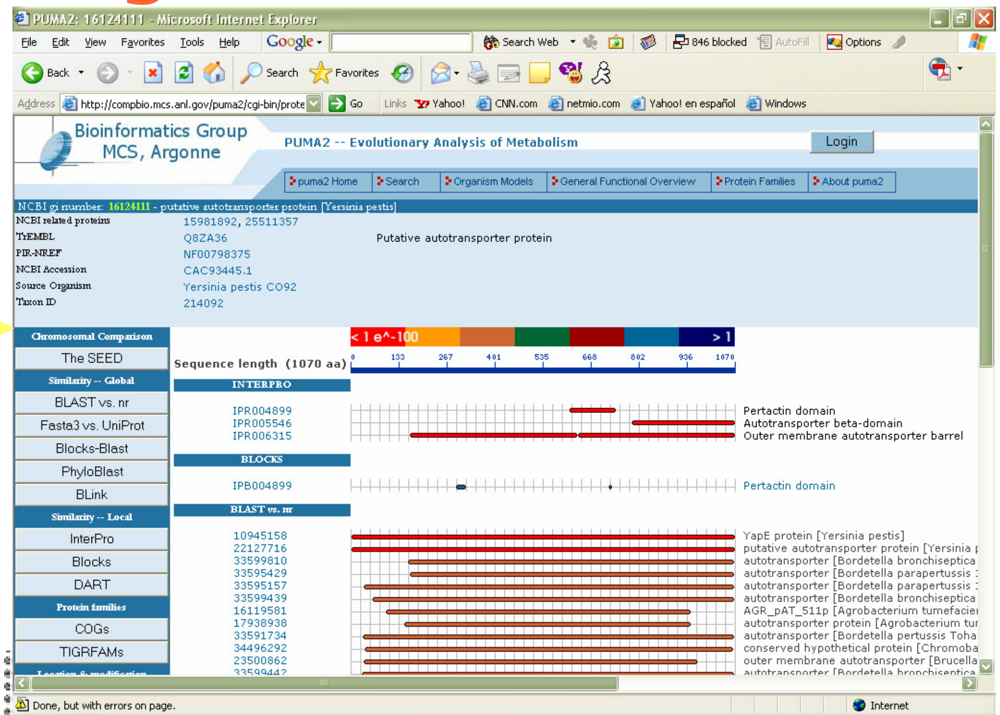

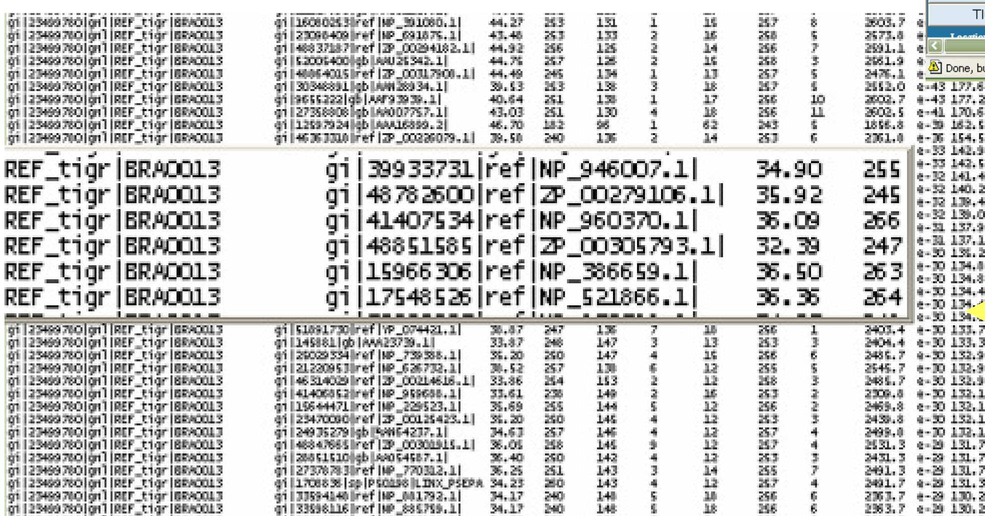

Back Office Analysis on Grid

Millions of BLAST, BLOCKS, etc., on OSG and TeraGrid

Natalia Maltsev et al., http://compbio.mcs.anl.gov/puma2

Second-Generation Grids

- Empower many more users by enabling on-demand access to **services**
- Science gateways (TeraGrid)
- Service oriented science
- Or, "Science 2.0"

"Service-Oriented Science", *Science*, 2005

"Web 2.0"

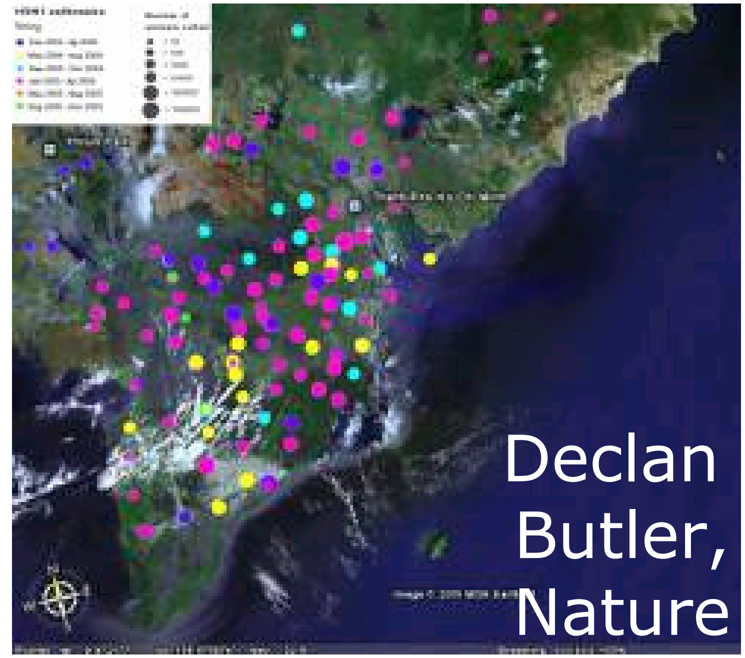

- Software as services
  - Data- & computation-rich network services
- Services as platforms
  - Easy composition of services to create new capabilities ("mashups")—that themselves may be made accessible as new services
- Enabled by massive infrastructure buildout
  - Google projected to spend $1.5B on computers, networks, and real estate in 2006
  - Many others are spending substantially
- Paid for by advertising

Really Big Data Centers

- Some are public
  - Google: Lenoir, NC
  - Microsoft: San Antonio, TX
  - Yahoo: Wenatchee, WA
- Generic attributes
  - 20-60 MW power
  - Near lights out
  - 200K+ square feet
  - Visible from orbit

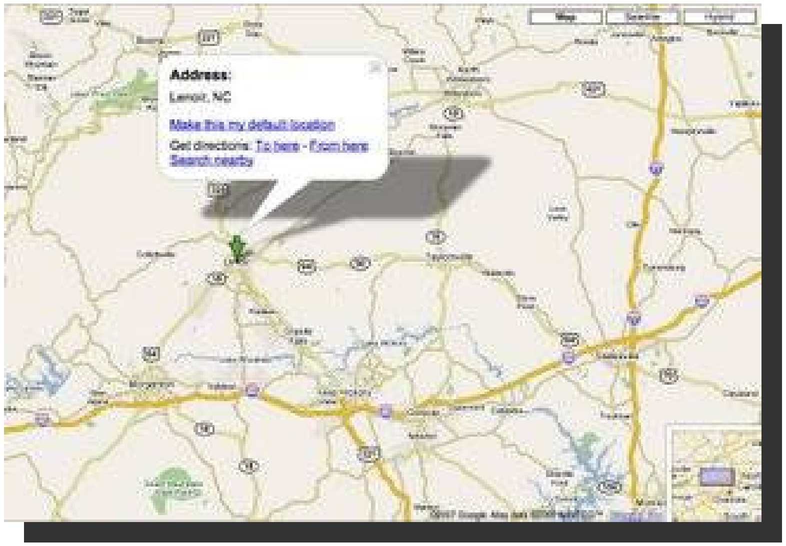

(Slide credit: Dan Reed)

Web Services and Mashups

- Amazon Services
  - Simple Storage Service (S3)
    - web-accessible data
  - Elastic Compute Cloud (EC2)
    - Abstract Machine Image (AMI)
- Yahoo Pipes
  - Compositional mashup
- Google Mashup editor (beta)
  - Service composition
- Microsoft Popfly (alpha)
  - Leverages Silverlight

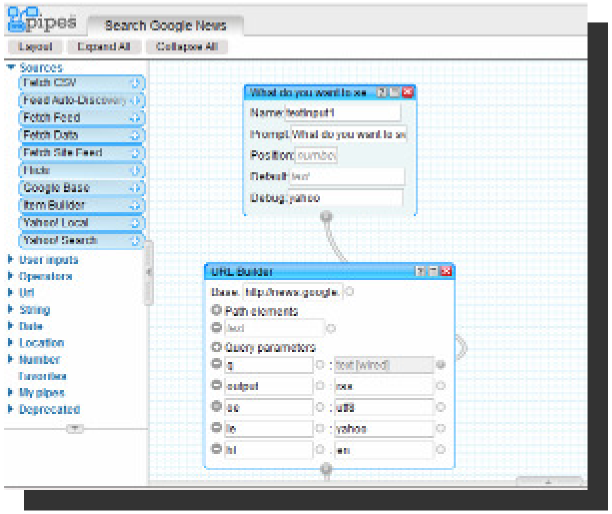

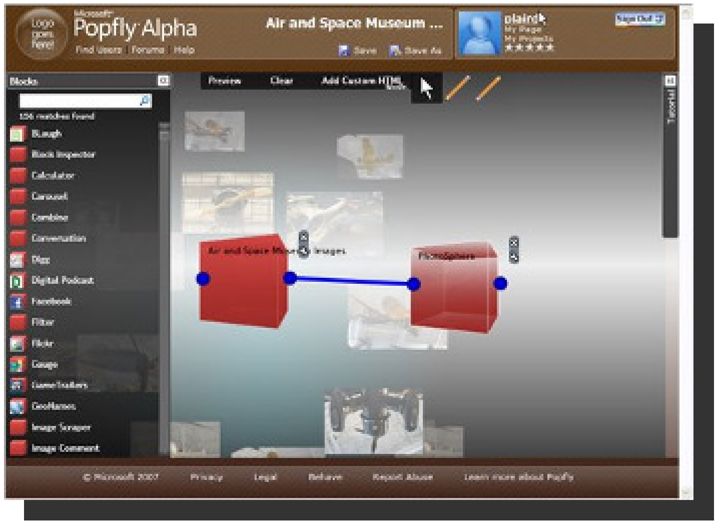

(Slide credit: Dan Reed)

Science 2.0:
E.g., Virtual Observatories

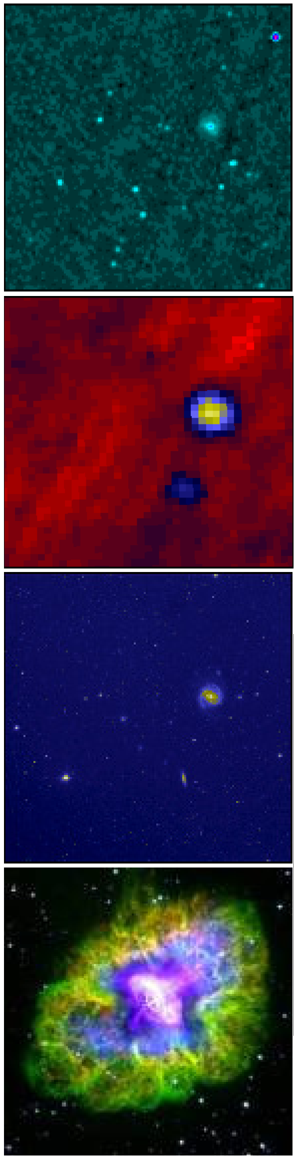

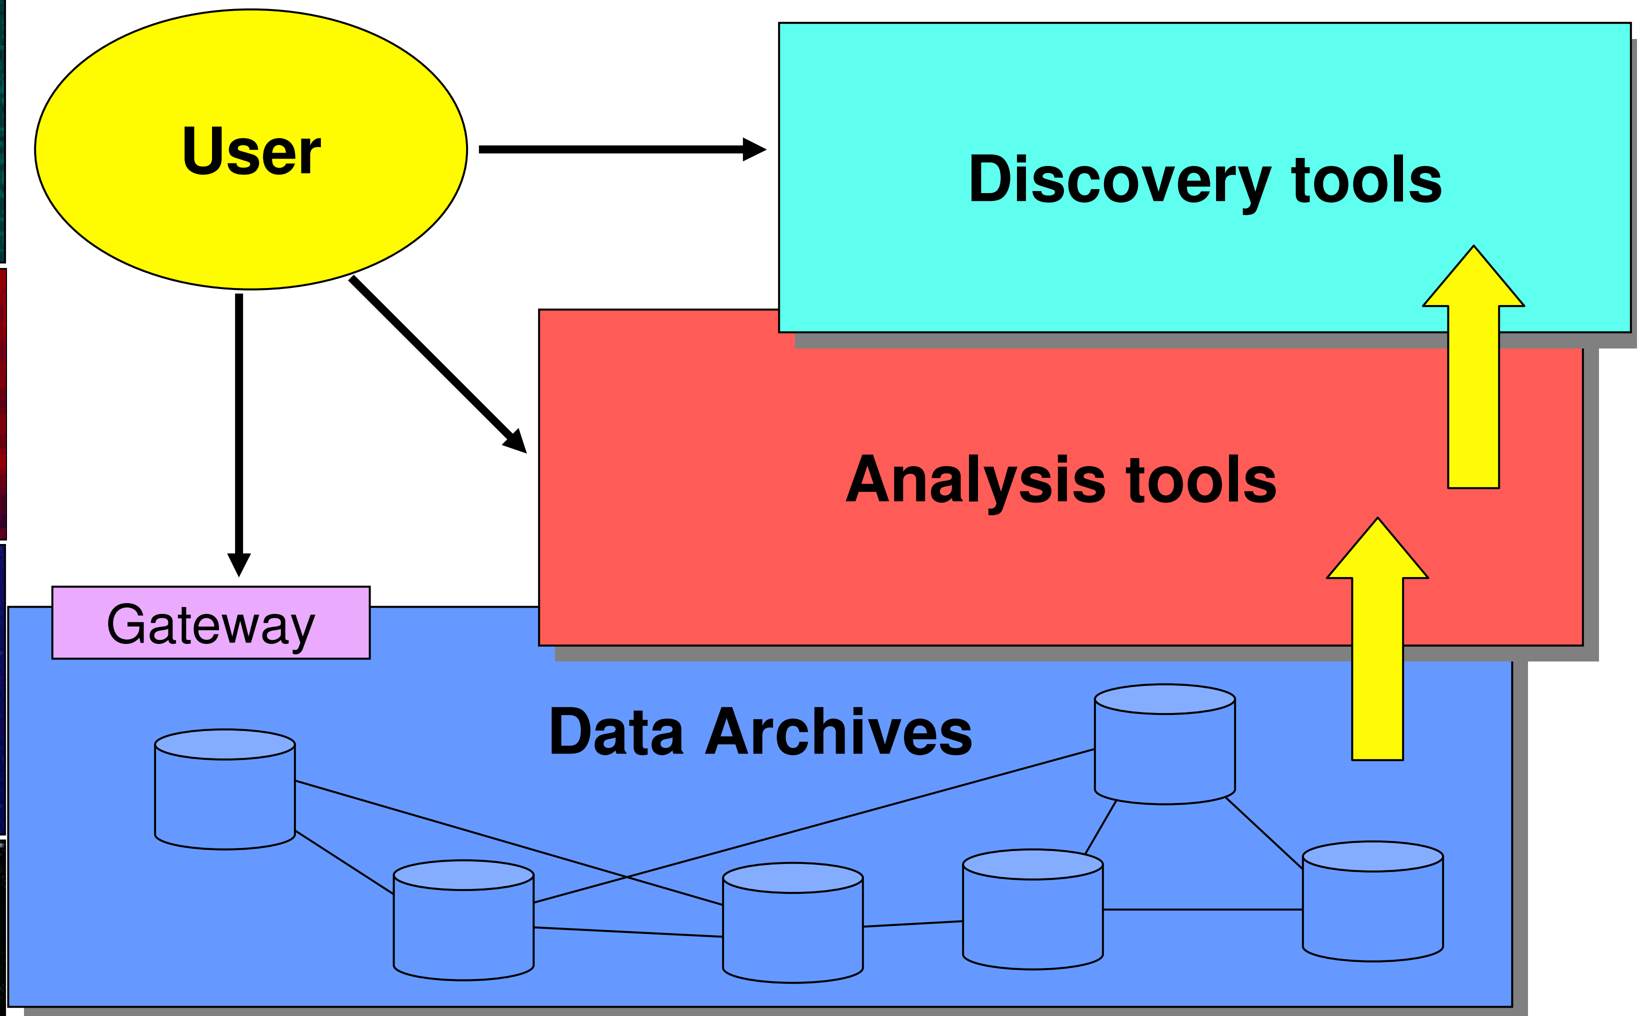

Figure: S. G. Djorgovski

Service-Oriented Science

People **create** services (data or functions) …
which I **discover** (& decide whether to use) …
& **compose** to create a new function ...
& then **publish** as a new service.

→ *I find "someone else" to **host** services, so I don't have to become an expert in operati services & computers!*

!! → *I hope that this "someone else" can **manage** security, reliability, scalability, …*

"Service-Oriented Science", *Science*, 2005

Service-Oriented Science & Cancer Biology

caBIG: sharing of infrastructure, applications, and data.

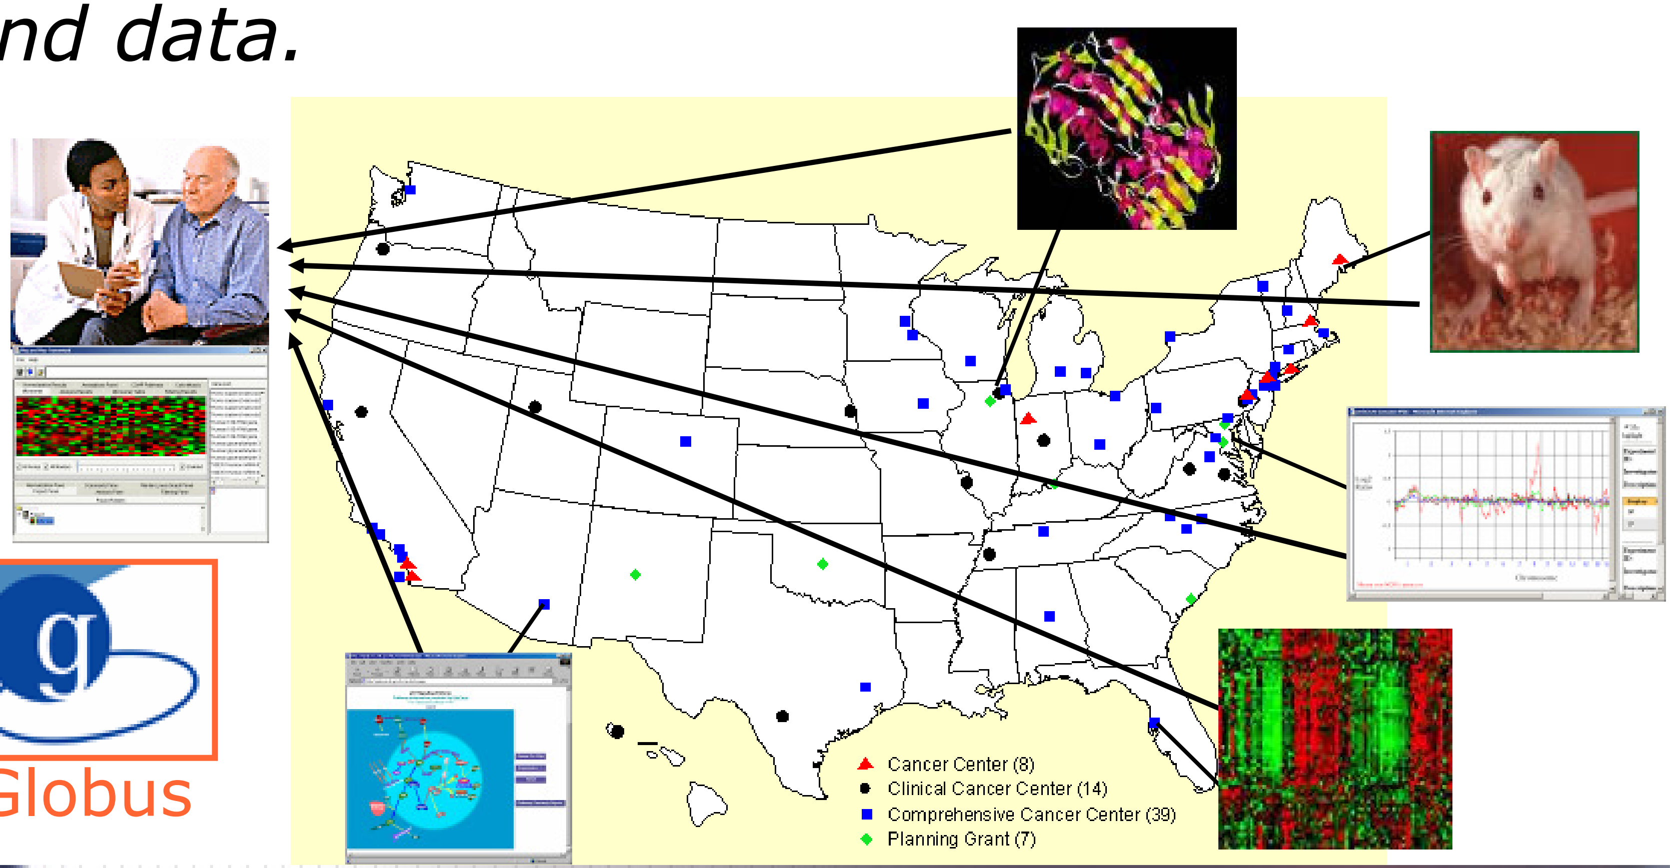

Globus

caBIG cancer Biomedical Informatics Grid

Integrating Thoughts

- Cross-cutting trends
  - High performance via massive parallelism
  - On-demand computing in grid cloud
  - Enormous increases in available data
  - Urgent need to understand social dynamics


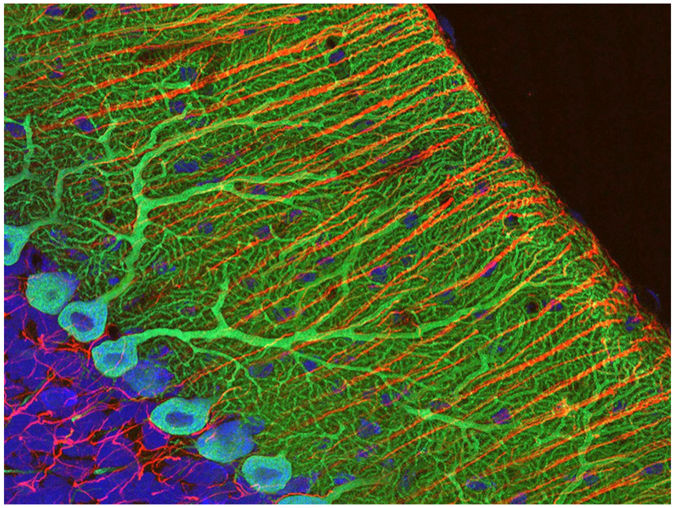


A Golden Age for the Social Sciences?

"Cosmology used to be regarded as a pseudo science, an area where wild speculation, was unconstrained by any reliable observations"

— Steven Hawking, 1999

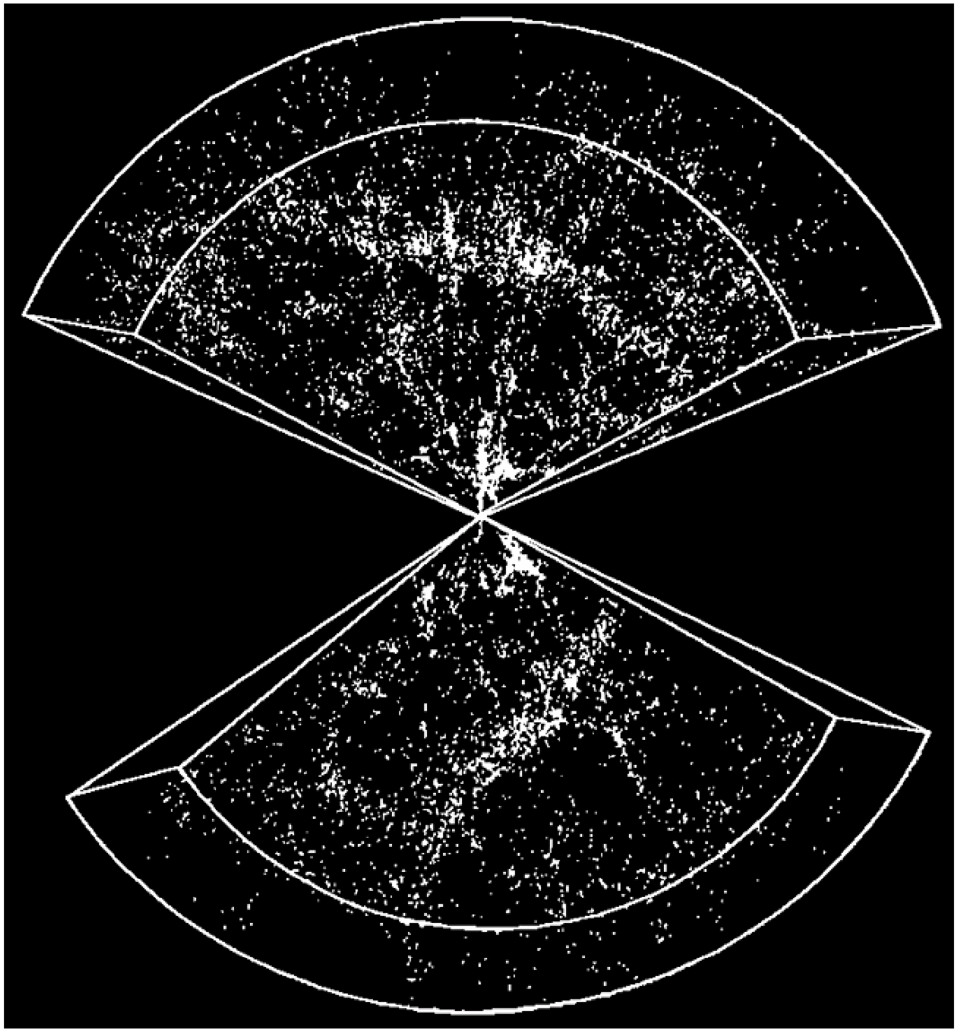

“We can turn meteorology into a science”
— Edward Teller

(advocating his “Brilliant Eyes” satellite program)

Approach

- Develop modeling methods & frameworks able to use the largest supercomputers
  - Large numbers of agents, sensitivity studies
  - Hybrid discrete-continuum methods

- Construct ultra-high-resolution datasets for validation & analysis
  - Include benchmark problems to enable comparison of alternative approaches

Scaling Parallel Discrete Event Simulation on IBM BG/L

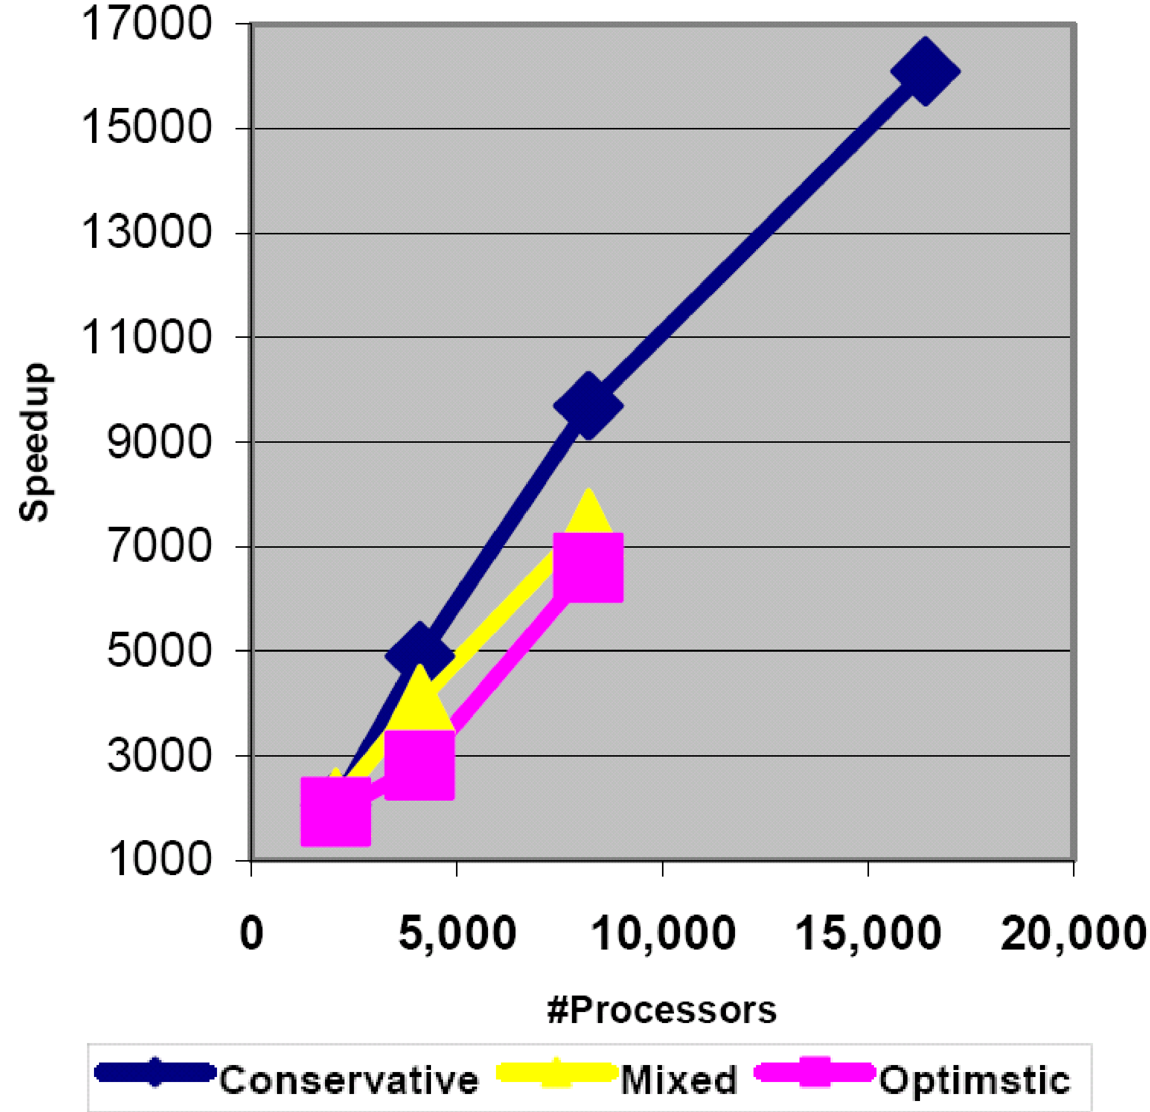

Results for PHOLD with 10^6 entities, 10^7 events

Scaling relative to 2000 nodes

530 million events per wall clock second on 16384 processors

(Kalyan Perumalla, Oak Ridge Nat Lab)

Argonne
Exascale ABMS Project

- Target future exascale (10^{18} op/s) systems
- Develop scalable agent modeling framework
- Multiple target applications
    - Microbial ecosystems
    - Cybersecurity response
    - Energy system dynamics

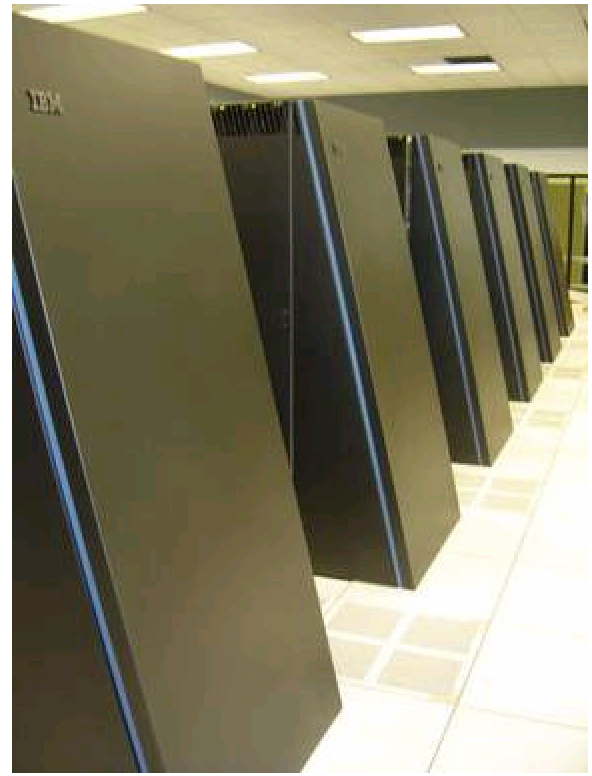

(North, Stevens, Macal, Papka, Sallach)

Microbial Ecosystem Problem

- 10^6 to 10^9 individuals
- 10^2 to 10^3 possible internal states
- 10^2 to 10^4 types
- Interacting via 10^2 to 10^4 messages and/or compounds and substrates
- 10^2 spatially varying resource types

(North, Stevens, Macal, Papka, Sallach)

For Example: Argonne Exascale ABMS Project

Number of Processors	Bacteria Agents		Network Agents		Human Agents	
	Basic	Detailed	Basic	Detailed	Basic	Detailed
1	10^5	10^5	10^3	1	10^4	1
10^6	10^{11}	10^{11}	10^8	10^7	10^{10}	10^6
10^7	10^{12}	10^{12}	10^9	10^8	10^{11}	10^7
10^8	10^{13}	10^{13}	10^{10}	10^9	10^{12}	10^8

(North, Stevens, Macal, Papka, Sallach)

Concluding Thoughts

- Exponentials are changing what we can compute & measure in fundamental ways
- It may become feasible to study previously inaccessible systems and phenomena
- Will require "grand challenge" efforts with both computational & discipline scientists
- Employ service oriented science methods to link communities and perform outreach
- **A golden age for computational social sciences and agent-based modeling?**