NSFNET: Impact on Science and Engineering

Kenneth A. Bishop, Ph.D.
Professor Emeritus of Chemical & Petroleum Engineering
The University of Kansas
kbishop@ku.edu

November 29, 2007

NSFNET
The Partnership That Changed The World

The World According to K. A. B.

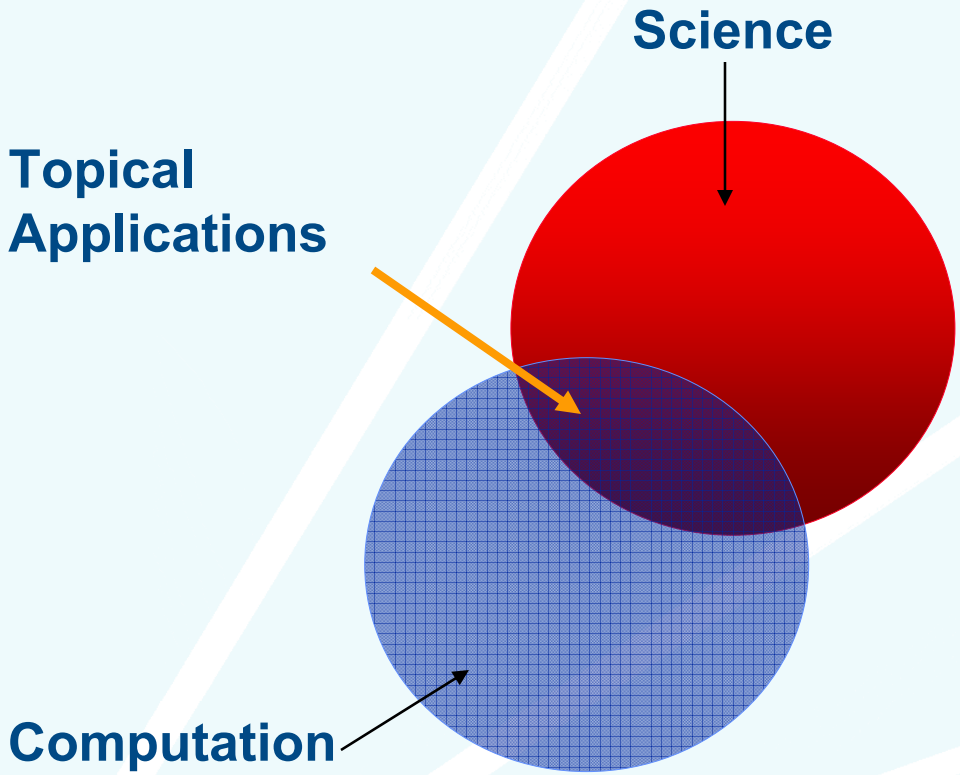

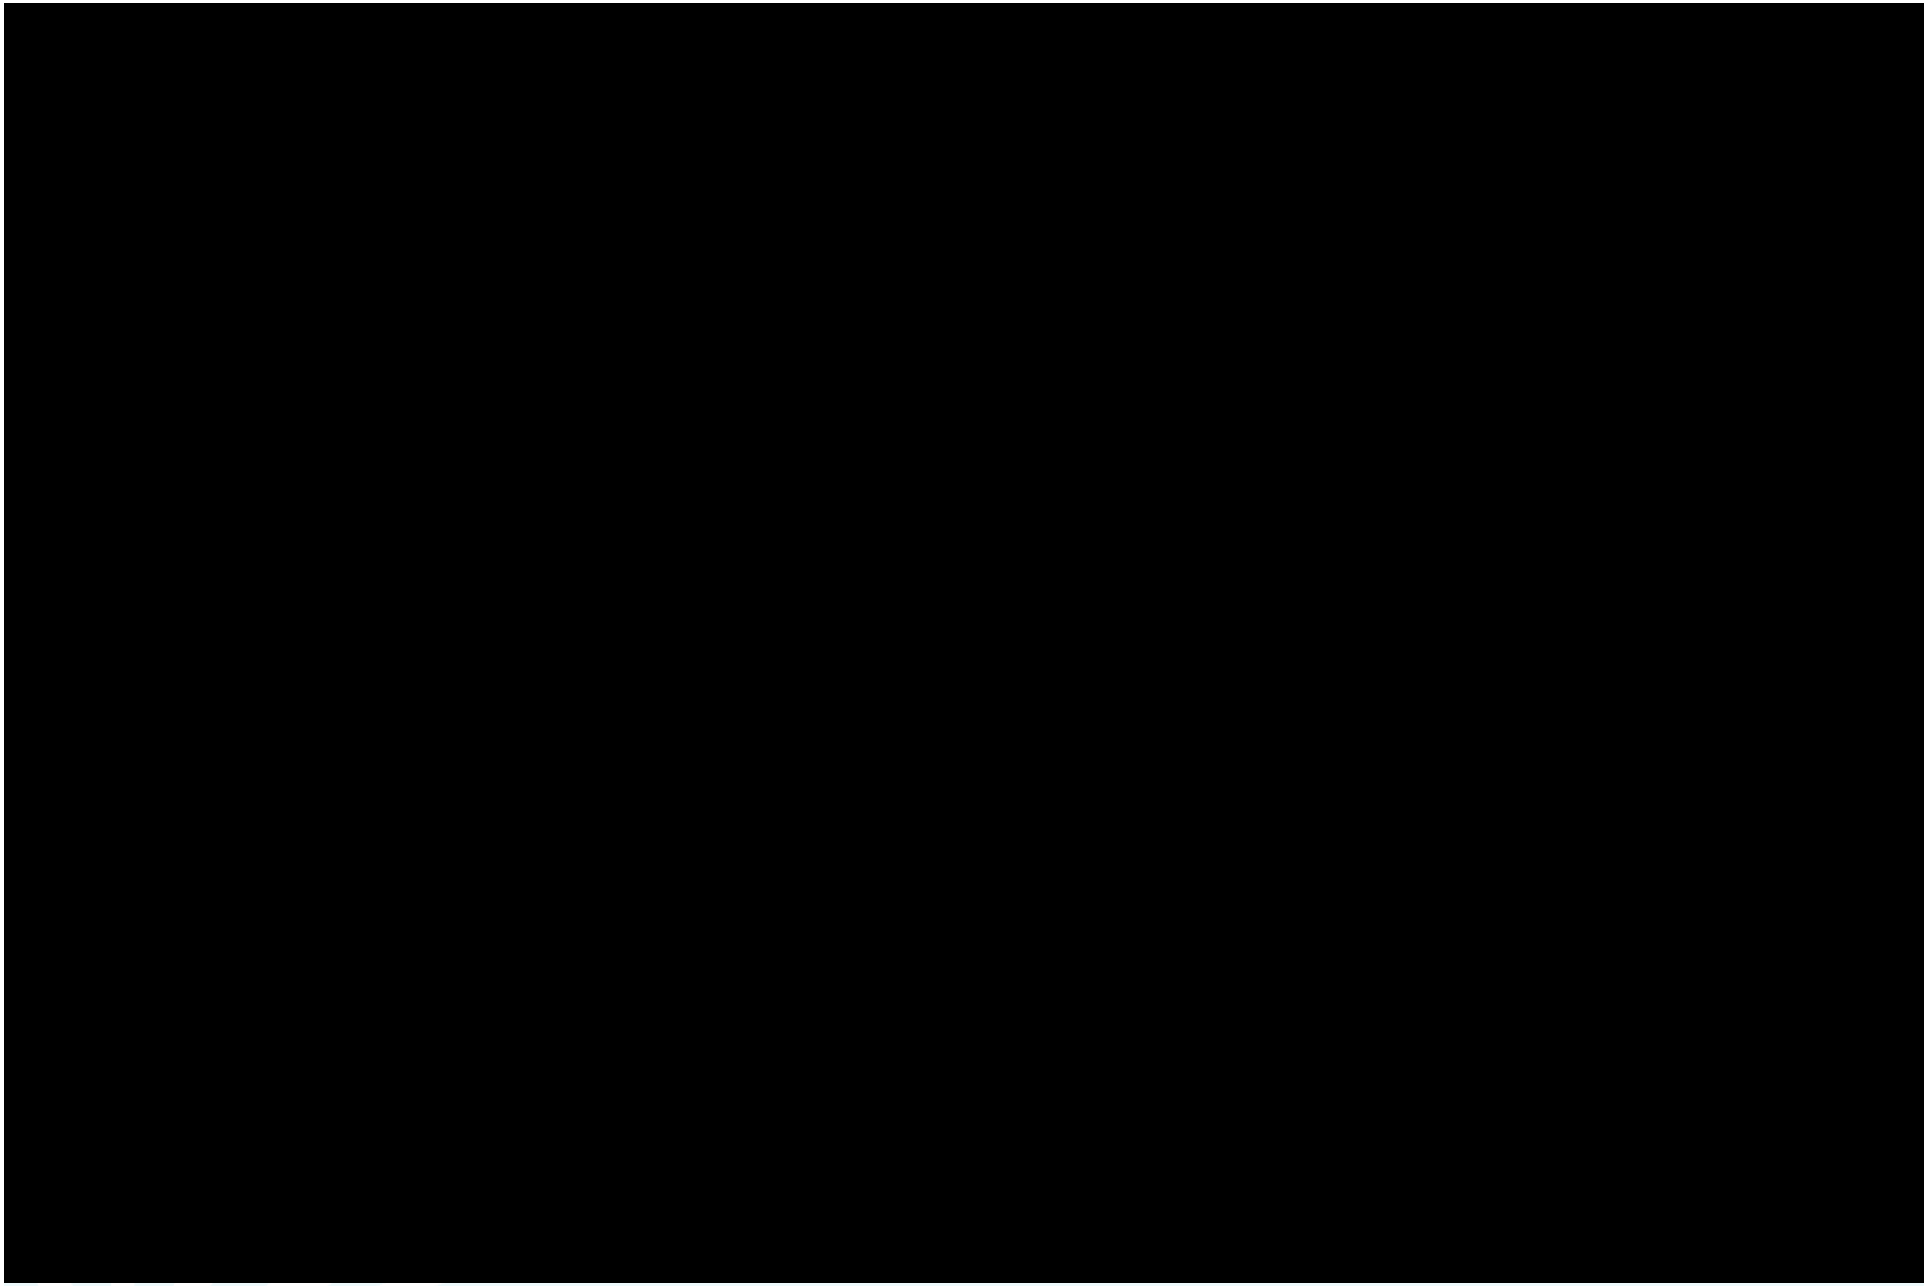

NSFNET

The Partnership That
Changed The World

Prototype for the 21st Century Information Infrastructure

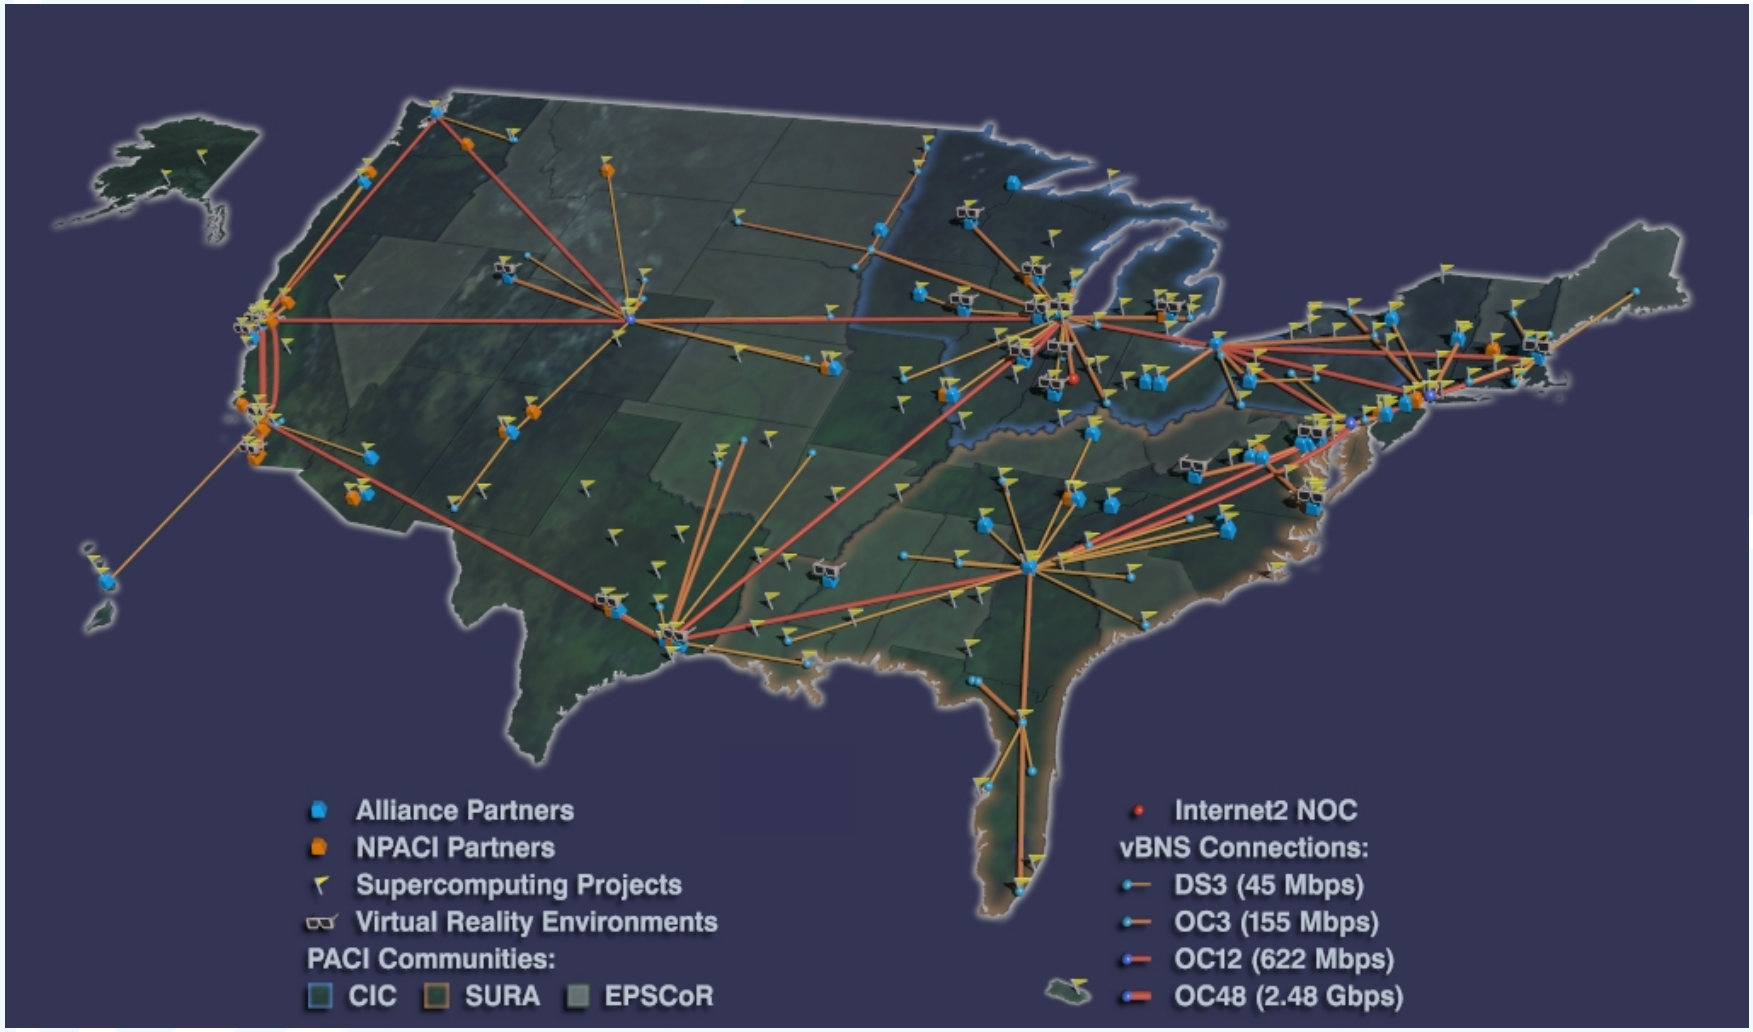

NSFNET
The Partnership That Changed The World

The National Technology Grid

The World According to K. A. B.

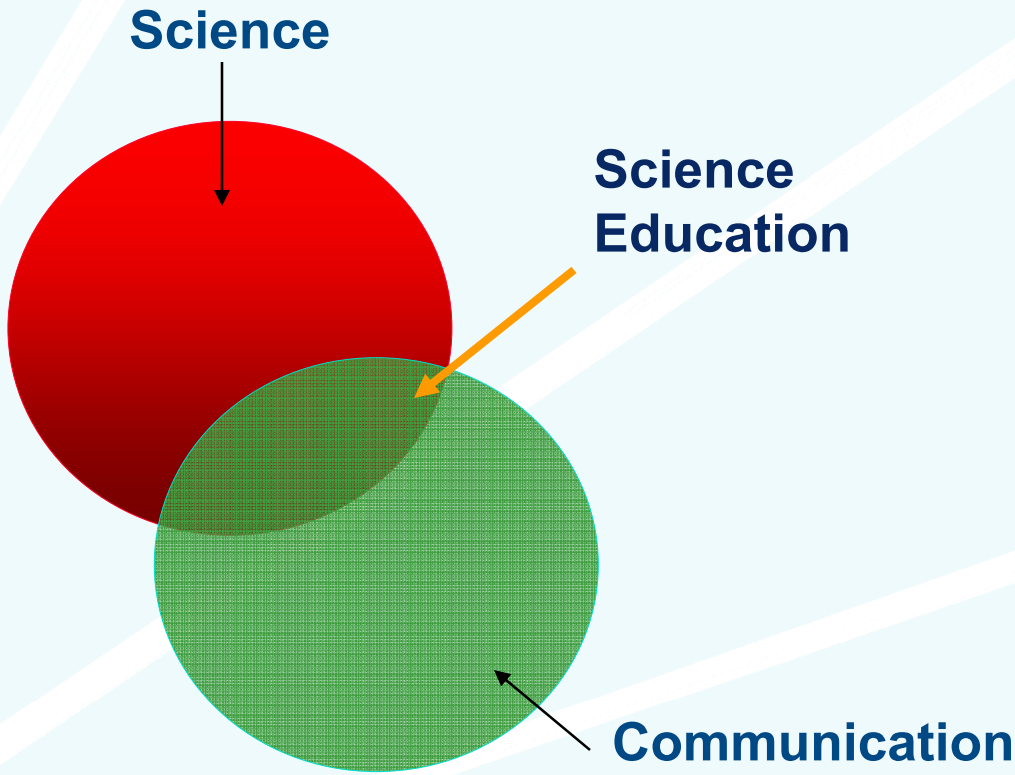

NSFNET
The Partnership That
Changed The World

Introduction to the Access Grid

- A set of hardware and software technologies used to enable multisite audio and video conferencing over the internet
- Typical uses include:
  - distributed seminars/conferences
  - informal meetings
  - collaborative research

NSFNET
The Partnership That Changed The World

Distributed Seminars/Conferences

NSFNET
The Partnership That
Changed The World

ALLIANCE

CHAUTAUQUA 2000

Innovations in Science, Computation, and Access Grid Technologies

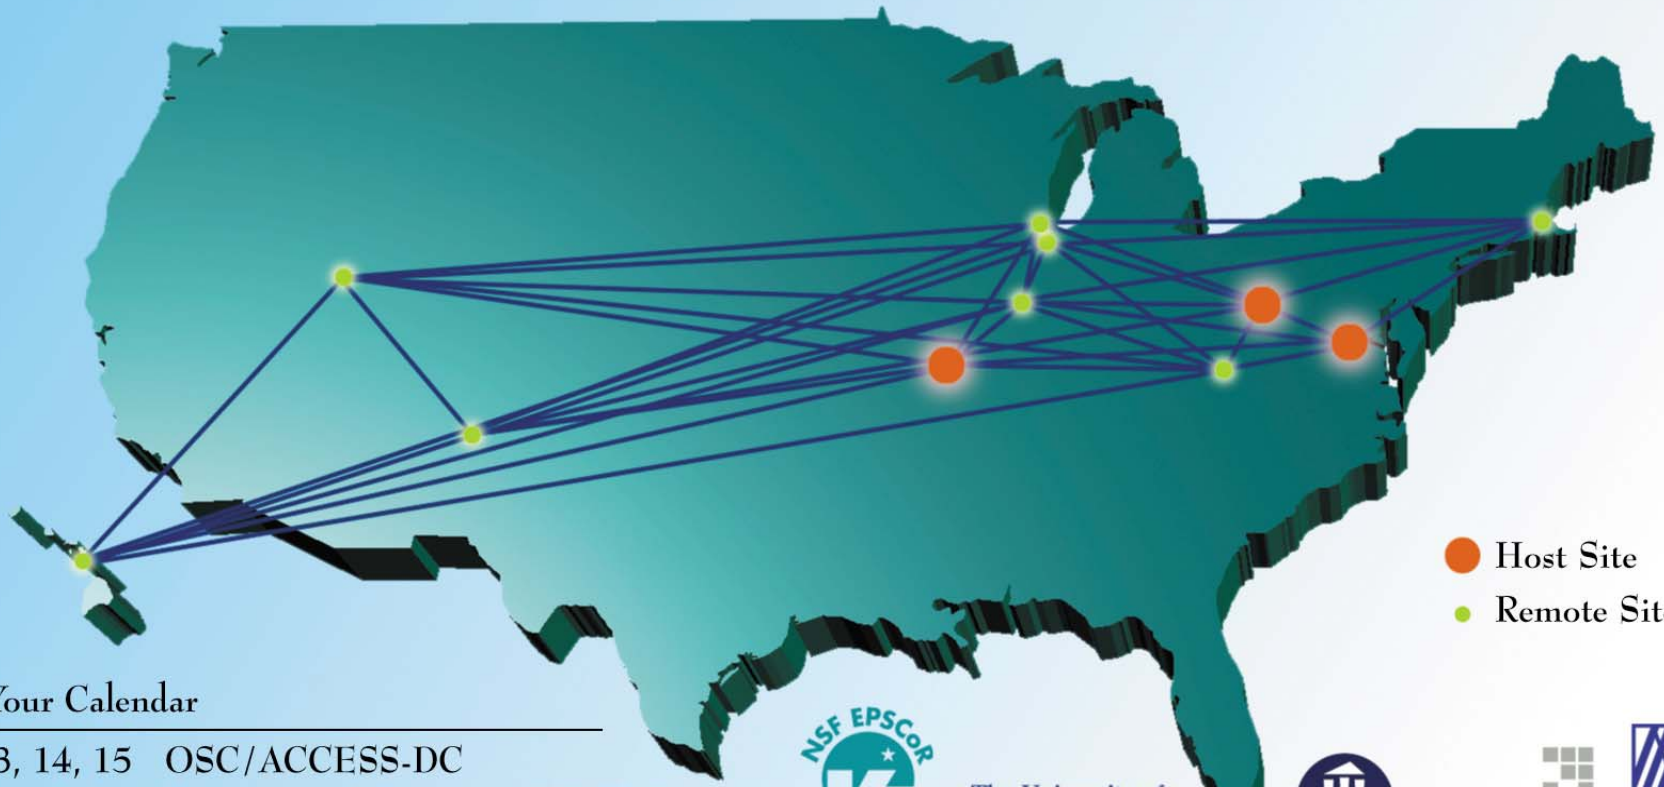

Mark Your Calendar

June 13, 14, 15 OSC/ACCESS-DC
August 1, 2, 3 University of Kansas/ACCESS-DC
http://www.ncsa.uiuc.edu/alliance/chautauqua/

NSF EPSCoR K STAR

The University of KANSAS

OSC

ALLIANCE

NSFNET
The Partnership That Changed The World

Informal Meetings

NSFNET
The Partnership That
Changed The World

Aquifer Simulation

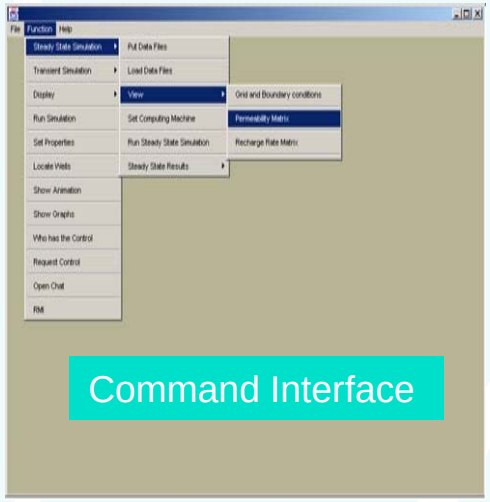

Command Interface

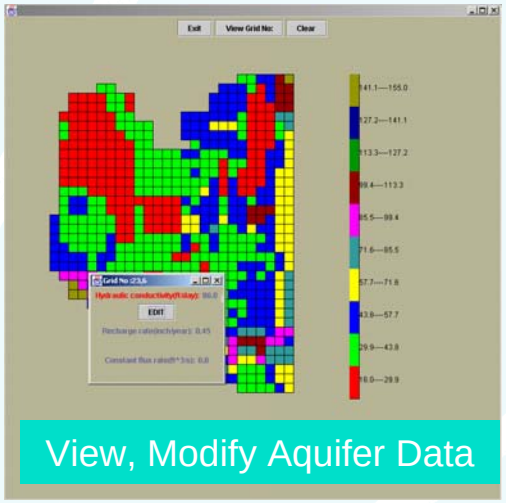

View, Modify Aquifer Data

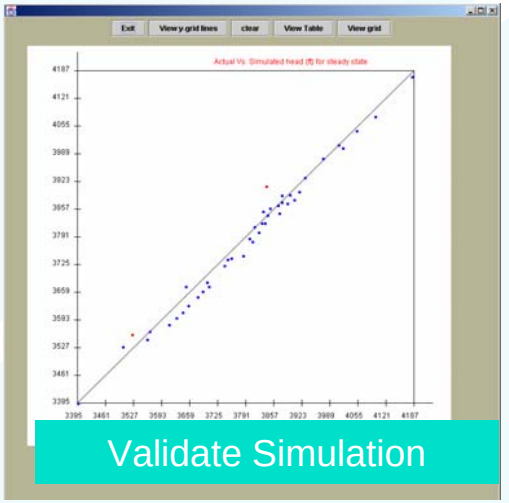

Validate Simulation

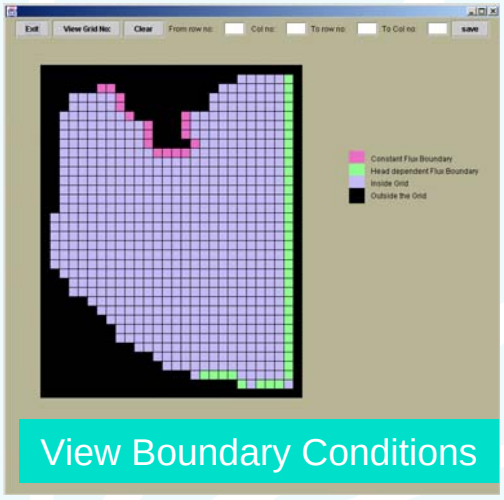

View Boundary Conditions

comparison_table

Row Number	Column Number	Measured head(ft)	My Simulated head...	%error	User's val...
7	6	4255.0	4174.3	1.9	0.0
8	9	4183.0	3848.3	8.0	0.0
9	5	4355.0	4045.1	7.1	0.0
9	11	4111.0	3900.8	5.1	0.0
9	22	3830.0	3873.0	-1.1	0.0
10	19	3905.0	3913.4	-0.2	0.0
11	11	4100.0	3788.1	7.6	0.0
11	14	4006.0	3881.8	3.1	0.0
12	19	3857.0	3803.2	1.4	0.0
12	19	3863.0	3724.8	3.6	0.0
13	10	4093.0	3979.3	2.8	0.0
13	12	4050.0	3738.3	7.7	0.0
14					
14					
14	19	3810.0	3747.8	1.6	0.0
14	19	3802.0	3663.1	3.7	0.0
15	16	3880.0	4011.0	-3.4	0.0

Tabular Results

Simulated Water Level

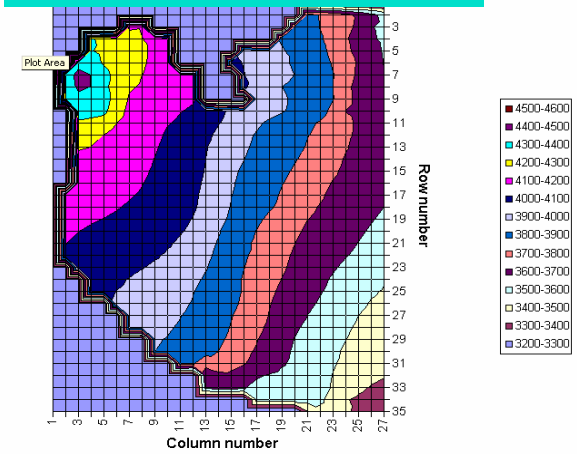

NSFNET
The Partnership That Changed The World

Externally Coupled Calculations

Finite Difference PETSc Code (Continuum Calculations)

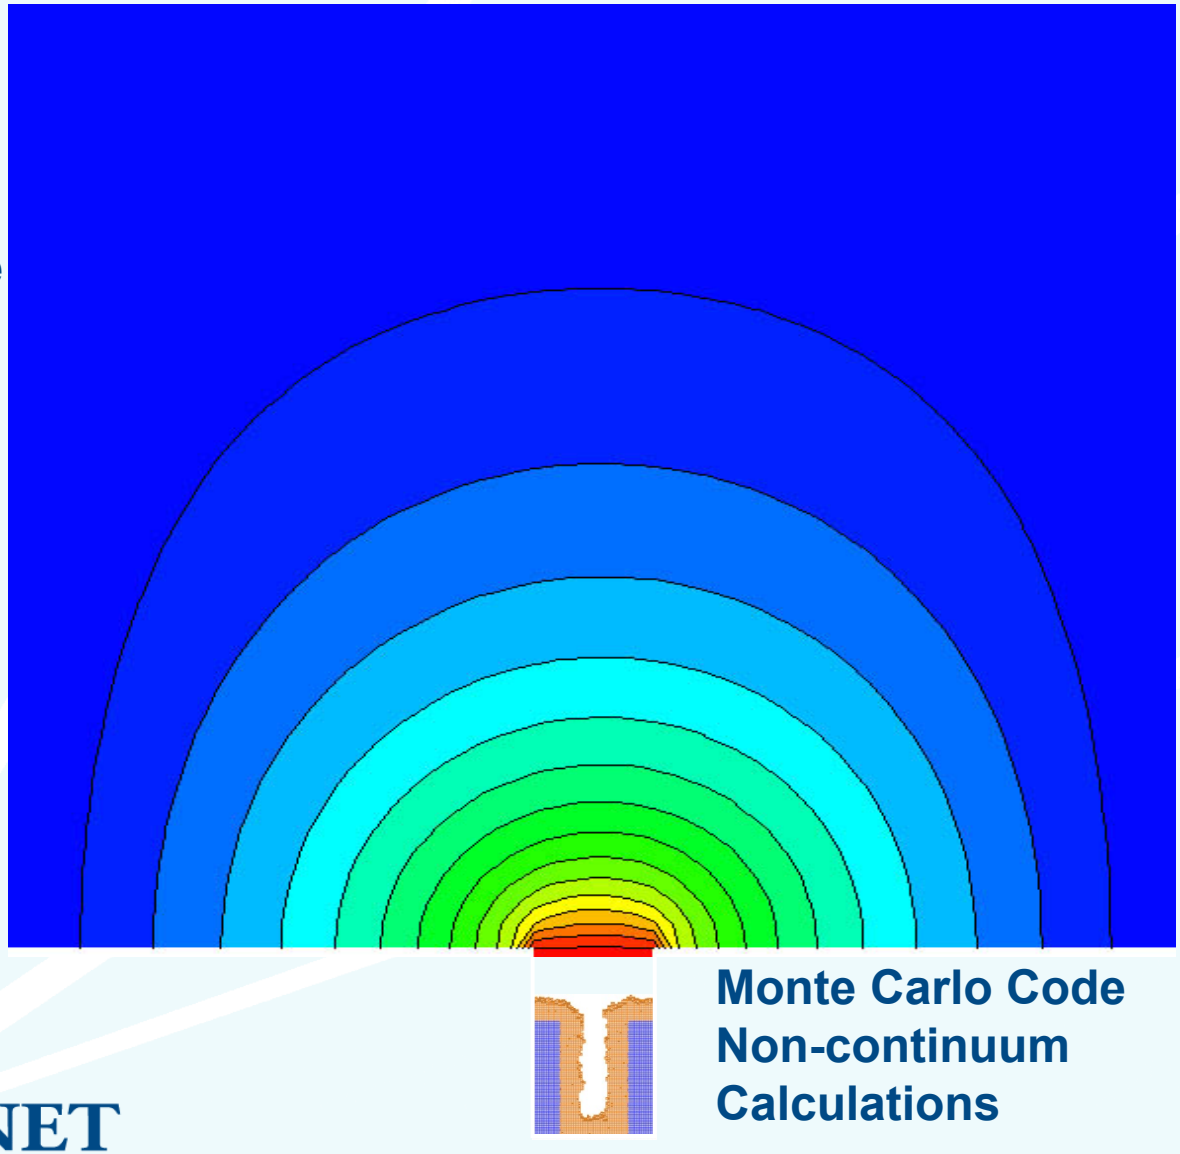

Monte Carlo Code Non-continuum Calculations

NSFNET
The Partnership That Changed The World

Collaborative Research

NSFNET
The Partnership That Changed The World

Chemical Reactor Design

Reaction Conditions:
Temperature: 640 ~ 770 K
Pressure: 2 atm

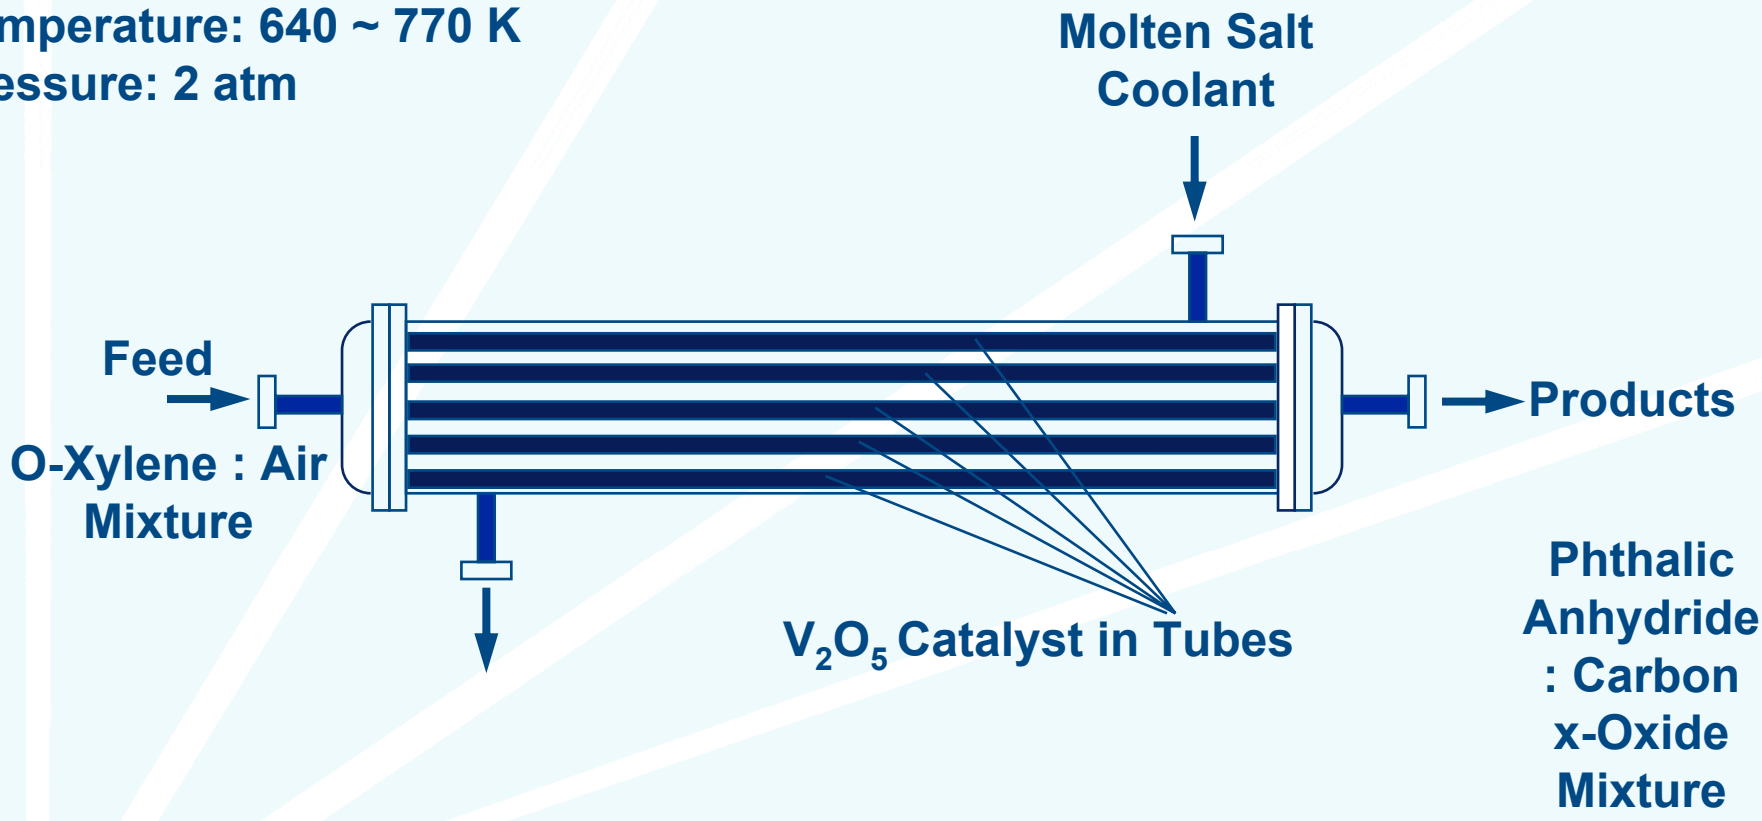

Steady State Temperature Distribution

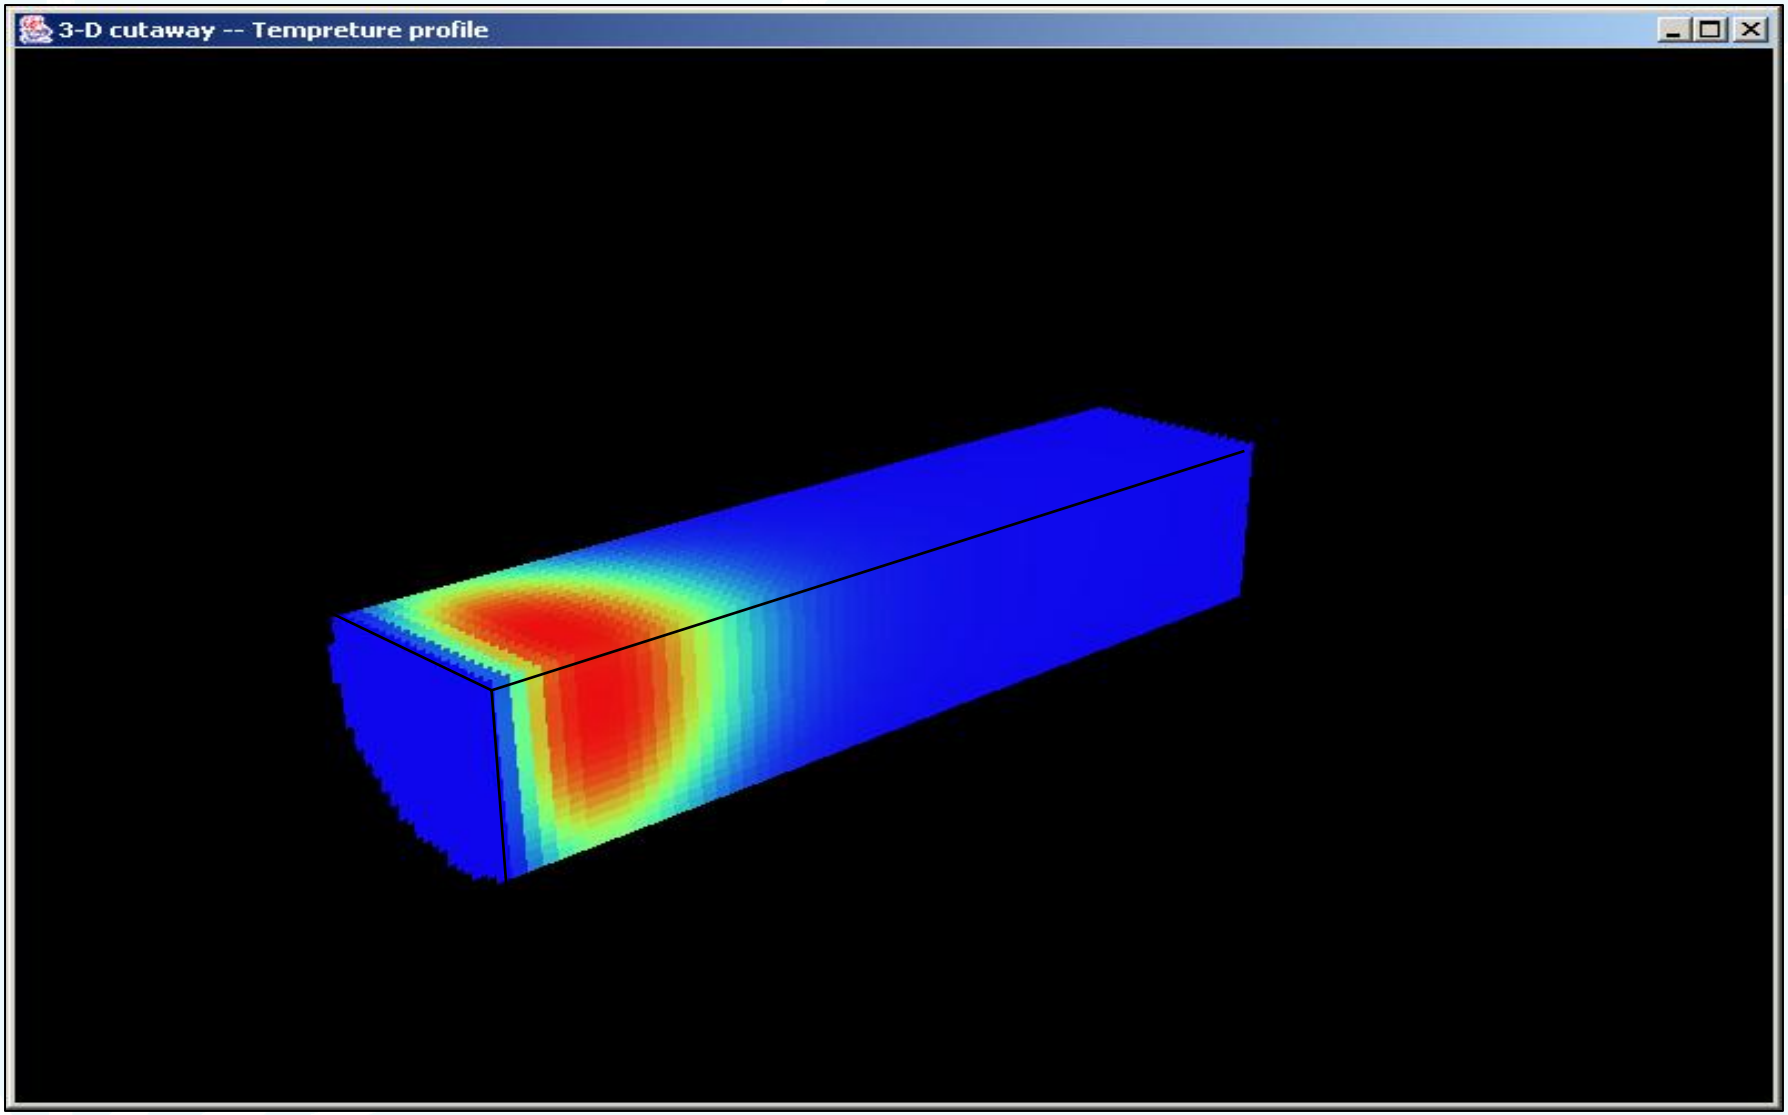

NSFNET
The Partnership That
Changed The World

Cut-away Visualization

Reactor Performance: $t = \infty$

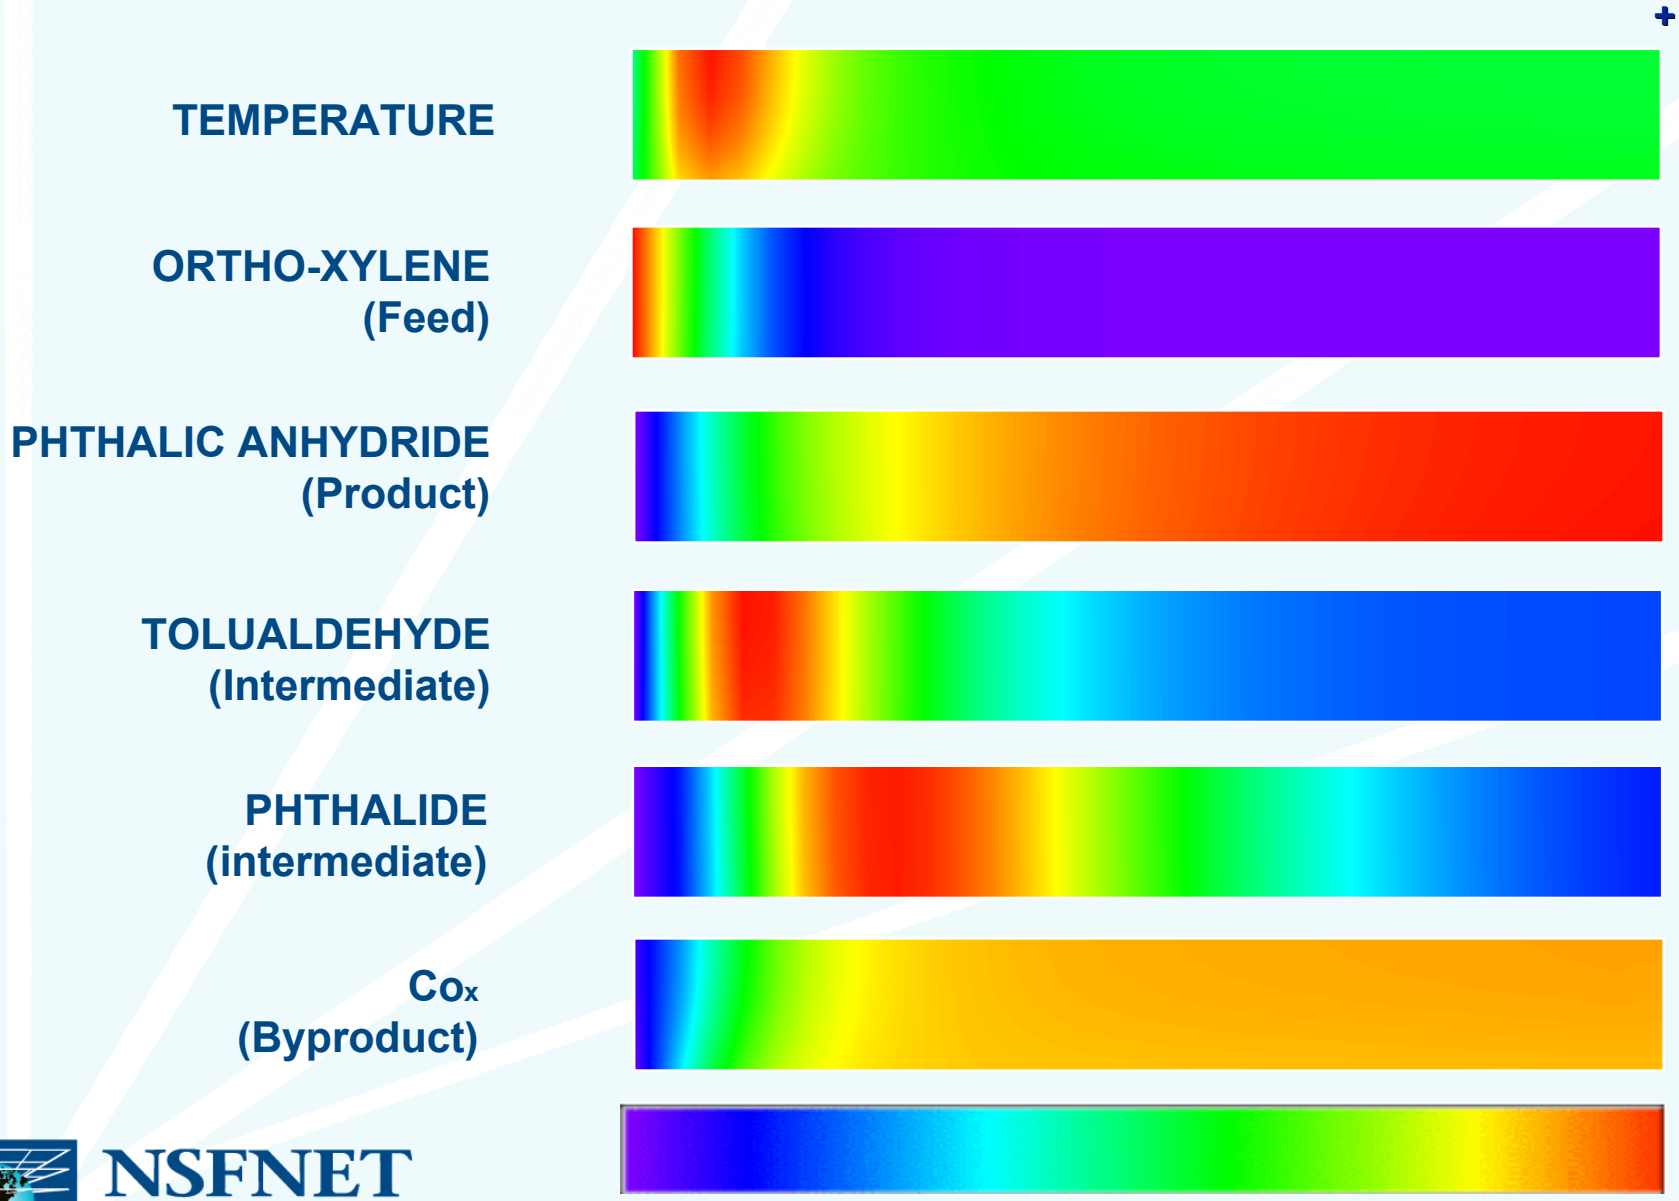

Collaboration Among Software

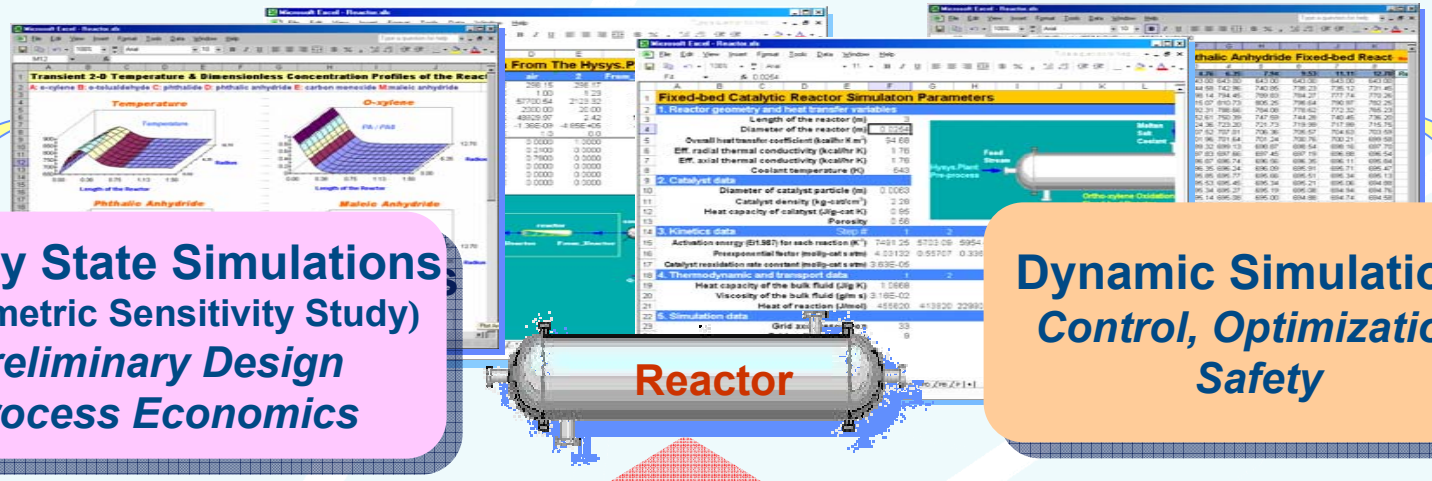

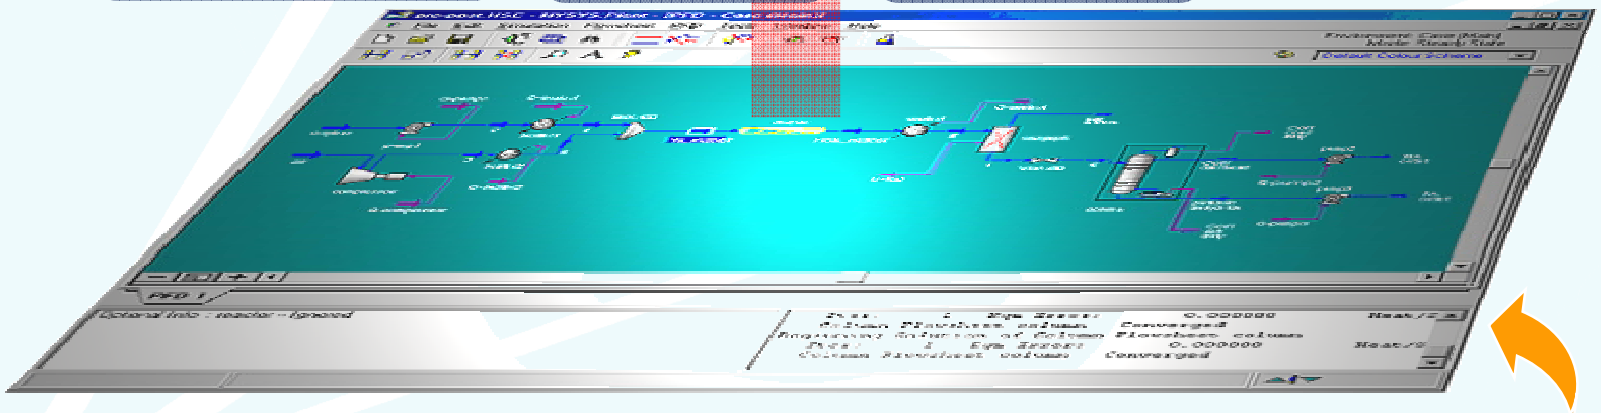

Process Design and Optimization

NSFNET
The Partnership That
Changed The World

The World According to K. A. B.

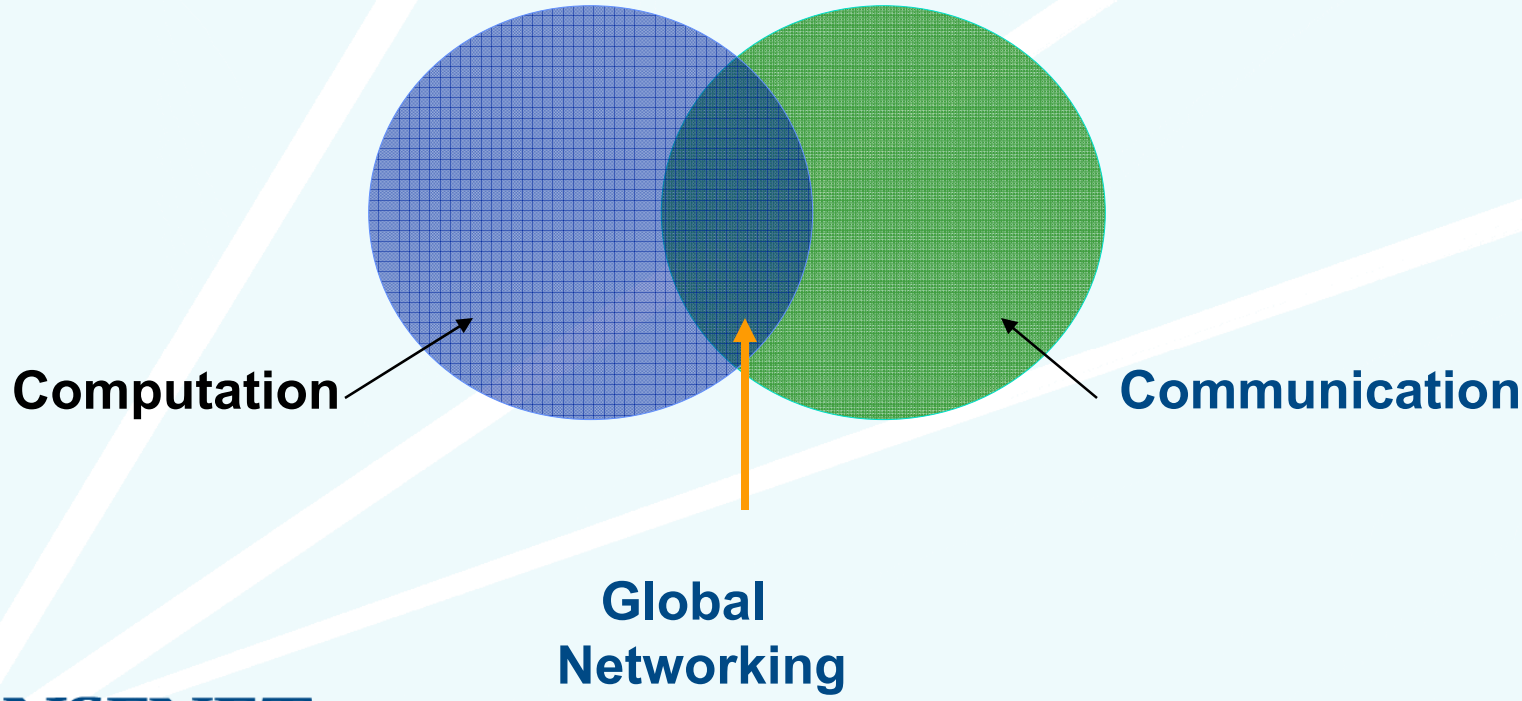

NSFNET
The Partnership That
Changed The World

From the NSFNET to Today's Commercial Internet

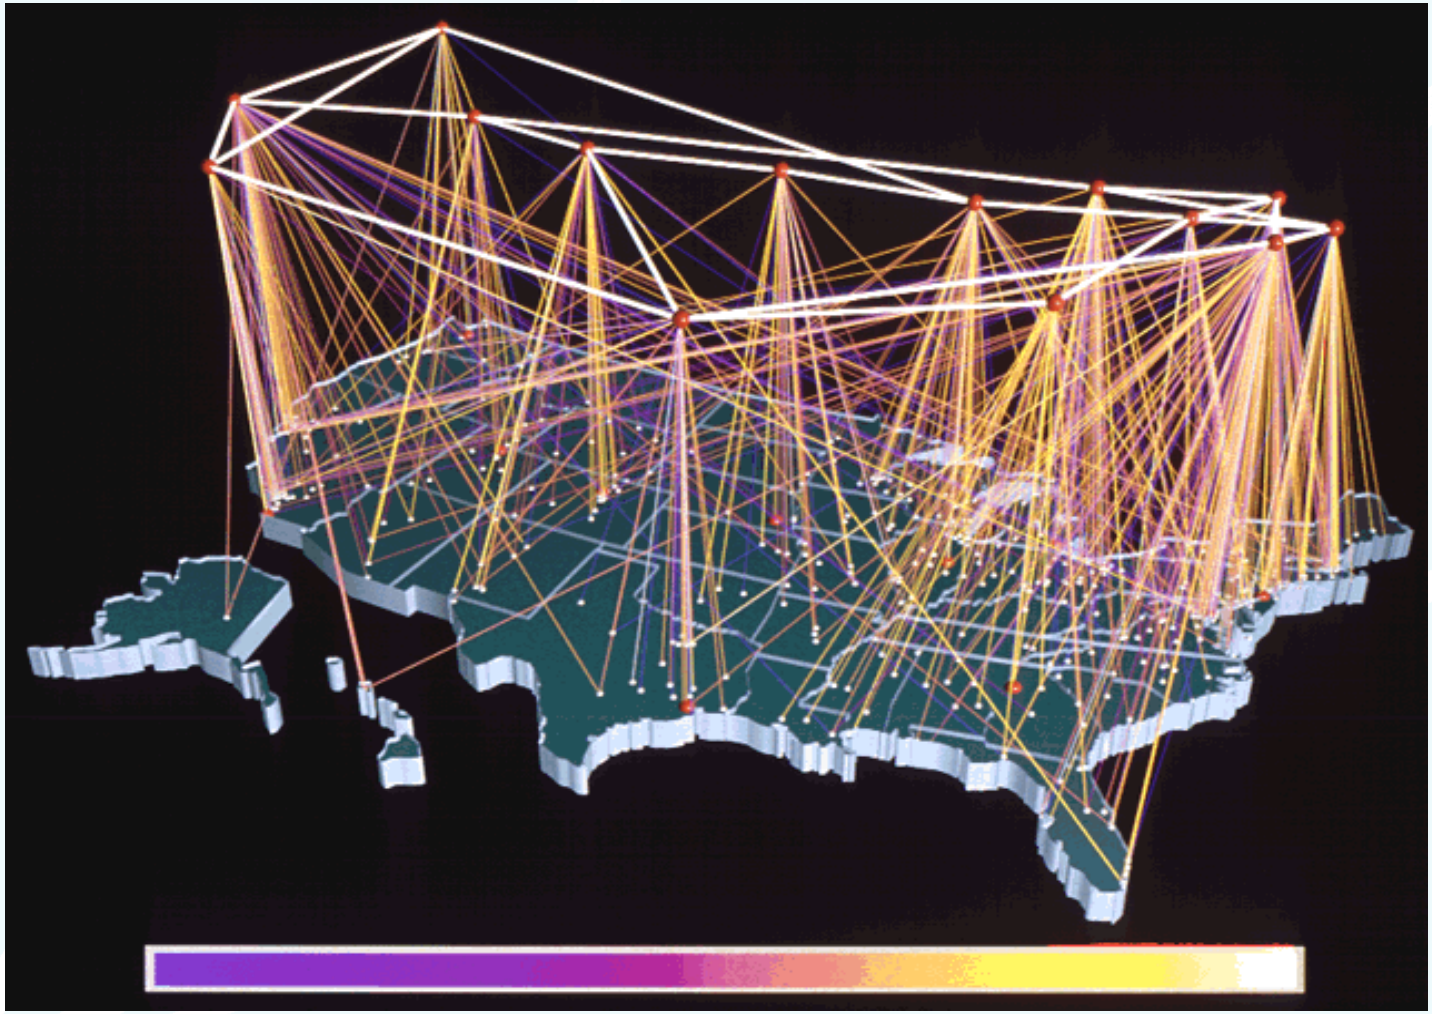

NSFNET
The Partnership That
Changed The World

KU-centric Introduction to GPN

- **Regional Network Service Provider.**
  - Connects the Kansas Research and Education Network (KU) to the Internet.
- **Established 1998 as MIDnet replacement**
  - National Computational Science Alliance
  - Internet2
- **Network Infrastructure Development**
  - Middleware
  - Application Support

NSFNET
The Partnership That Changed The World

Network Infrastructure Development

- **GPN Cluster Project**
  - NSF Middleware Initiative grant (2004?)
    - **Shibboleth enabled identification, certification, authentication for regional sharing of computational resources**
    - **Remote use of large cluster facilities (OSCER)**
    - **Pilot project for clustering small cluster computer installations**
  - Sun Microsystems equipment grant (2006)
    - **Four 8 node (computation) clusters (Sun x2100)**
      - **U Arkansas, U Kansas, U Missouri, U Nebraska**
    - **1.2 TB Data Storage Facility**
      - **U Missouri**

NSFNET
The Partnership That Changed The World

GPN Cluster Project

- **Purpose:**
  - Investigate and ameliorate the operational issues associated with using geographically distributed clusters for demanding applications.
- **Partners:**
  - Great Plains Network: Service provider/administration
  - University of Arkansas: Network software/hardware
  - University of Kansas: Topical applications/hardware
  - University of Missouri: Network software/hardware
  - University of Nebraska: Network hardware
  - University of Oklahoma: Network hardware/applications

NSFNET
The Partnership That Changed The World

Logical Structure of Future Effort

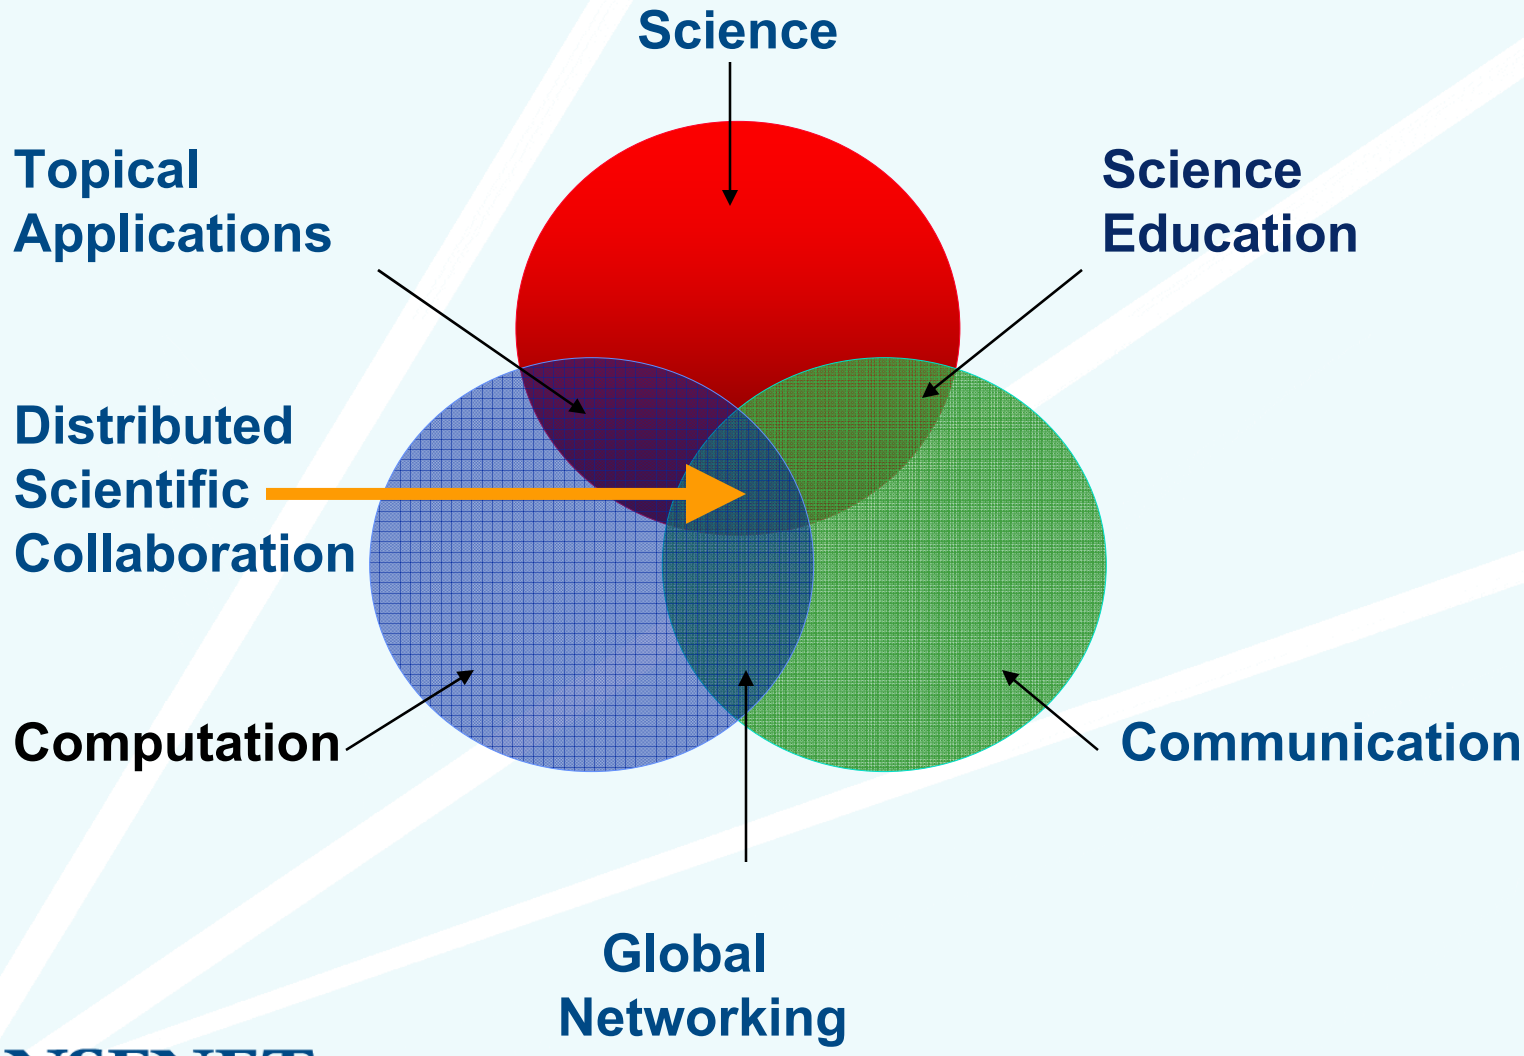

NSFNET
The Partnership That Changed The World

Recognized Colleagues

A short list of colleagues whose work is closely connected with projects described herein. Thank you.

- Jay Alameda, National Center for Supercomputing Applications
- Dick Alkire, University of Illinois
- Karen Camarda, Washburn University
- Kyle Camarda, University of Kansas
- John Connelly, University of Kentucky
- Donna Cox, National Center for Supercomputing Applications
- Larry Smarr, California Institute for Telecommunications and Information Technology
- Rick Stevens, Argonne National Laboratory

NSFNET
The Partnership That Changed The World

KU Chemical Engineering Team

- Kyle Camarda, Assistant Professor CPE
- Karen Camarda, Post-Doctoral Research Associate CPE
- Yuan “Eric” He, Ph.D. ChE
- John Eslick, M.S. & Ph.D. (Candidate) ChE
- Abhijith Halikhedkar, M.S. EECS
- Hong Chen, M.S. ChE & EECS
- Li Cheng, M.S. ChE & EECS
- J. P. Pakalapati, M.S. ChE
- Neela Shubhakar, M.S. ChE
- Sachin Siddhaye, M.S. ChE
- Ryan Gwaltney, Undergraduate ChE

NSFNET
The Partnership That
Changed The World

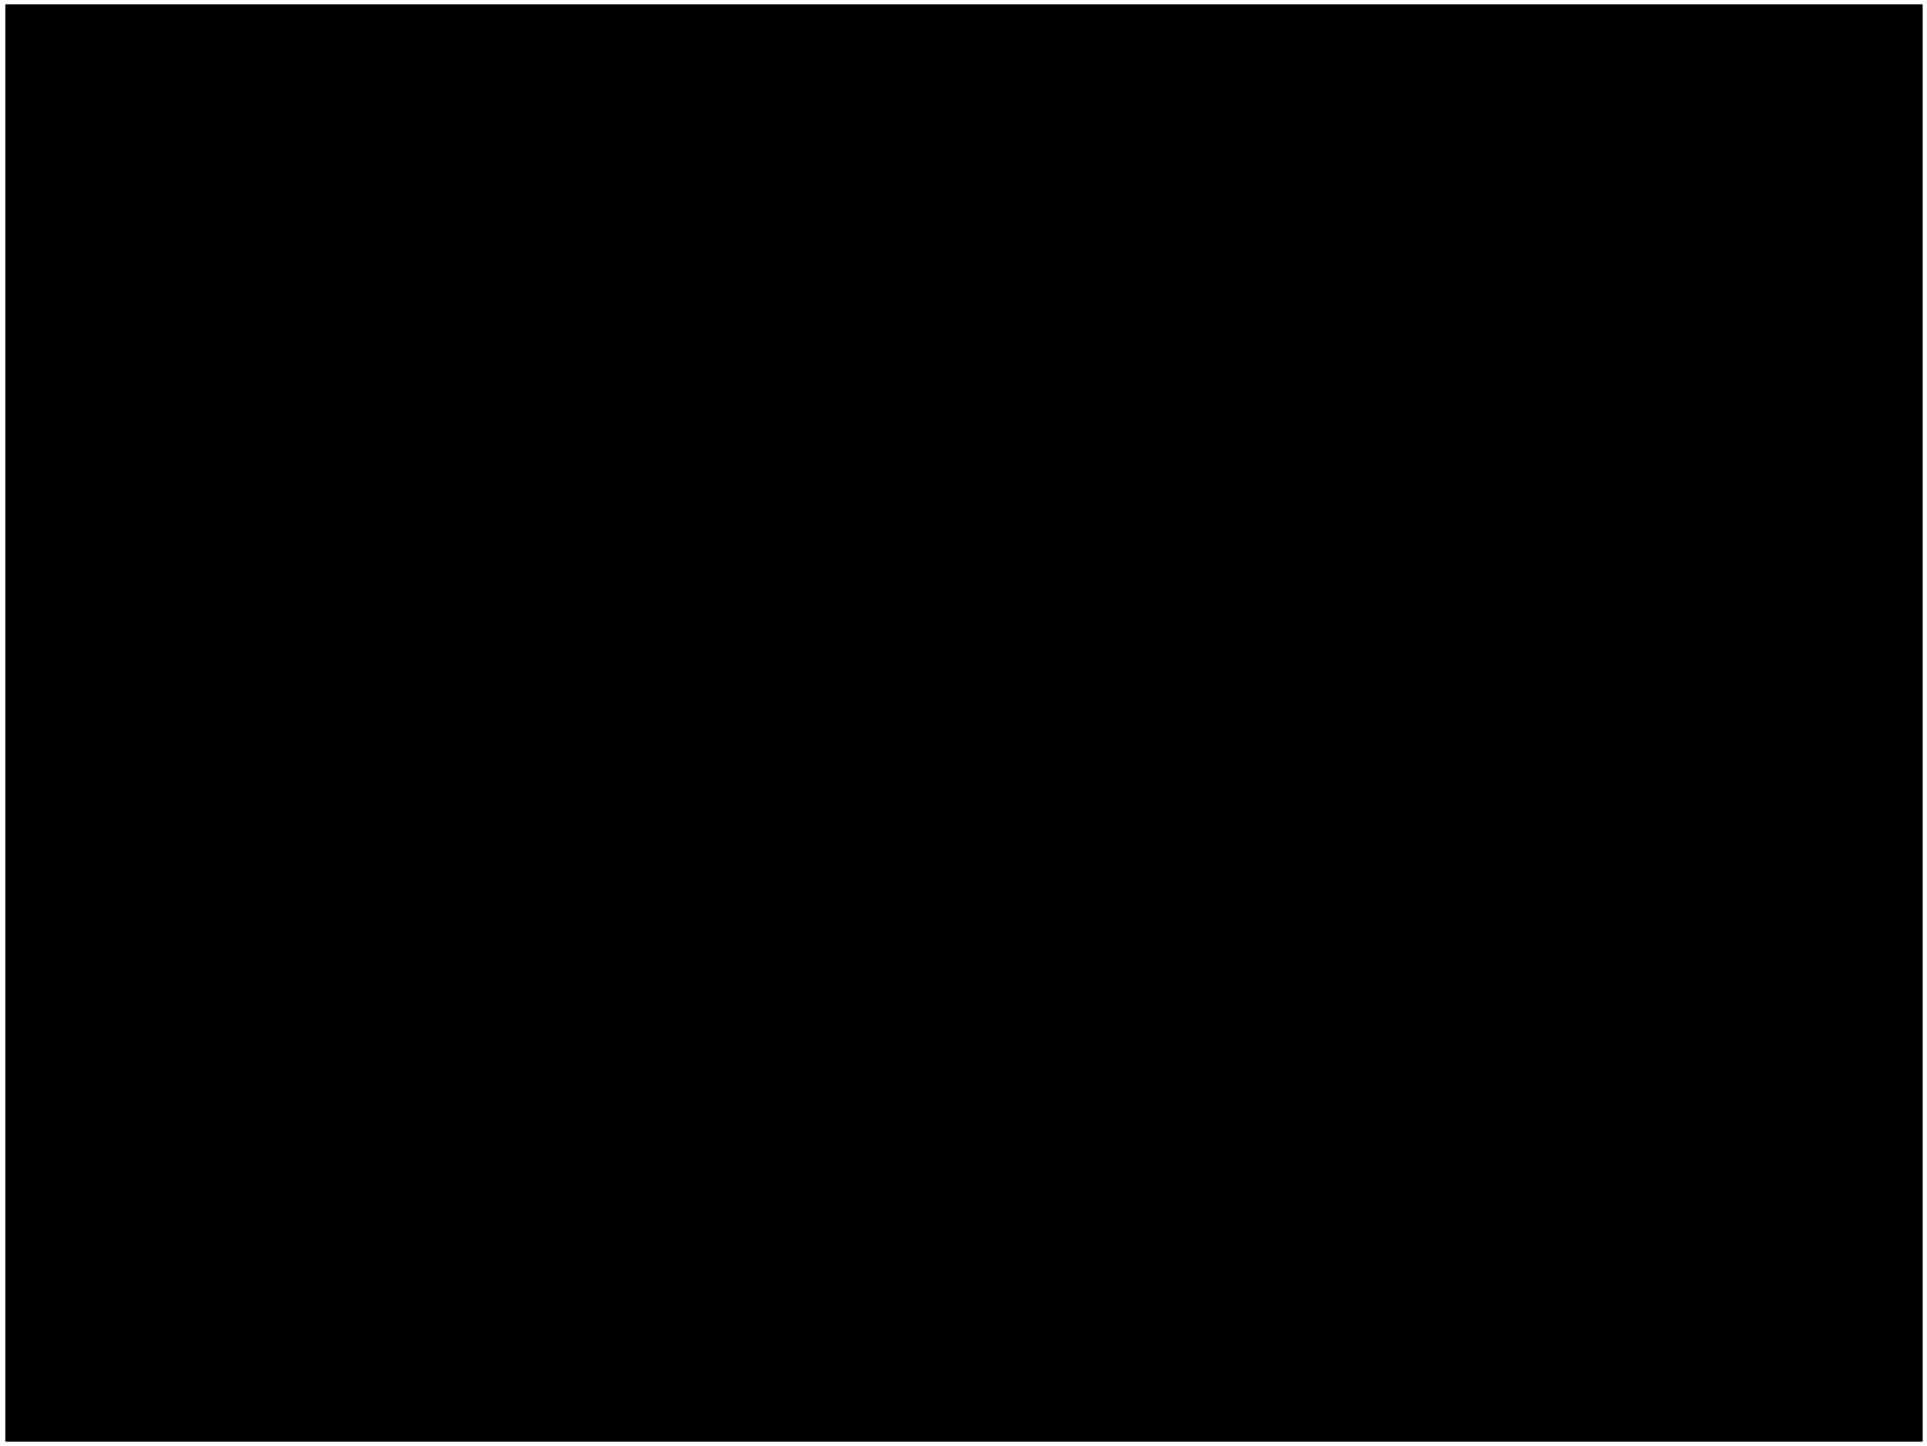

Science in 5 Dimensions

Chris Greer
National Coordination Office
Networking and Information Technology Research and Development

NSFNET
The Partnership That Changed The World

Email: greer@nitrd.gov

Growth of GenBank

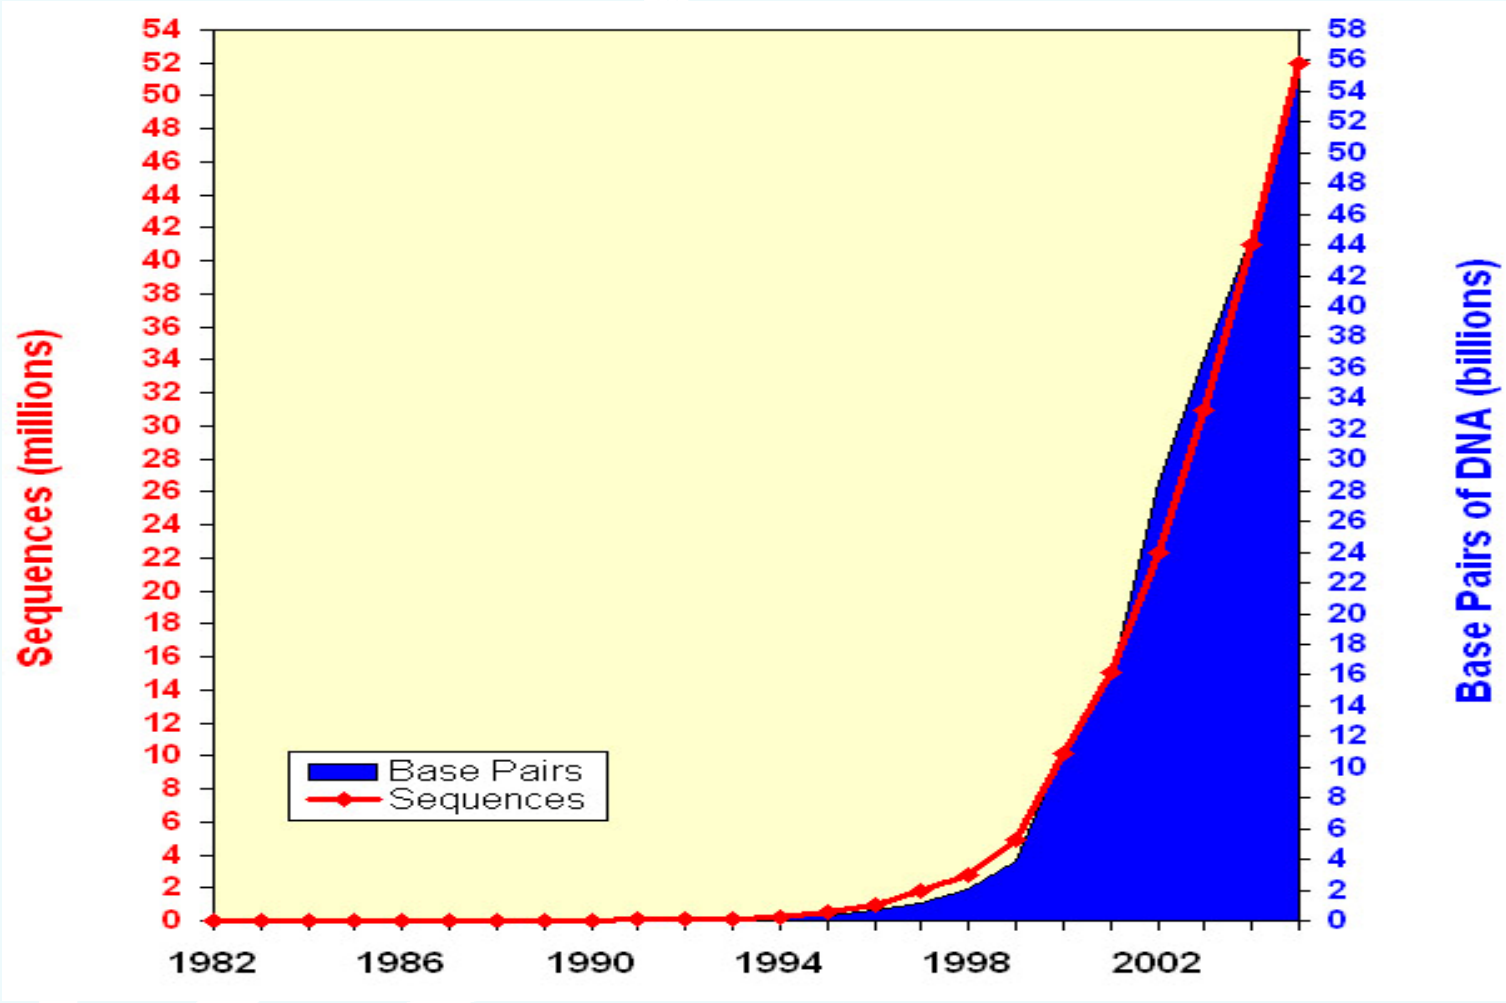

NSFNET
The Partnership That Changed The World

www.ncbi.nlm.nih.gov/Genbank/genbankstats.html

"In 2006, the amount of digital information created, captured, and replicated was 1,288 x 10^{18} bits (or 161 exabytes) … This is about 3 million times the information in all the books ever written"

The Expanding Digital Universe
IDC White Paper sponsored by EMC; March, 2007

NSFNET
The Partnership That Changed The World

Cyberinfrastructure

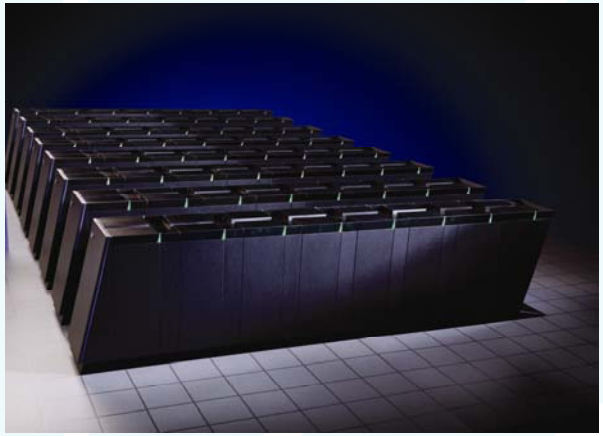

Computational capacity and capability

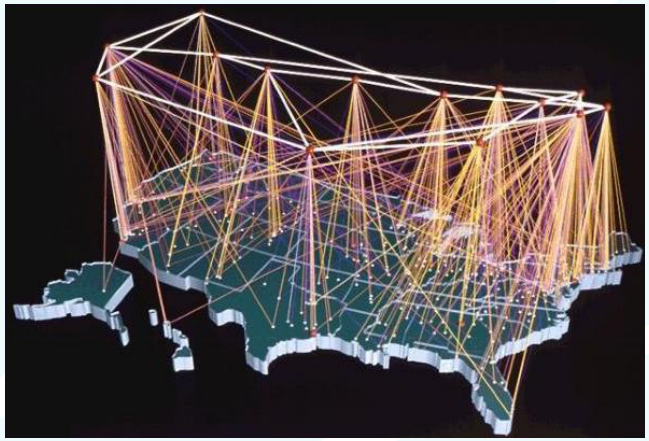

Connectivity for access and interaction

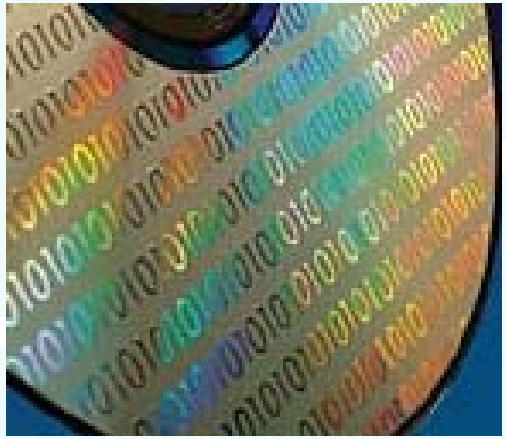

Information for innovation and discovery

NSFNET
The Partnership That Changed The World

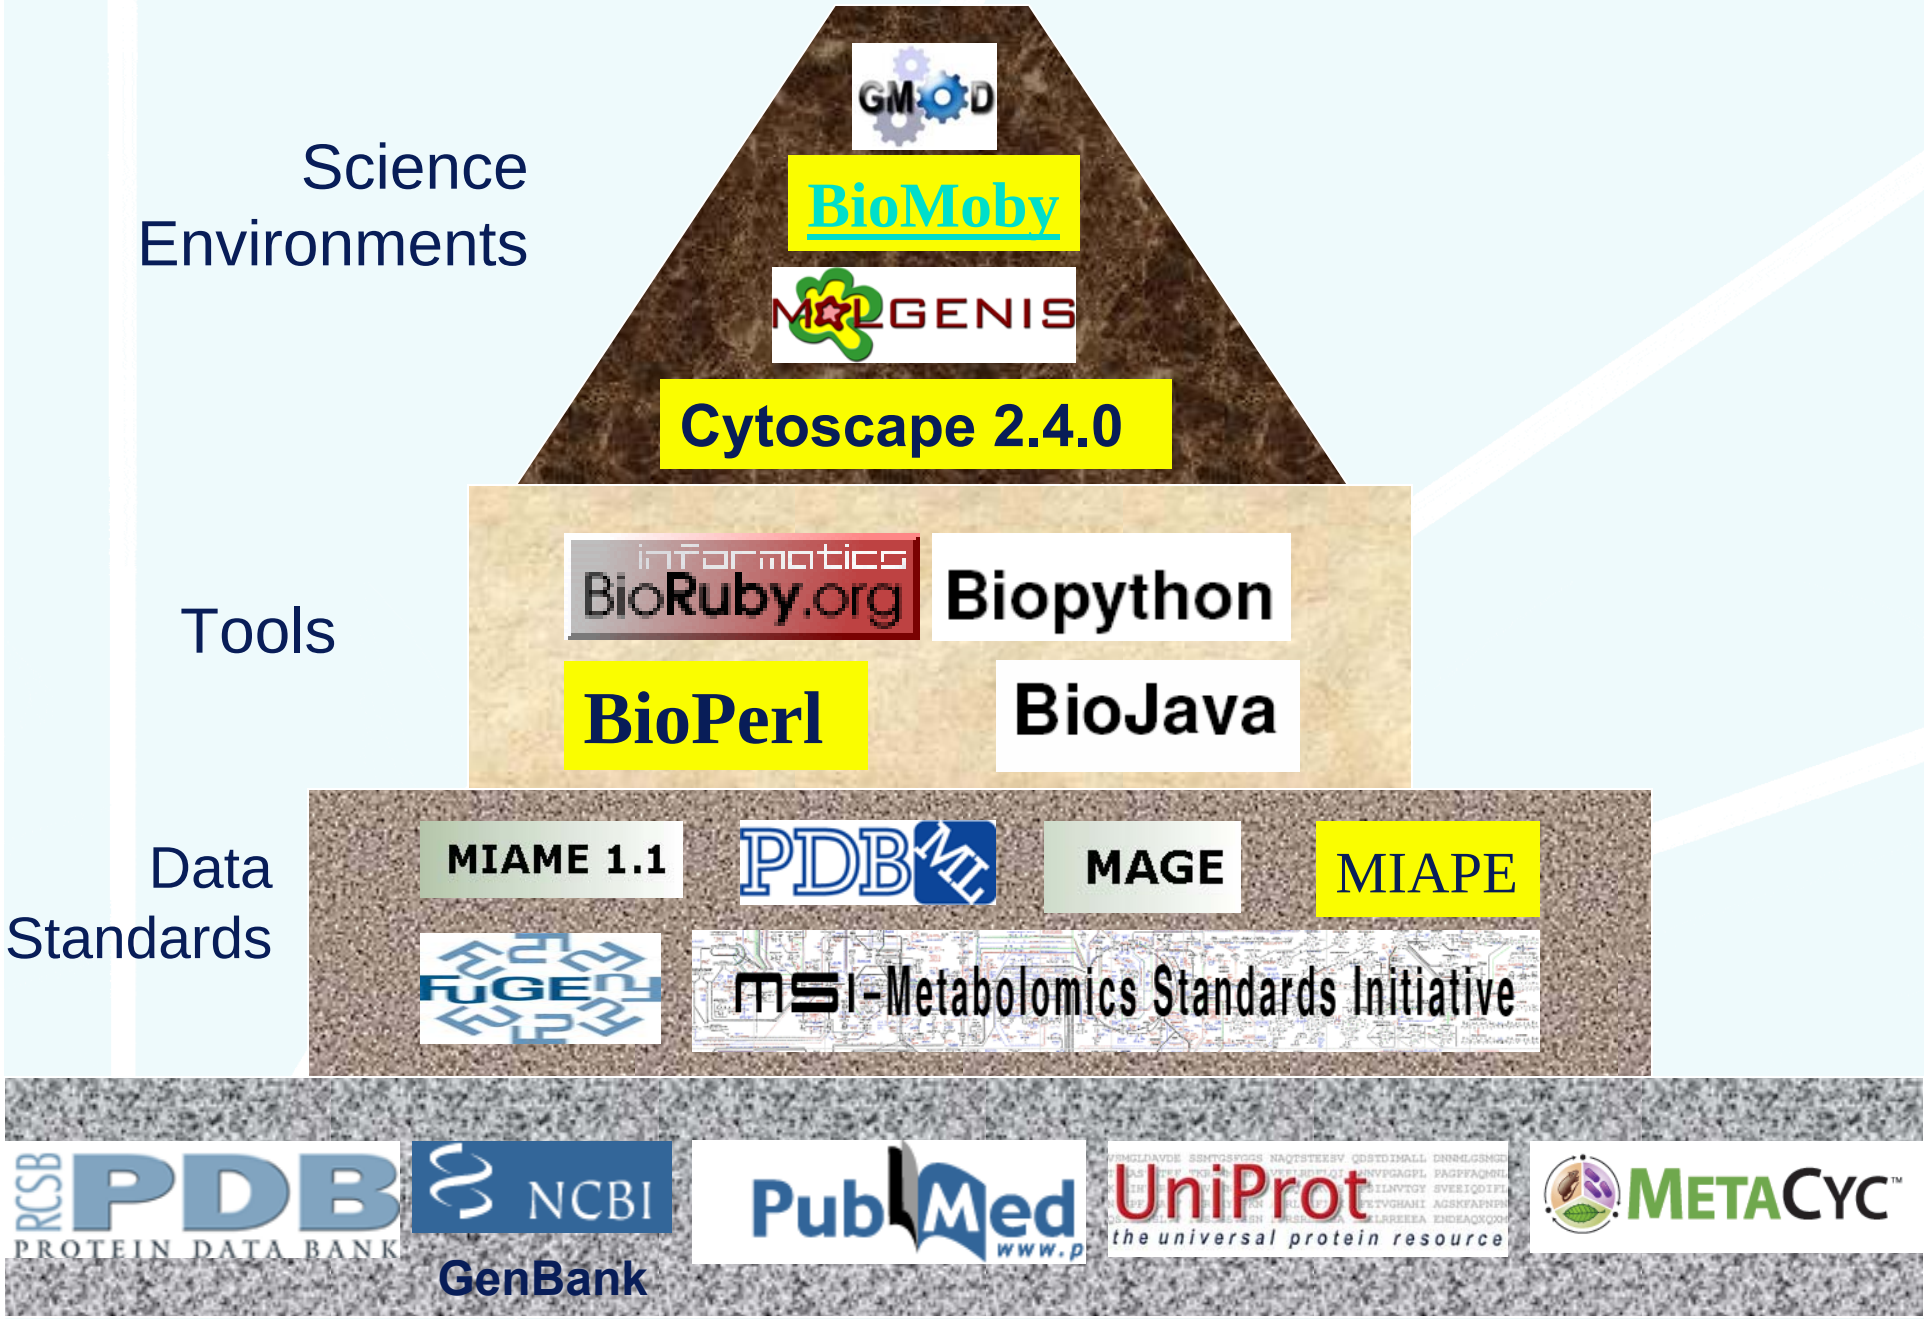
Science Environments
GMOD
BioMoby
MOLGENIS
Cytoscape 2.4.0
Tools
BioRuby.org informatics
Biopython
BioPerl
BioJava
Data Standards
MIAME 1.1
PDBML
MAGE
MIAPE
FuGE
MSI-Metabolomics Standards Initiative
RCSB PDB PROTEIN DATA BANK
NCBI
GenBank
PubMed www.p
UniProt the universal protein resource
METACYC
Digital Preservation/Access Organizations

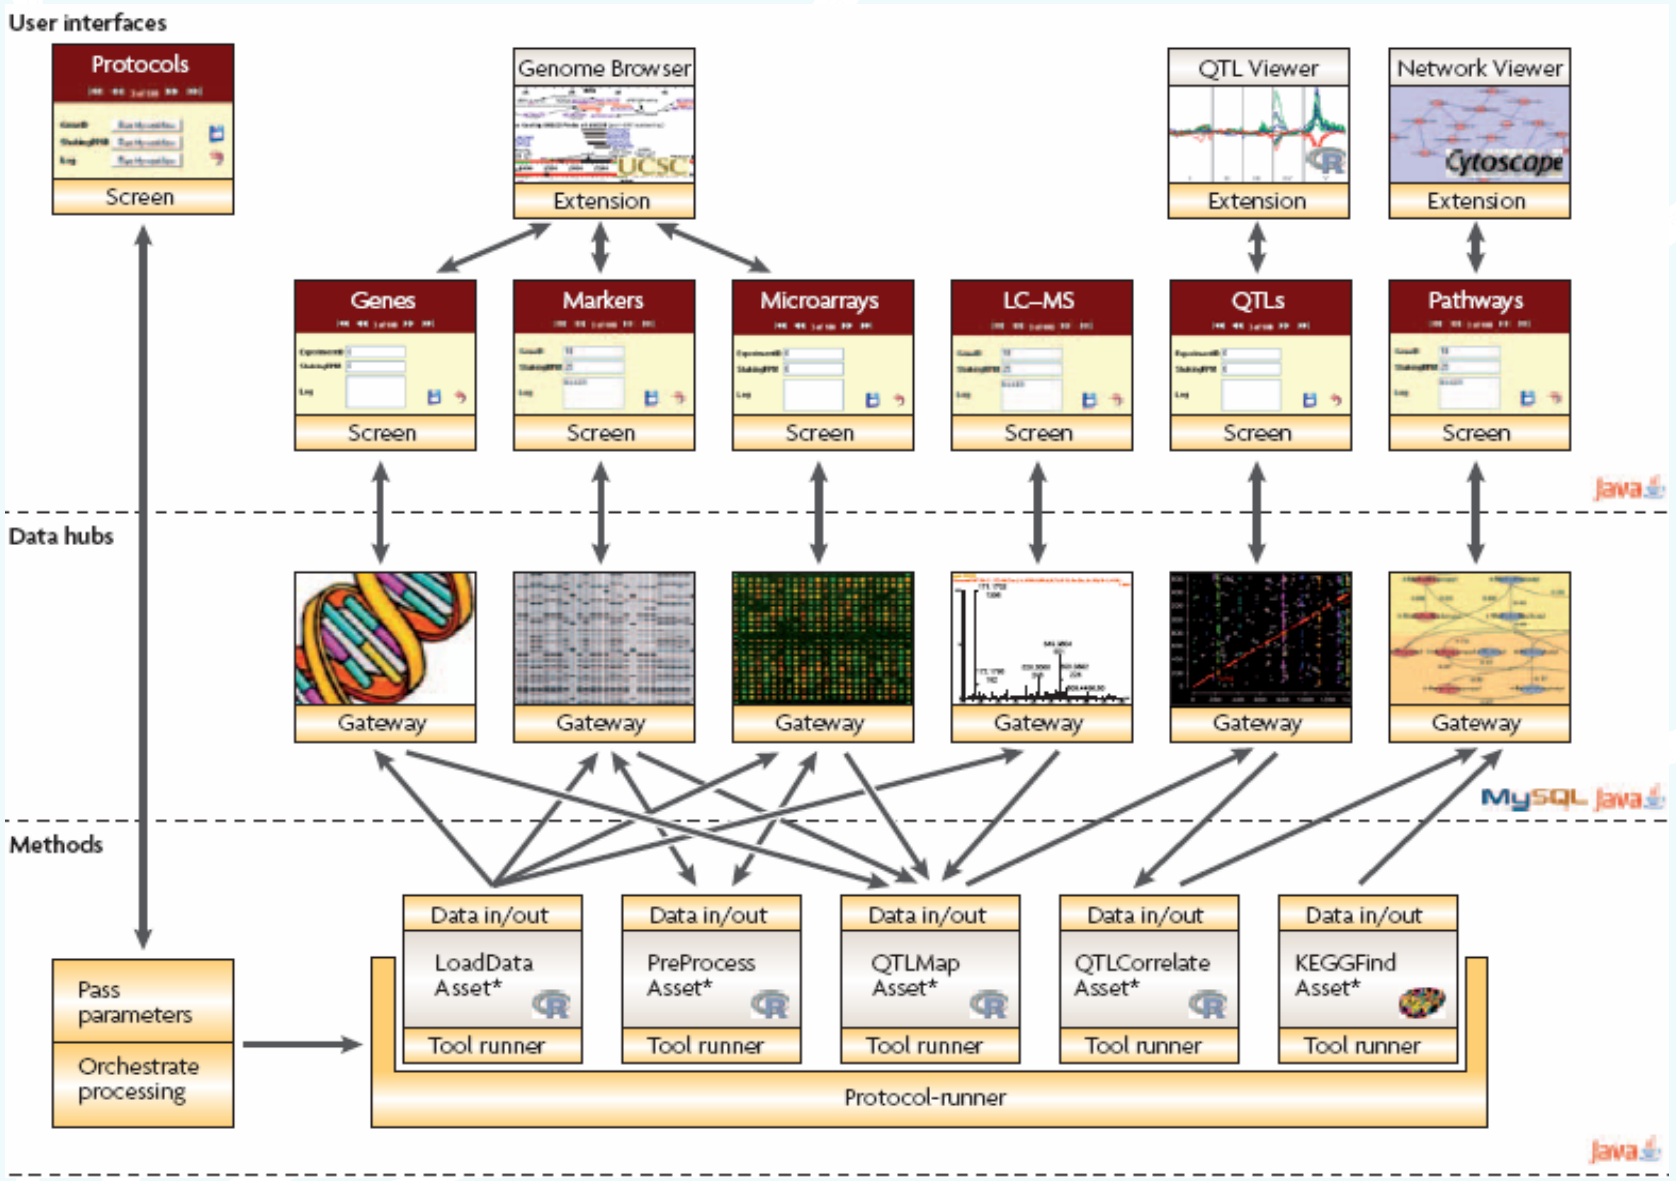

NSFNET
The Partnership That Changed The World

Source: Swertz, MA, Jansen, RC (2007) Nature Rev Gen 8:235 .

Connectivity → Synthesis

Human Accessible → Machine Accessible

Text → Context

NSFNET
The Partnership That Changed The World

The Impact on Research and Science

Nov 29, 2007, Mark Luker, EDUCAUSE

NSFNET
The Partnership That Changed The World

Shifting gears, setting the stage

- 95-97: Hand over the keys to industry
  - ISPS, NAPS, Routing Arbiter
- New High-Performance Connections program
  - New regional and national networks
  - New campus networks
  - Higher Education “is connected”
- Harmonize federal research networks
- Promote “Next Generation Internet”

NSFNET
The Partnership That Changed The World

Revolutionary impact of NSFNET on our research workforce

- 24x7 access to shared content
- Interactive distance learning
- Asynchronous distance (and on-site) learning
- Active problem solving in teams
- The open courseware movement
- The globalization of education
- Direct involvement of students in research
- Podcasts, mobile learning
- STEM education initiatives
- Virtual organizations

NSFNET
The Partnership That Changed The World

Preparing our future researchers

- Problem-solving
- Flexibility and ability to work in teams
- Proven skills in real-life situations
- Develop expert thinking and strategies

NSFNET
The Partnership That Changed The World

Virtual observatory

BRADFORD ROBOTIC TELESCOPE

Sunday 25 March

MAIN PAGE | SYSTEM STATUS | USE THE TELESCOPE | WEATHER REPORTS | WEBCAMS | FORUMS | PROJECT NEWS | CONTACT US | TELESCOPE STATISTICS | IMAGE GALLERY | PHOTO GALLERIES | EDUCATIONAL MATERIAL | REGISTER FOR AN ACCOUNT | INFORMATION | FUNDERS

Login

Username

Password

Login

Forgotten your details? Click here

Click here to register for a new account

Welcome to the Bradford Robotic Telescope

The Bradford Robotic Telescope is unique.
If you want to wonder at the grandeur and beauty of your star sign it will take a colour image for you covering all the stars in your constellation.

It will show you the majesty of the sky seen by our grandparents before the age of light pollution. Look at our image gallery!

Image from the Gallery

MESSIER 63 | Avg. rating 9.5

Educational Material

Teachers

General Interest

Remote instruments

Accelerometer

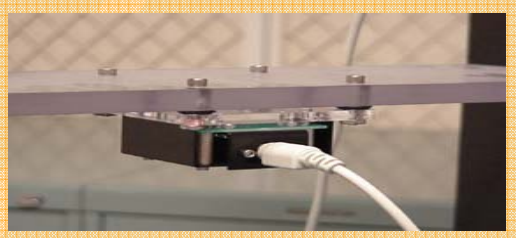

Data Acquisition / Control Devices

Lab Server

NSFNET
The Partnership That Changed The World

—Lerman, 2006

Haptics

- Users feel force, pressures and temperature while interacting with virtual environment

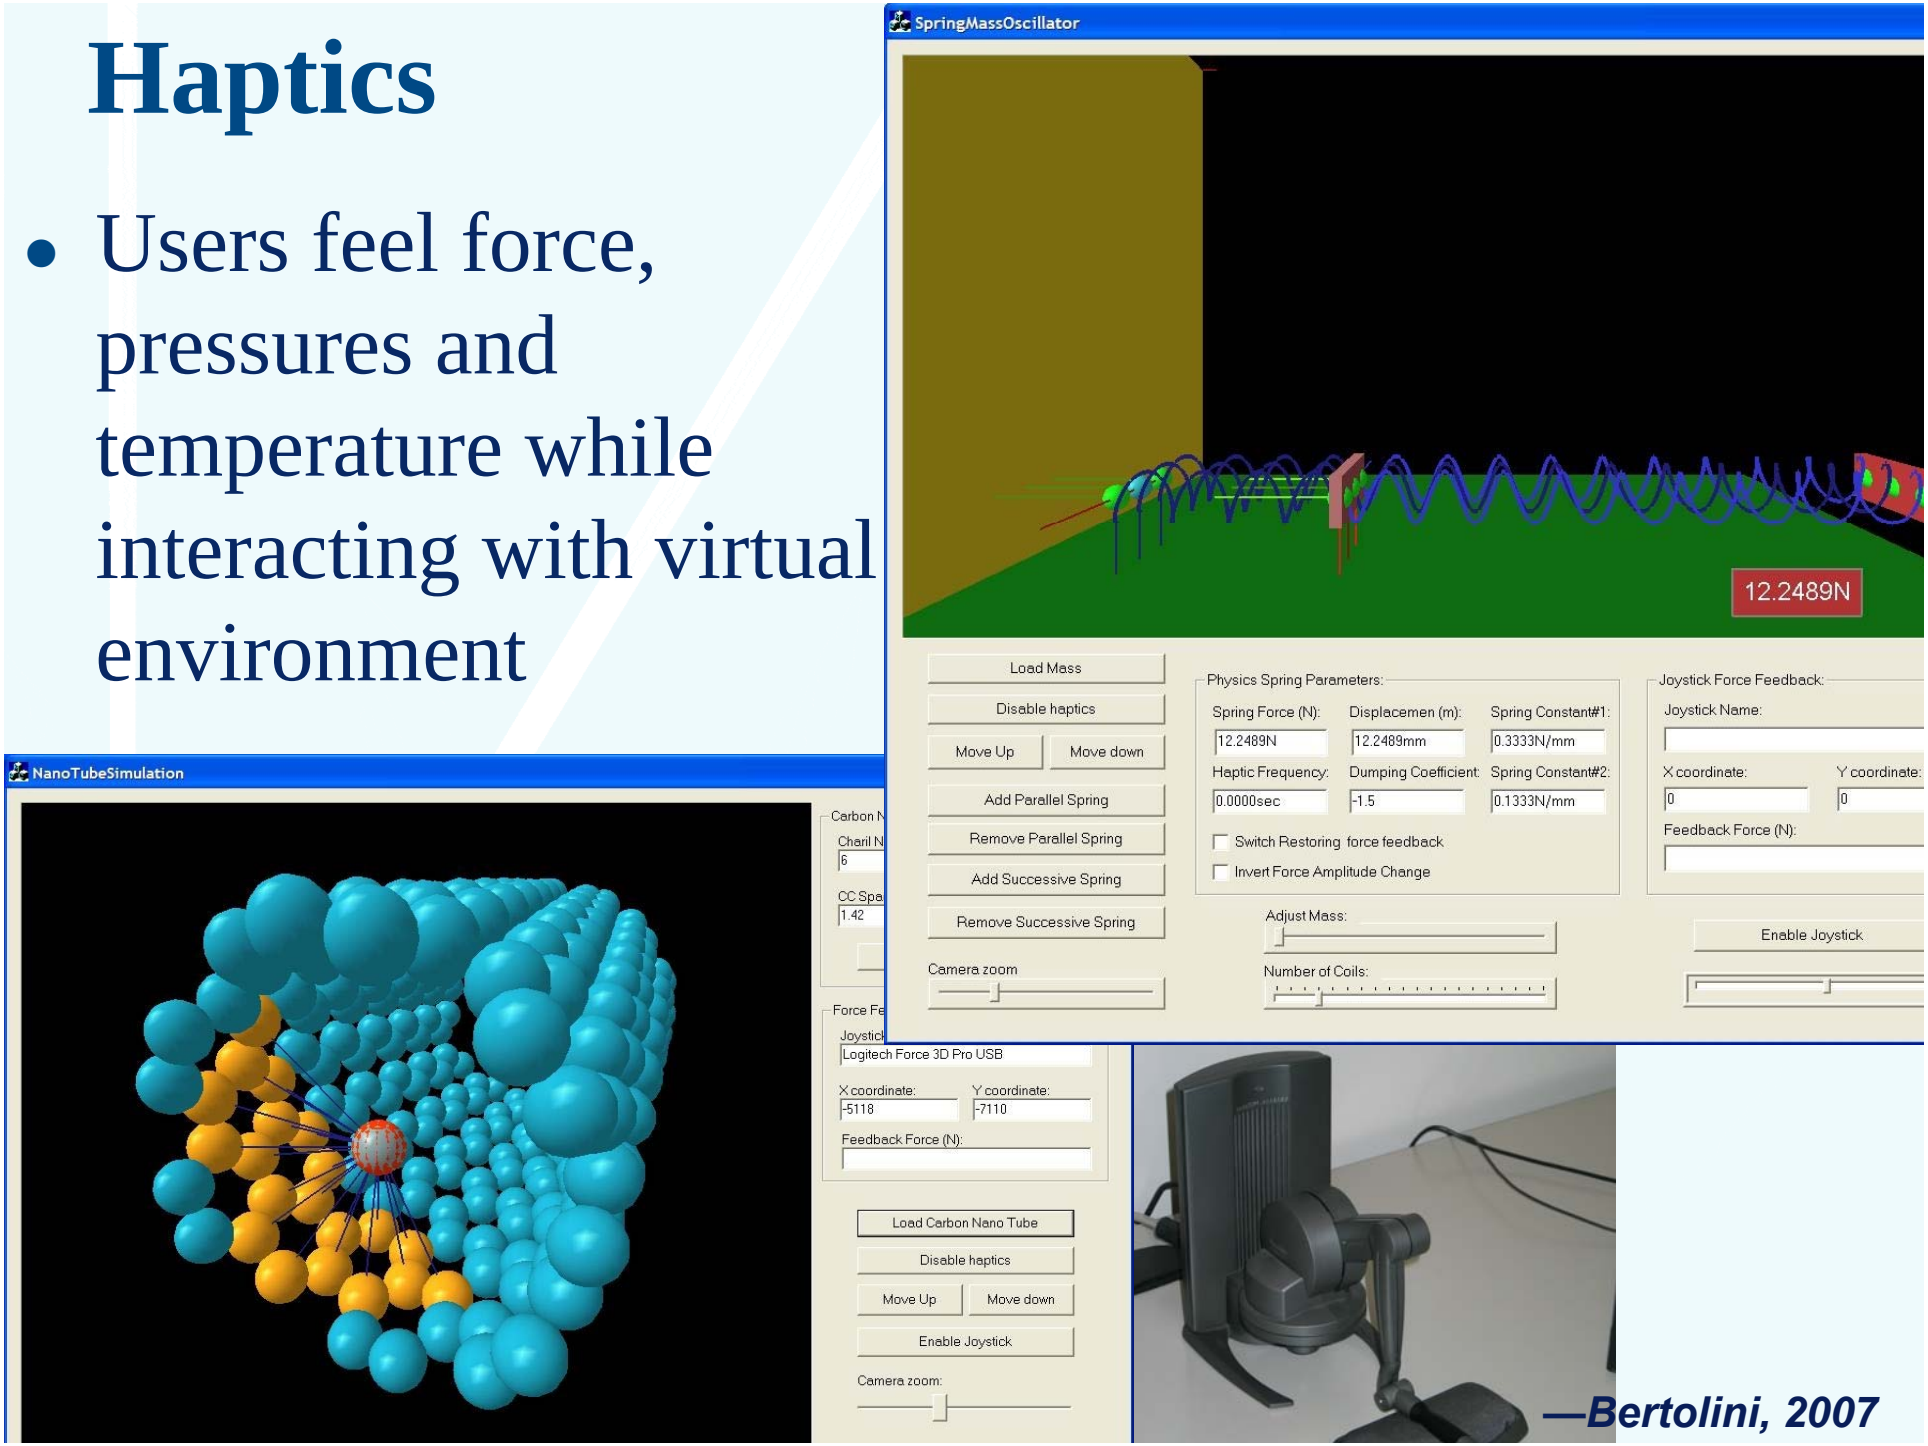

—Bertolini, 2007

Learning-to-be

- National Ecological Observatory Network
- Remote & collaborative environments
- Widely distributed sensors
- Real-time data collection and analysis

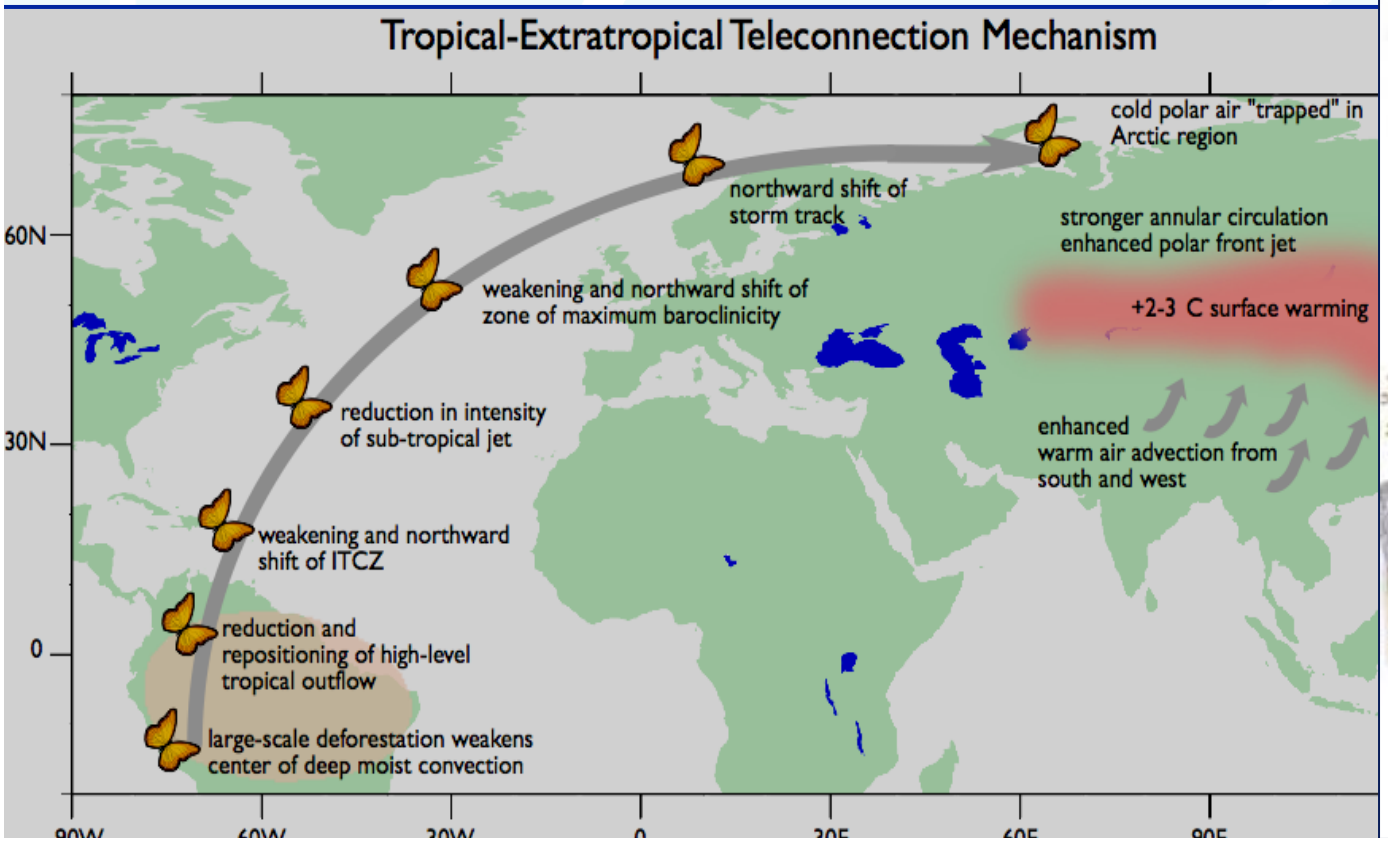

--NSF, 2006

Education for research transformed

- Collaborative problem solving
- Cross-disciplinary engagement
- Integration of learning and research

NSFNET
The Partnership That Changed The World

www.educause.edu

NSFNET

The Partnership That
Changed The World

NSFNET: The Impact on Research and Science. Panelist Comments.

11/29/07 David Nelson

NSFNET
The Partnership That
Changed The World

Impacts More-or-Less Proven
(Roughly chronological)

- Remote supercomputing and collaborative computational science - starting 1974
- Remote access to experimental data
- Critical mass for TCP/IP-based Internet (especially NSFNET)
- Remote operation of equipment and facilities (but reliability issues)
- Distributed running of codes (only partial success because of latency)

NSFNET
The Partnership That Changed The World

Present or Ongoing

- Distributed collaboration
  - Email, video, shared data, etc.
  - Video window vs. data window
- Remote access to research results
  - Informal access works (preprints, research web sites, reports, etc.)
  - Access to published papers far from solved problem - serious IP issues

NSFNET
The Partnership That Changed The World

Future and Under Development

- Very wide, dynamically adjustable, bandwidth with high reliability
  - Third generation Internet at 10-100 Mb/s end-to-end: “Effortless video to the home and desk”
  - A new Turing test: “Is it a person or an image?
- Trusted network with adequate security and privacy

NSFNET
The Partnership That Changed The World

Networks for HEP and Data Intensive Science: and the Digital Divide

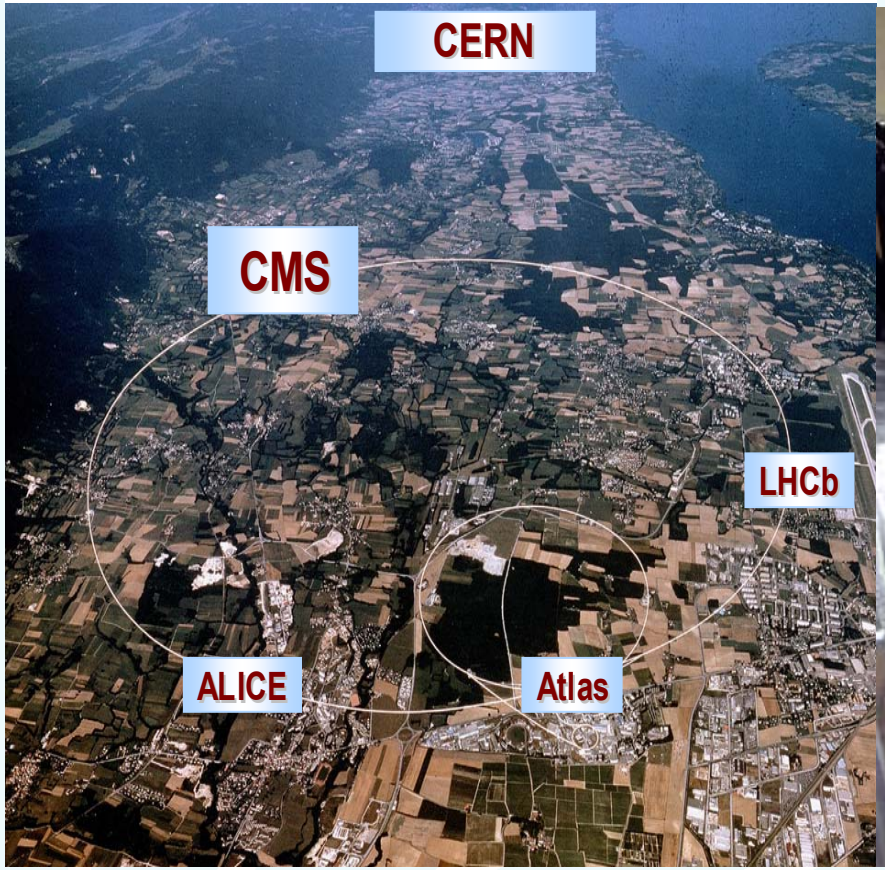

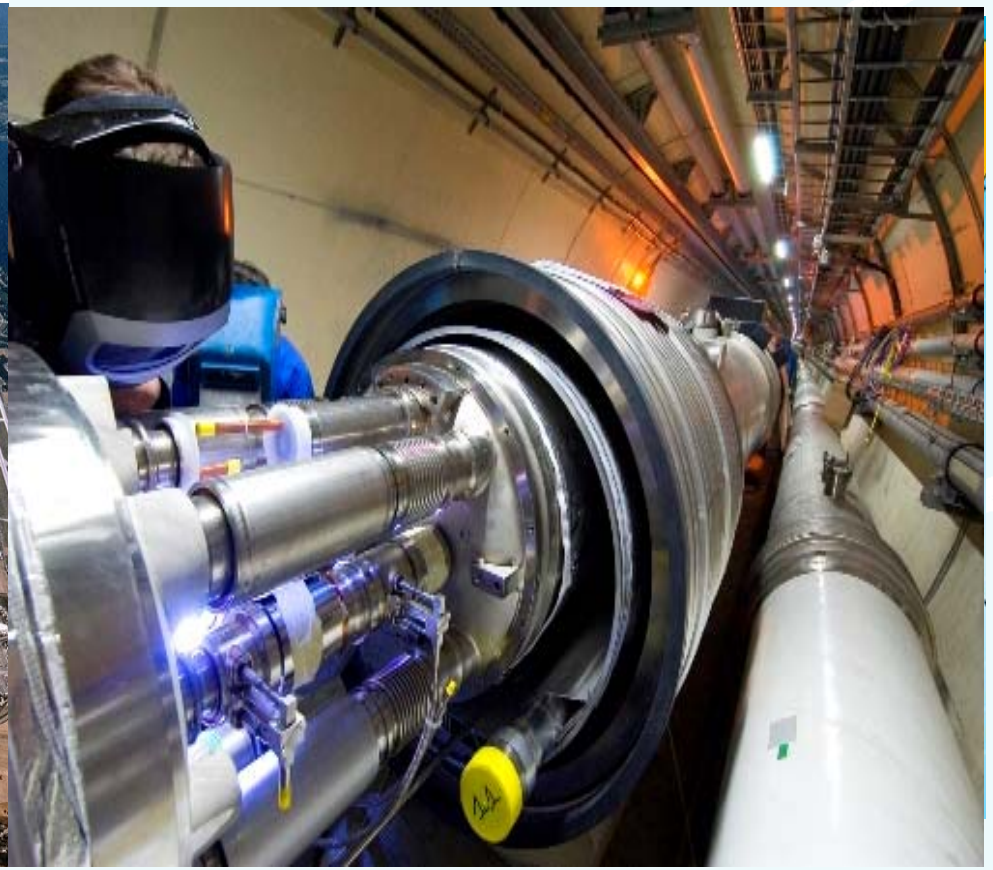

Harvey B Newman, Caltech

NSFNet Event

November 29, 2007

NSFNET
The Partnership That Changed The World

ICFA SCIC

Large Hadron Collider CERN, Geneva: 2008 Start

- pp √s = 14 TeV L = 10^34 cm^-2 s^-1
- 27 km Tunnel in Switzerland & France

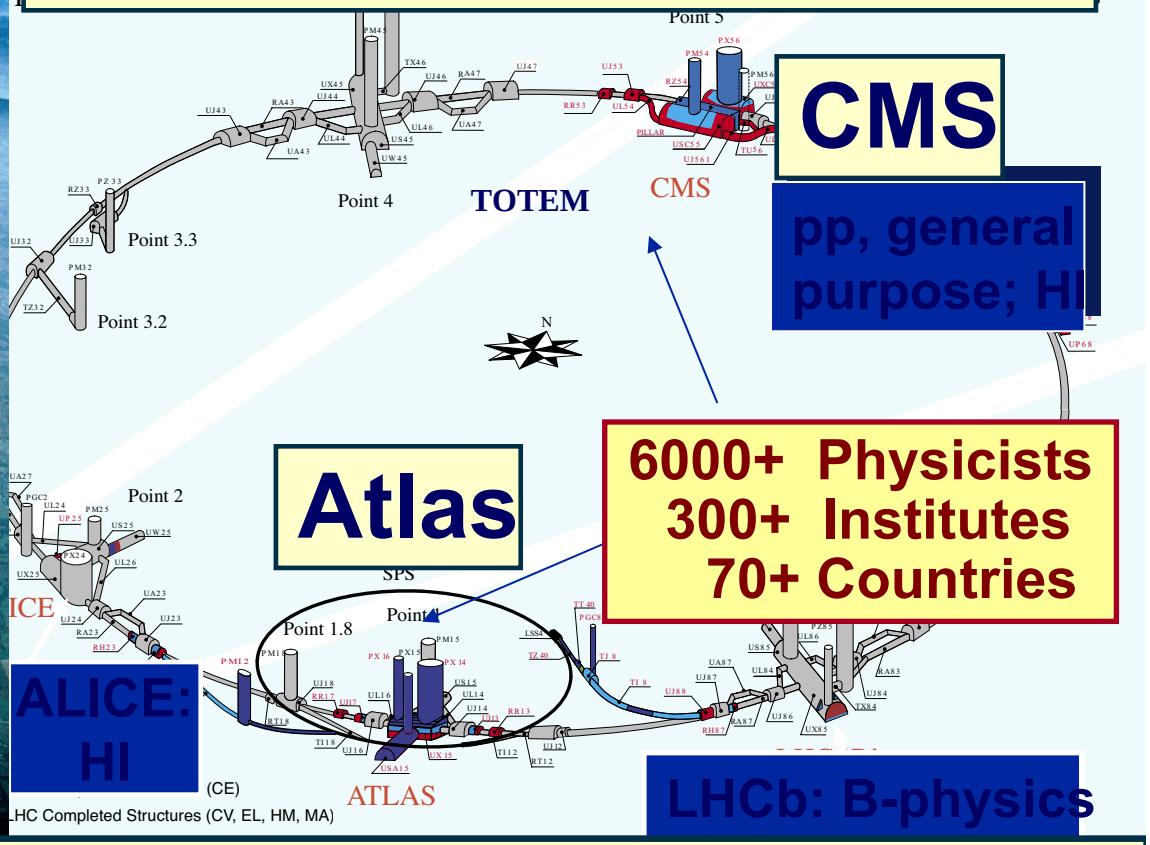

CMS

pp, general purpose; HI

Atlas

6000+ Physicists
300+ Institutes
70+ Countries

ALICE: HI

LHCb: B-physics

Challenges: Analyze petabytes of complex data cooperatively

Harness global computing, data & network resources

Parton-parton Luminosities

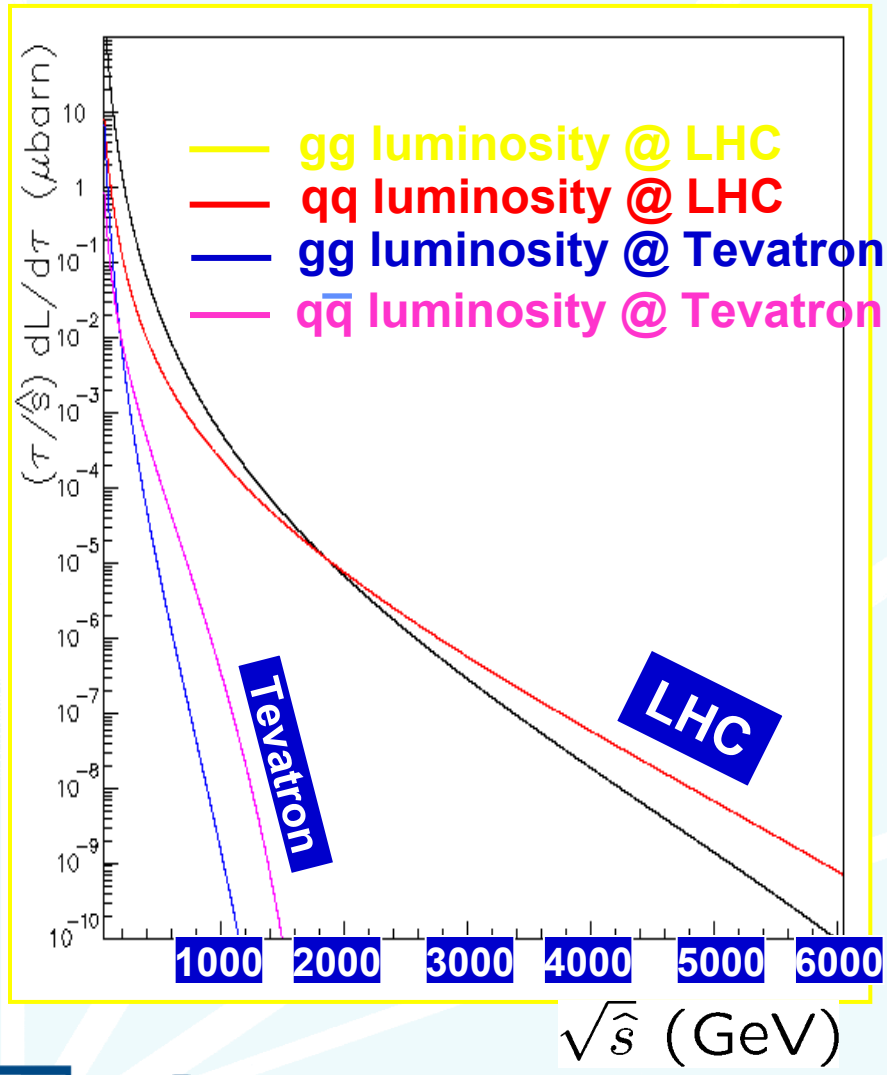

- The LHC is a Discovery Machine
- The first accelerator to probe deep into the Multi-TeV scale
- Many reasons to expect new TeV-scale physics

Parton-Parton CM Energy

NSFNET
The Partnership That Changed The World

Making Black Holes: $M_{BH} > M_*$(Gravity) (Saved by Hawking Radiation)

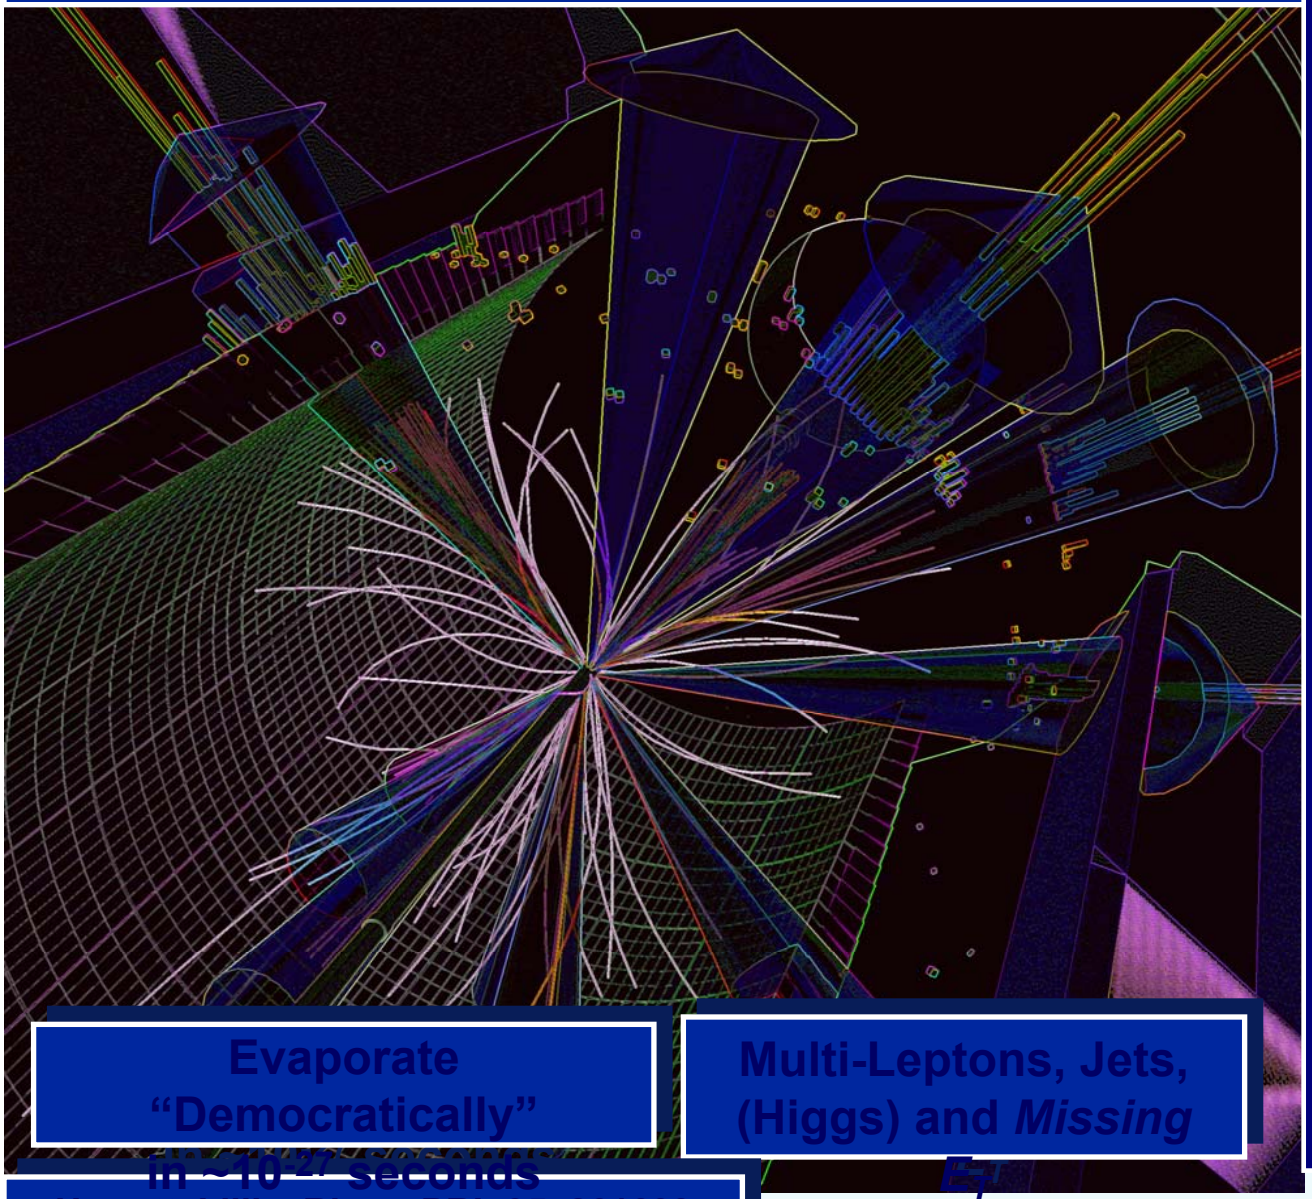

Evaporate "Democratically"

Multi-Leptons, Jets, (Higgs) and *Missing* E_T

In ~10^{-27} seconds

Hewett, Lillie, Rizzo, PRL 95, 261603 (2005)

Black Holes may be copiously produced:

BH Events per year at Low Lumi.

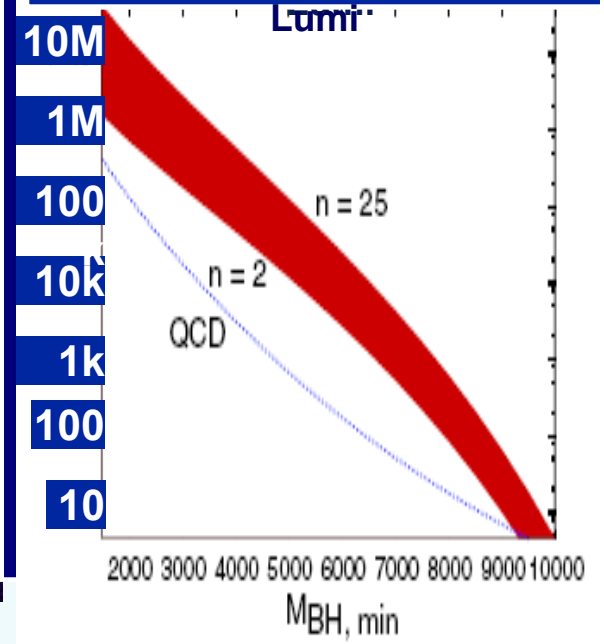

MONARC ➡ CMS & ATLAS Models, DISUN

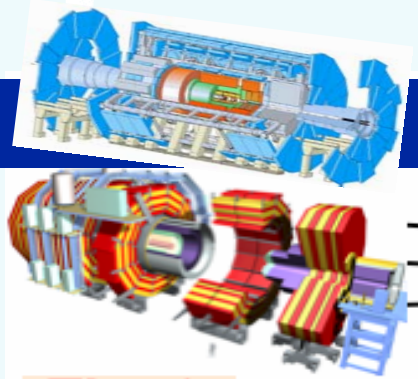

CERN/Outside Ratio ~1:4 T0/(ΣT1)/(ΣT2)
~1:2:2
~40% of Resources in Tier2s

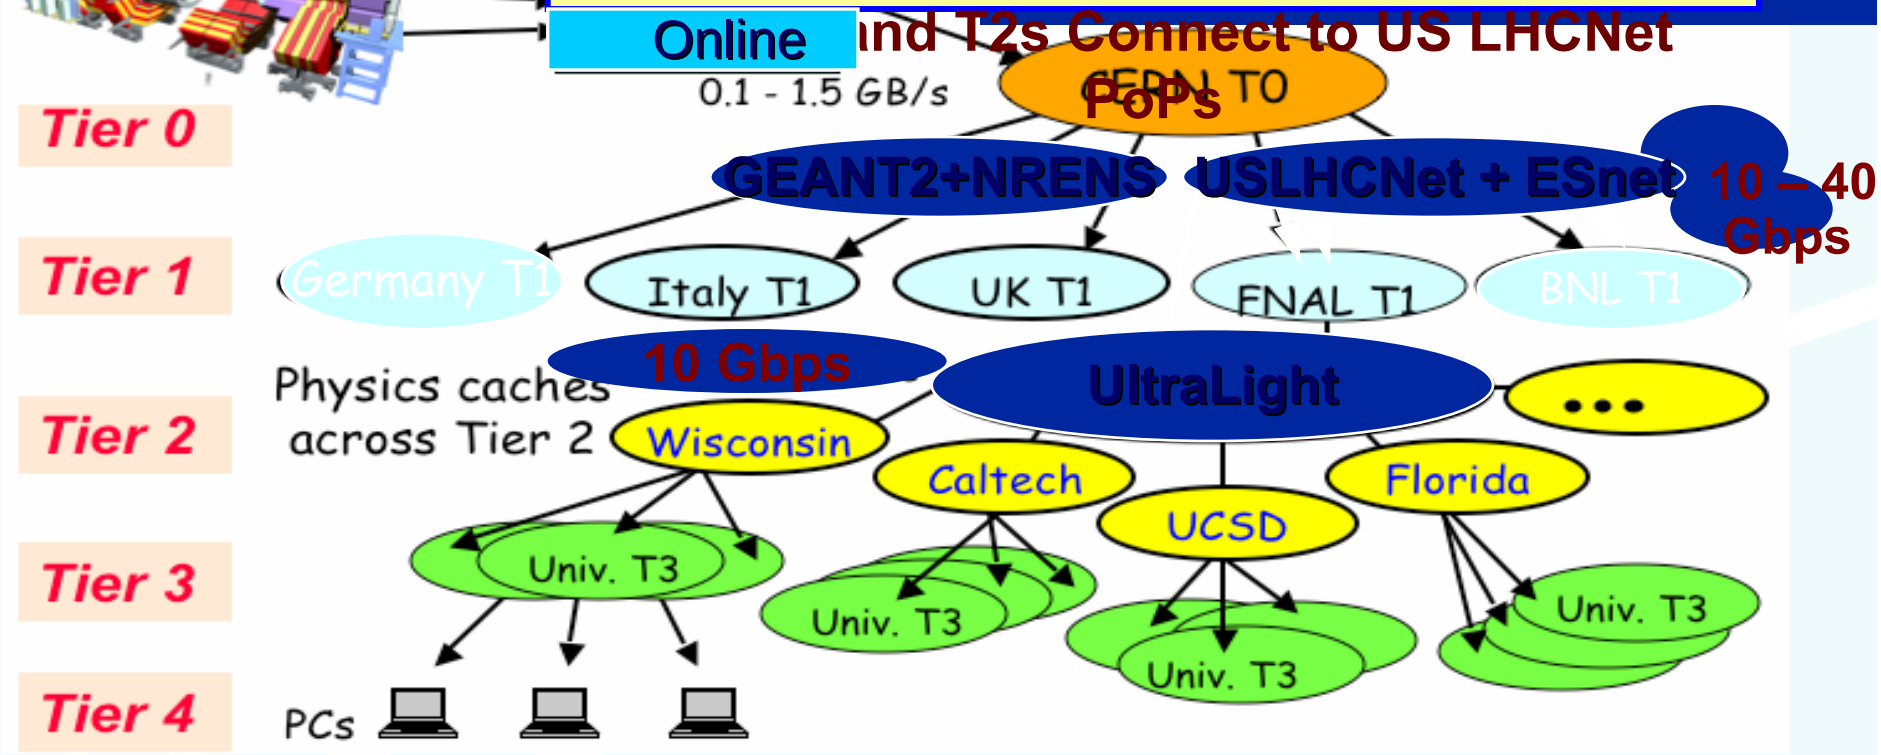

Outside/CERN Ratio Larger; Expanded Role of Tier1s & *Tier2s: Greater Reliance on Networks*

for Major Links (in Gbps)

Year	Production	Experimental	Remarks
2001	0.155	0.622-2.5	SONET/SDH
2002	0.622	2.5	SONET/SDH DWDM; GigE Integ.
2003	2.5	10-20	DWDM; 1 + 10 GigE Integration
2005	10-20	2-10 X 10	λ Switch; λ Provisioning
2007	3-4 X 10	~10 X 10; 100 Gbps	1st Gen. λ Grids
2009	~6 X 10 or 100	~20 X 10 or ~2 X 100	100 Gbps λ Switching
2011	~20 X 10 or 2 X 100	~10 X 100	2nd Gen λ Grids Terabit Networks
2013	~Terabit	~MultiTbps	~Fill One Fiber

ontinuing Trend: ~400-1000 Times Bandwidth Growth Per Deca
Paralleled by ESnet Roadmap for Data Intensive Sciences

Links
US-CERN Example 1985 - 2010

- *Rate of Progress >> Moore's Law*
  - *9.6 kbps Analog* *1985*
  - *64-256 kbps Digital* *1989 - 1994* *[X 7 – 27]*
  - *1.5 Mbps Shared [All TCP T1]* *1990-3; IBM* *[X 160]*
  - *2 -4 Mbps* *[Dedicated E1s]* *1996-1998* *[X 200-400]*
  - *12-20 Mbps* *1999-2000* *[X 1.2k-2k]*
  - *155-310 Mbps* *2001-2* *[X 16k – 32k]*
  - *622 Mbps* *2002-3* *[X 65k]*

20-80Gbps US-CERN, ESnet MANs

US-LHCNet:
Wavelength Quadrangle
NY-CHI-GVA-AMS
2007-10: 30, 40, 60, 80 Gbps

ESnet4
SDN Core:
30-50G

ESnet IP Core

SEA
SNV
SDG
DEN
ELP
ALB
ATL
CHI
FNAL
BNL
NYC
DC
Europe
GEANT2
SURFNet
IN2P3

10Gb/s
10Gb/s
30Gb/s
2 x 10Gb/s

US-LHCNet
Data Network
(3 to 8 x 10 Gbps
US-CERN)

ESNet MANs to FNAL & BNL; Dark Fiber to FNAL; Peering With GEANT

NSFNET
The Partnership That Changed The World

ESnet4: 40-50 Gbps by 2009-2010, 500 Gbps by

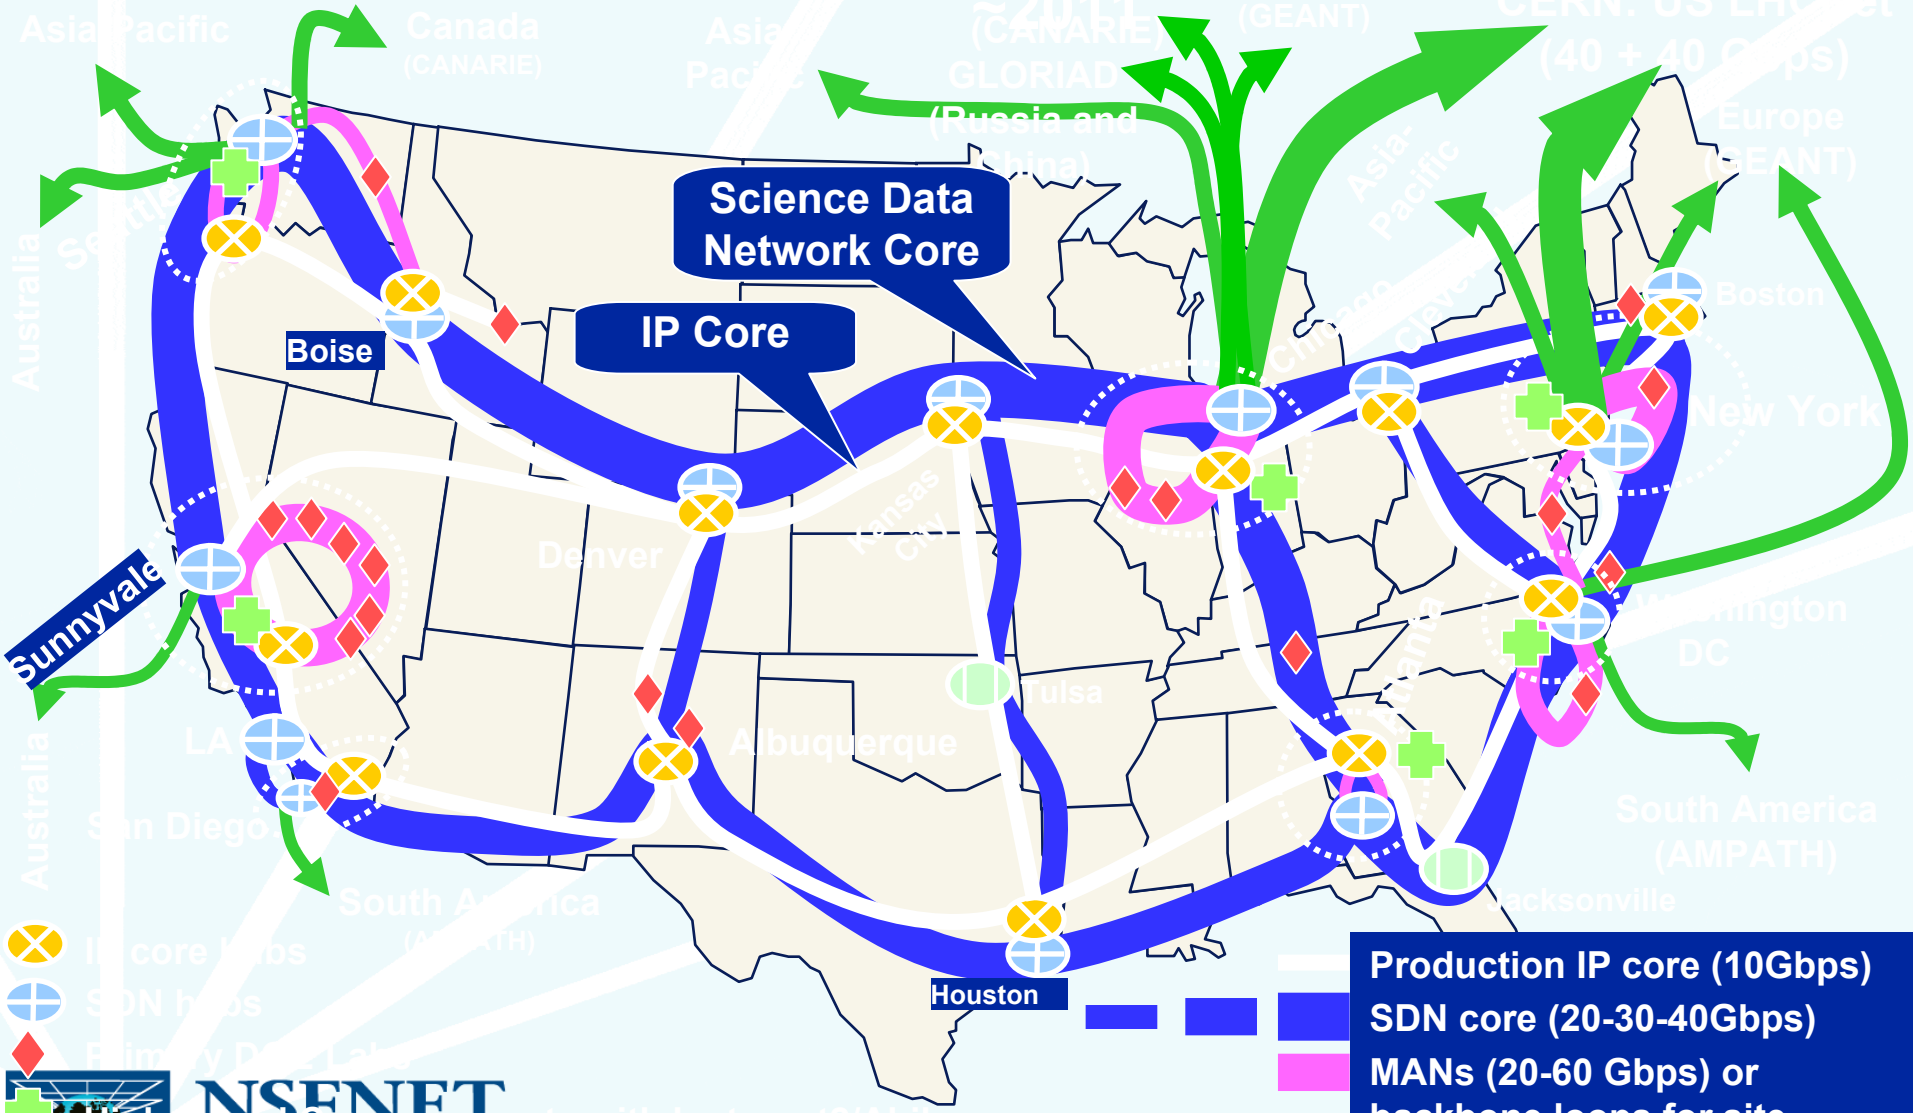

NSFNET
The Partnership That Changed The World

2nd Revolution: Networks for Research & Education and Data Intensive Science

- Current generation of 10 Gbps R&E network backbones and major
  Int'l links arrived in 2001-6 in US, *Europe*, Japan, Korea; Now *China*
  - Bandwidth Growth: from 4 to 2500 Times in 5 Yrs; >> Moore's Law
- Rapid Spread of "Dark Fiber" and DWDM; emergence of Continental, Nat'l, State & Metro "Hybrid" Networks in Many Nations
  - Cost-effective *10G or N X 10G Backbones*, complemented by
    Point-to-point "Light-paths" for "Data Intensive

2nd Quiet Revolution in Science and R&E Networks Continues

- 2000-2007: We developed the knowledge to use long distance networks efficiently, at high occupancy, for the first time
  - "Demystification" of large long range data flows
    with TCP:
    - Up to 10 Gbps per flow; 151 Gbps aggregate
    - 17.7 Gbps disk to disk on one 10 Gbps link (bi-directionally)
  - Major advances: in the TCP stack (e.g. FAST TCP), Linux Kernel (2.6.19-20), end

NSFNET
The Partnership That Changed The World

CERN, FNAL, BNL, SLAC, UM, UF, PUB, ESNet, I2…

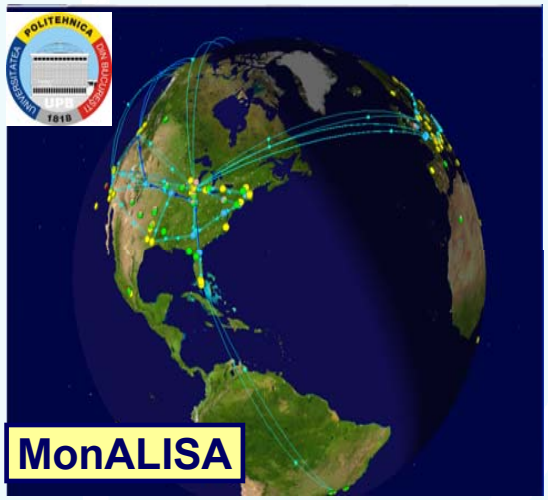

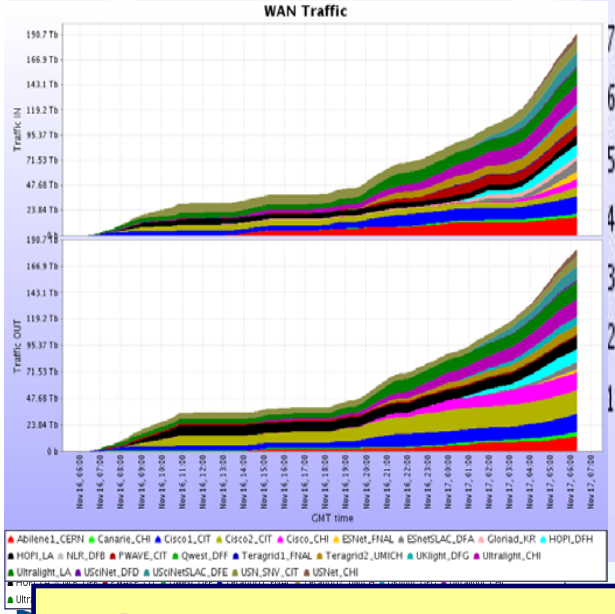

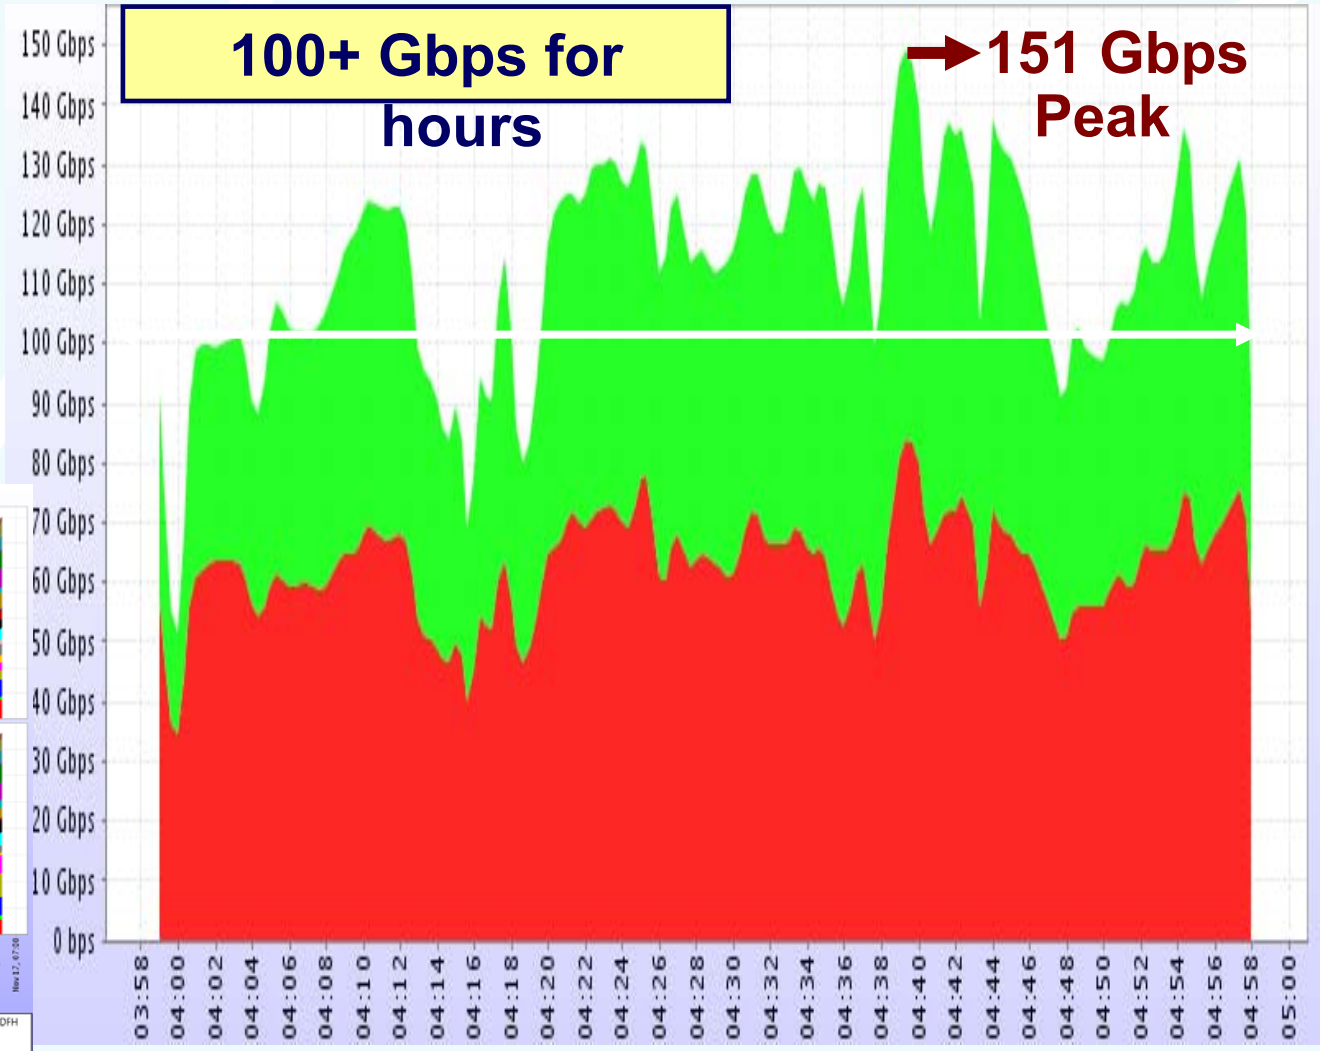

475 TB Total in < 24h; Sustained Rate of 1.1 Petabyte Per Day

Changed The World

FDT: Fast Data Transport
SC06 Results 11/14 – 11/15/06

- Stable disk-to-disk flows Tampa-Caltech:
  Stepping up to 10-to-10 and 8-to-8 1U Server-pairs 9 + 7 = 16 Gbps; then Solid overnight. Using One 10G link

I. Legrand

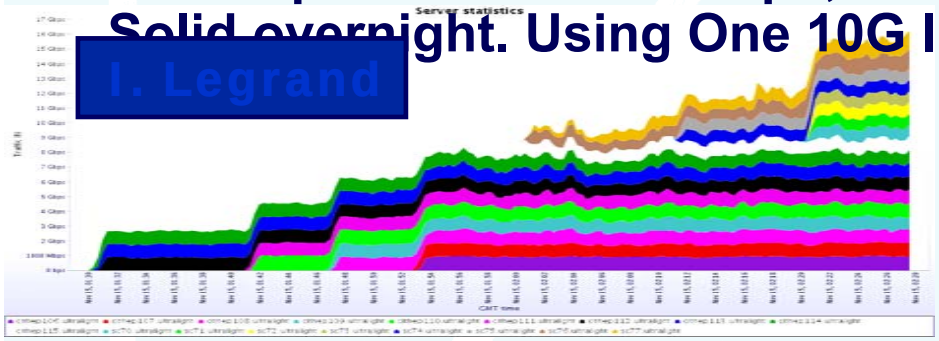

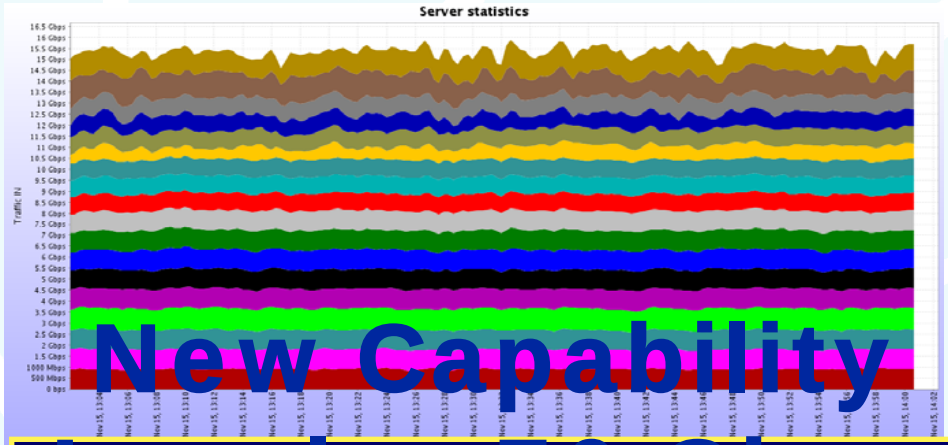

New Capability Level: ~70 Gbps per rack of low

Efficient Data Transfers

- Reading and writing at disk speed over WANs (with TCP) for the first time
- SC06 Results: 17.7 Gbps on one link; 8.6 Gbps to/from Korea
- In Java (NIO libraries): runs on all major platforms.
- Based on an asynchronous, multithreaded system
- Streams a *dataset* (list of files) continuously, from a managed pool of buffers at a moderated rate, through an open TCP socket
  - Smooth data flow from each disk to/from the network
  - No protocol start-phase

SC07 Results: 80+ Gbps, One Rack of Servers on 5 10G Waves 11/14/07

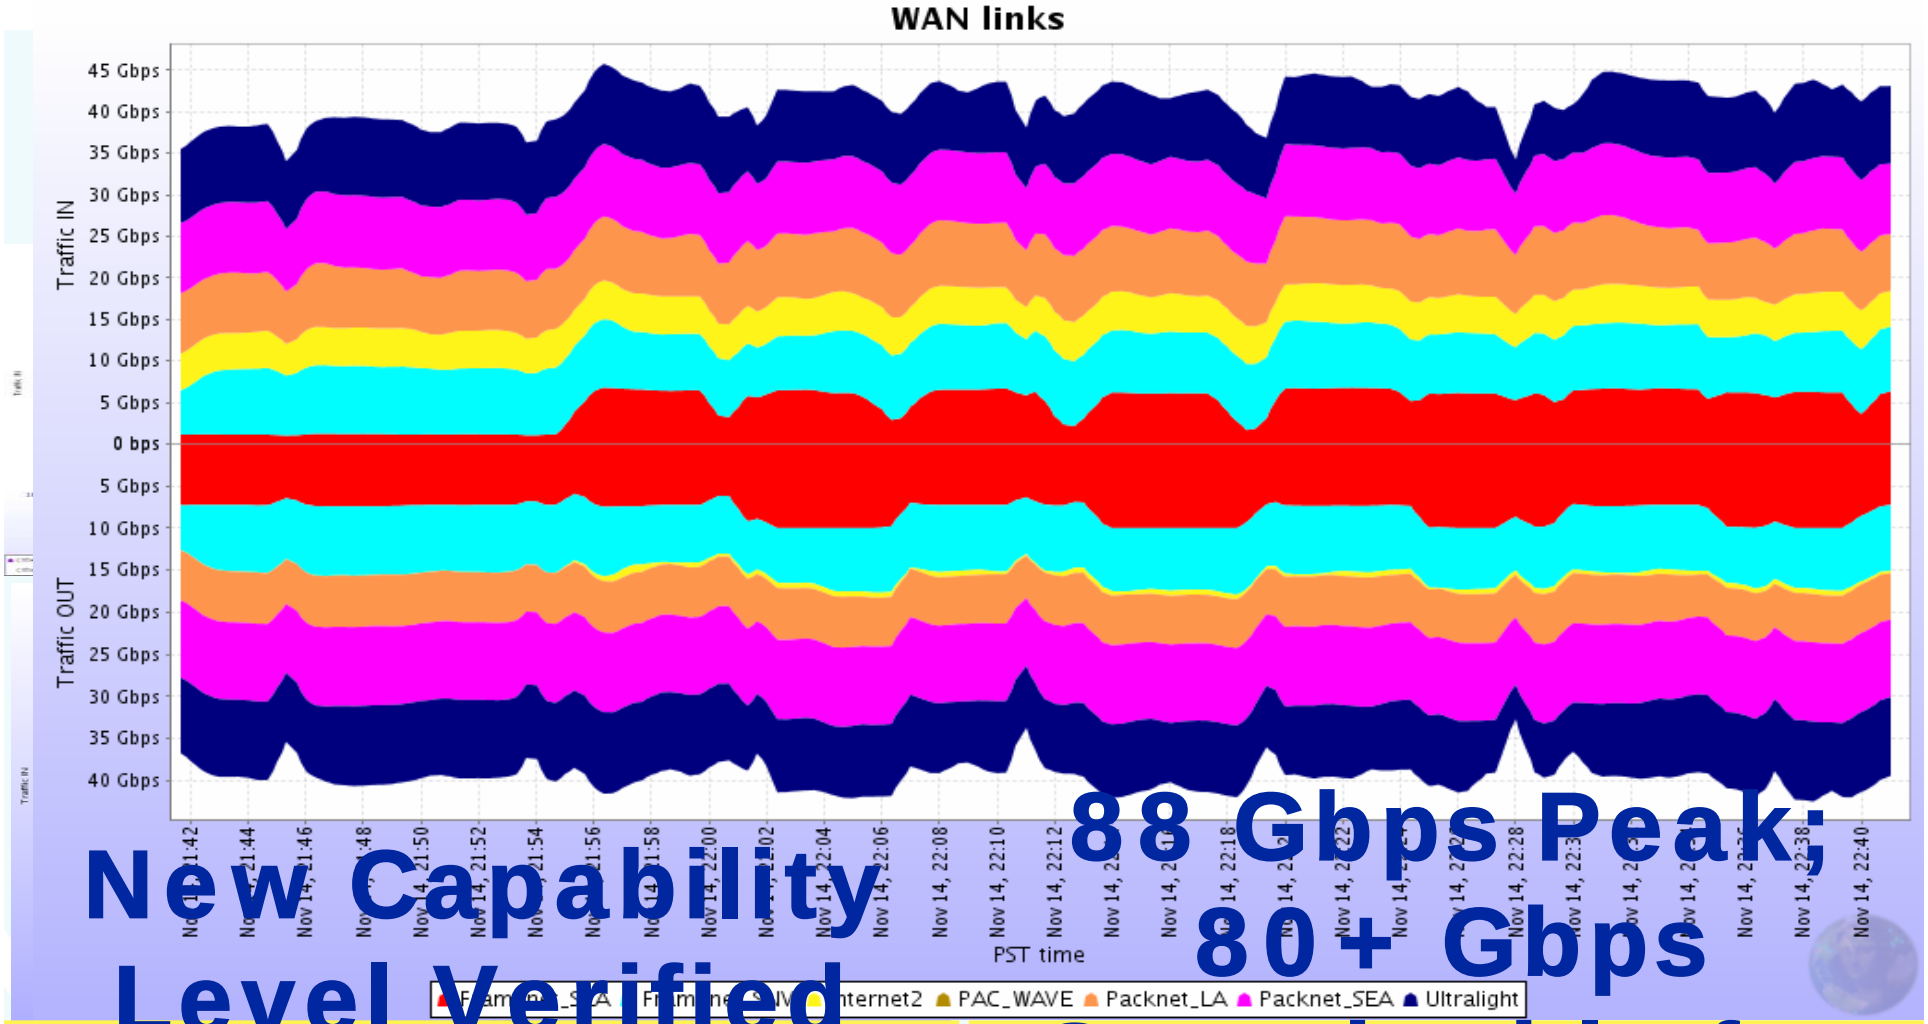

88 Gbps Peak; 80+ Gbps Sustainable for Hours, with

New Capability Level Verified

Mature Open

Networks, Compute Nodes, Running Jobs, Processes

- Built for speed and global scale: 5k messages/sec/server ; multi-threaded engine schedules ML services
- Monitors, controls and optimizes large-scale distributed systems

CENIC Innovation Award '06

Internet2 IDEA Award '07

time for a variety of tasks

MonALISA Today

Running 24 X 7 (5 Years)

Now at 340 Sites

- Collecting > 1,000,000 parameters in near real-time
- Update rate of >20,000 parameter-updates per sec
- Monitoring
  - > 40,000 CPUs
  - > 100 WAN Links
- Thousands of Grid jobs running

Major Communities

- OSG
- CMS
- ALICE
- D0
- STAR
- VRVS, EVO
- LGC RUSSIA
- SE Europe GRID
- APAC Grid
- UNAM Grid (Mx)
- ITU
- RoEduNET
- Internet2
- US LHCNet
- ULTRALIGHT
- GLORIAD
- Enlightened

A Sampling of End To End Monitoring and Global

FDT Automatic Path Recovery Field: Trial Fiber Cut Emulations

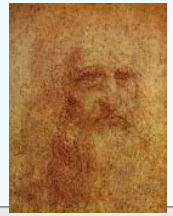

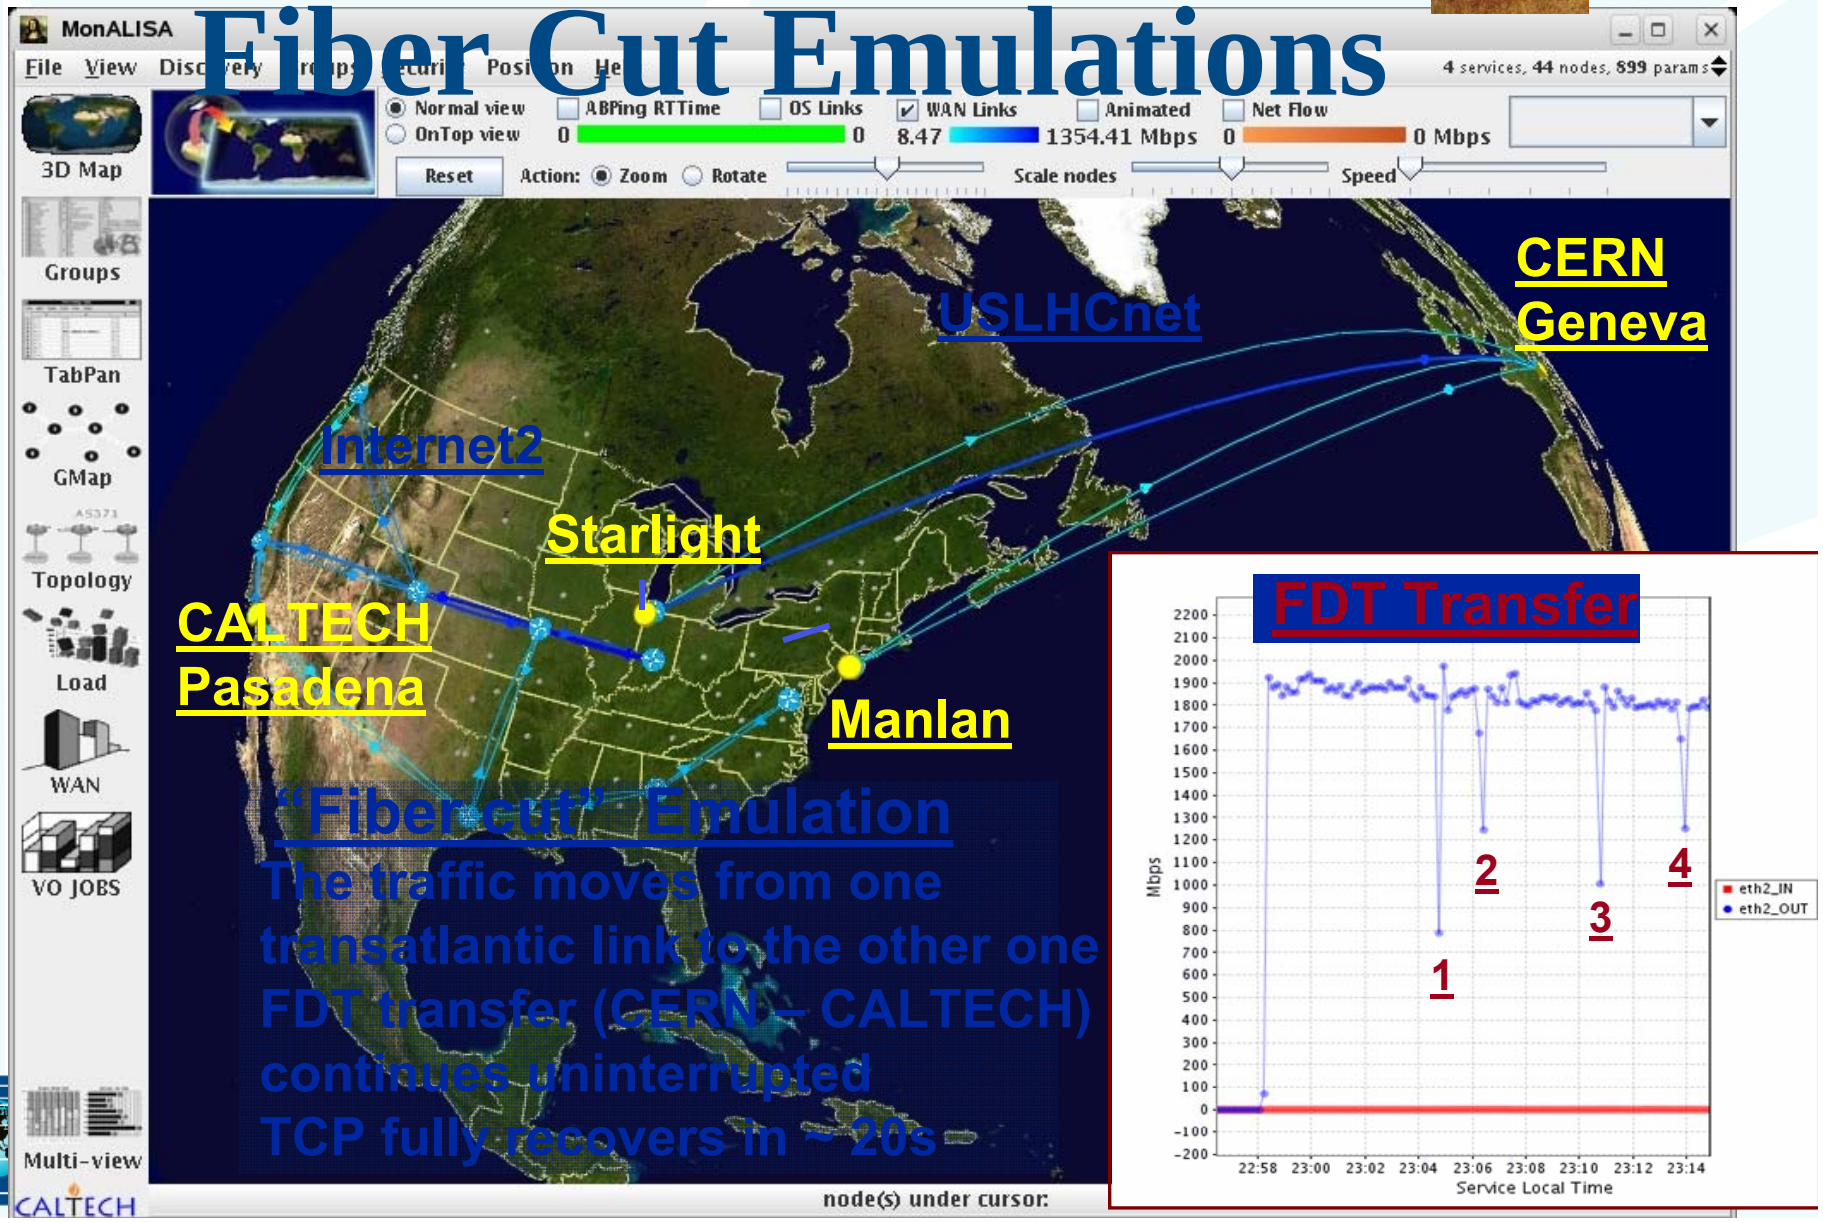

SCIC in 2006-2007

http://cern.ch/icfa-scic

Focus on the Digital Divide (Since 2002)

Rapid Progress; Deepening Divide

- *Main Report: "Networking for HENP"* [H. Newman, et al.]
  - Includes Updates on the Digital Divide, World
    Network Status; Brief updates on Monitoring
    and Advanced Technologies
  - 31 Appendices: A World Network Overview

Status and Plans for the Next Few Years of

Throughput Improvements 1995-2007

Progress: but Digital Divide is Mostly

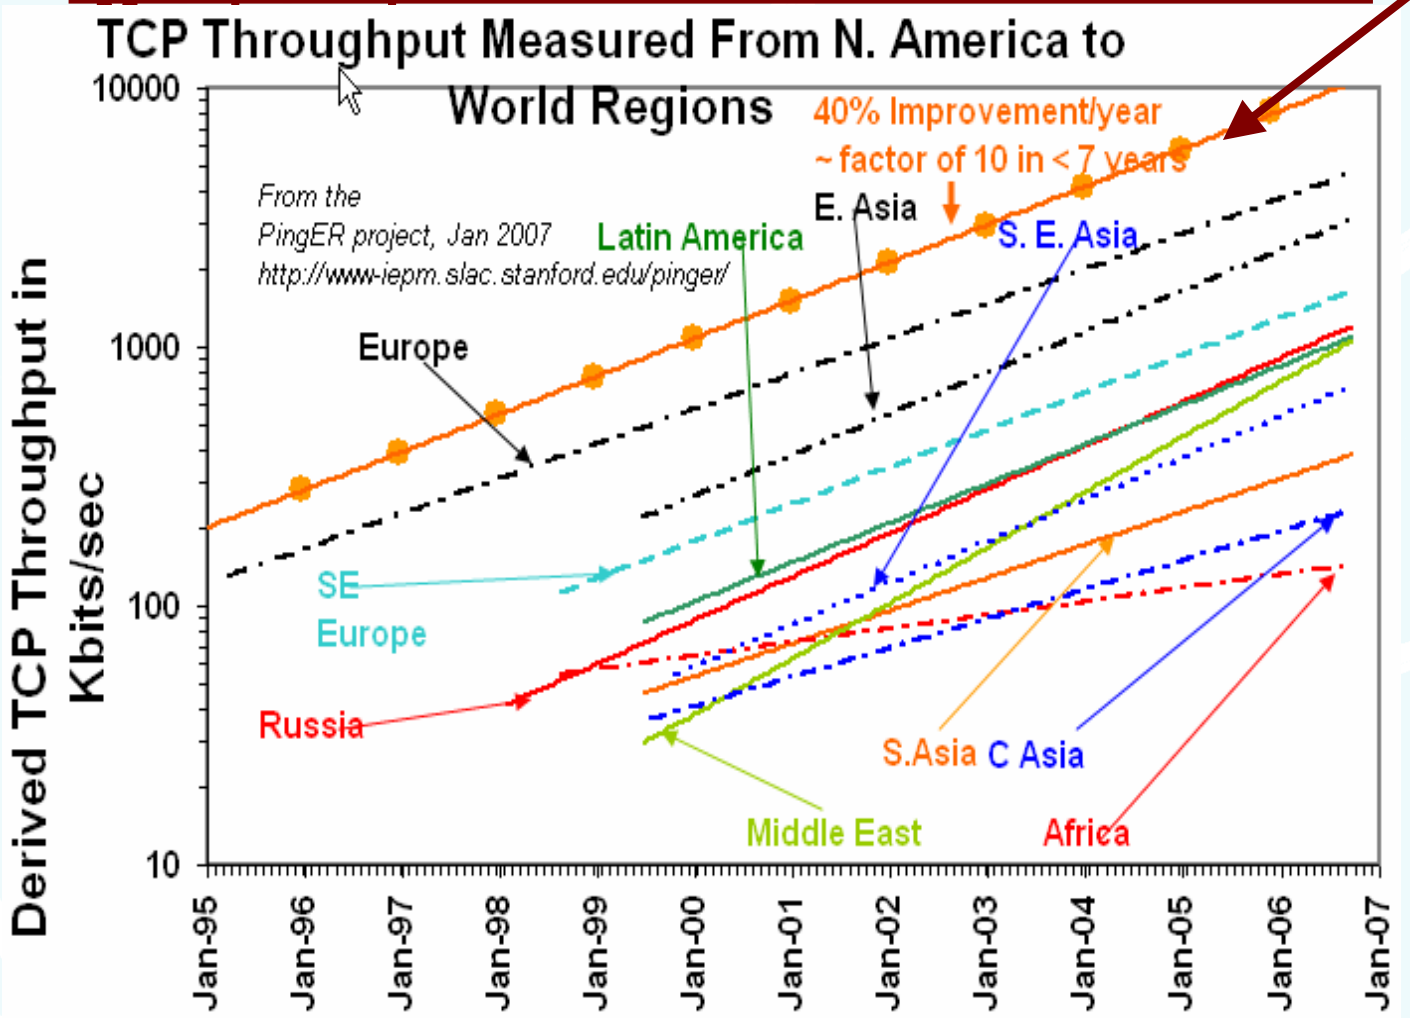

40% annual improvement ~10X Per 7 yrs

Behind Europe

- **6 Yrs: Russia, Latin America**
- **7 Yrs: Mid-East, SE Asia**
- **8-9 Yrs: So. Asia**
- **11 Yrs: Cent. Asia**
- **12 Yrs: Africa**

India, Pakistan, Central Asia, and Africa are in Danger of Falling Even Farther Behind

R. Cottrell

Bandwidth of TCP < MSS/(RTT*Sqrt(Loss))

Matthis et al., Computer Communication Review 27(3), July 1997