

# Riding the Tiger: Technology Trends and Scientific Breakthroughs

Dan Reed

Director, NCSA and the Alliance  
Chief Architect, NSF TeraGrid

Edward William and Jane Marr Gutgsell Professor  
University of Illinois



reed@ncsa.uiuc.edu



National Center for Supercomputing Applications



National Computational Science ALLIANCE

## Presentation Outline

- **Revolution and risks**

- technology trends
- community building



- **Applications, Grids, and beyond**

- NSF TeraGrid, ETF, ...
- emerging applications



- **Futures**

- short-term opportunities
- expeditions and settlements




National Center for Supercomputing Applications




National Computational Science ALLIANCE

# 21<sup>st</sup> Century Science & Engineering


- **The three fold way**
  - theory
  - experiment
  - computational simulation
- **Supported by**
  - multimodal collaboration systems
  - distributed, multi-petabyte data archives
  - leading edge computing systems
  - distributed experimental facilities
  - internationally distributed multidisciplinary teams
- **Collectively defining a new future**
  - creation of 21<sup>st</sup> century IT infrastructure
  - sustainable, multidisciplinary communities



Simulation




National Center for Supercomputing Applications

National Computational Science 

# Consilience

- **The unity of ideas**
  - integration and inclusive themes
- **The Standard Model (physics)**
  - quantum theory
    - electroweak and strong forces
  - seeking the “Grand Unified Theory”
    - gravity integration and rationale for mass
- **The Central Dogma (biology)**
  - DNA carries genetic information
  - DNA is transcribed to RNA
  - RNA is translated proteins



**Consilience (n):** The fact of ‘jumping together’ or agreeing; coincidence, concurrence; said of the accordance of two or more inductions drawn from different groups of phenomena.




National Center for Supercomputing Applications

National Computational Science 

# The Power of Exponentials

- **Non-intuitive exponentials**
  - they look linear for a long time
  - then a phase transition occurs
- **Understanding change**
  - ride the exponential
  - recognize the inflection points
- **Quantitative begets qualitative**
  - exponential change
  - a breakthrough opportunity
- **Change brings opportunity**
  - *and also danger*
    - *Andy Grove is right about paranoia*



IBM



National Center for Supercomputing Applications



National Computational Science ALLIANCE

# Riding the Exponentials

- **Non-intuitive exponentials**
  - they look linear, then a phase transition occurs
- **Understanding change**
  - ride the exponential/recognize inflection points
- **Change brings opportunity/danger**
- **Fiber optics, the dot.com gift**
  - 1000 wavelength DWDM
    - Lucent and NTT
  - 10 Tb capacity/pair
    - multiple 400 pair cables
- **Large format displays**
  - human scale presentation
  - commodity components
    - projectors, Linux clusters, aluminum frames
  - high resolution NCSA wall
    - 8192 x 3840 pixels



National Center for Supercomputing Applications




National Computational Science ALLIANCE

# Digital Reality: Riding Exponentials


- **Megabyte**
  - a small novel
- **Gigabyte**
  - a pickup truck filled with paper or a DVD
- **Terabyte: one thousand gigabytes**
  - the text in one million books
  - entire U.S. Library of Congress: ten terabytes of text
- **Petabyte: one thousand terabytes**
  - 1-2 petabytes equals *all* academic research library holdings
- **Exabyte: one thousand petabytes**
  - 2 exabytes of data generated annually
  - 5 exabytes of words spoken in the history of humanity



# Revolution and Risks



Revolution



Risks

“...it’s a lot better to *initiate* change while you can than it is to try to react and adjust to it.”

# Revolution and Risks: Failing Gloriously

## • PLATO

- early online community
- CAI courseware and support
- computer music and interactive displays
- lessons gave us Lotus Notes™ and browser Mosaic™



## • ILLIAC IV

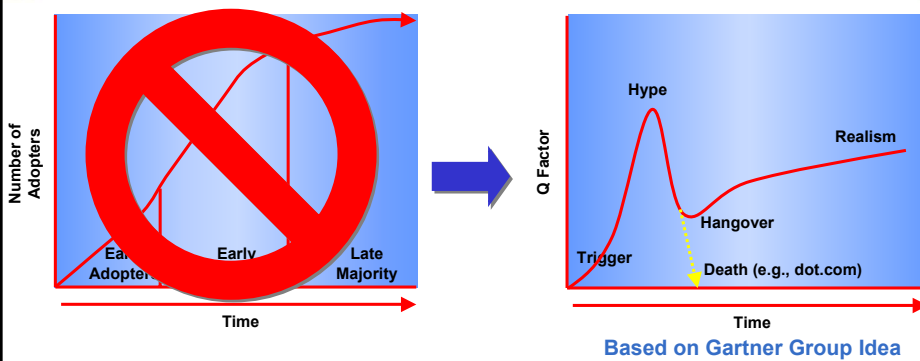
- world's fastest computer at design goal
- parallel programming languages
- first real-time fully CAD designed
- ARPANet for remote access
- lessons shaped generations of systems



Defining Experiences!



# Technology Adoption Cycles



***“A small group of thoughtful people could change the world. Indeed, it's the only thing that ever has.”***

***Margaret Mead***



# Lessons from the Dot.Com Bust

- **Pets.com**
  - founded in 1999
  - February 2000 IPO for \$82.5M
    - IPO at \$11/share
    - ultimately closed at \$0.22/share
  - national TV marketing
    - \$2M Super Bowl spot
    - \$20M total in 1999
  - 1.8M visitors/month by October 2000
  - 570,000 customers
- **Lessons**
  - outspent rivals 2:1
  - overestimated online customer base
  - ineffective value proposition
    - did not offer compelling rationale
    - sold below cost with free shipping



My grandmother told me, "Good judgment comes from experience. Experience comes from bad judgment."

Sock Puppet




National Center for Supercomputing Applications



National Computational Science ALLIANCE

# The Six Computing Eras

- **Big Iron (post WW II)**
- **Mainframe ('60s/'70s)**
- **Workstations ('70s/'80s)**
- **PCs ('80s/'90s)**
- **Internet ('90s)**
- **Implicit computing (21st century)**
  - *embedded intelligence in everyday objects*
  - *broadband wireless networking*
  - *number of processors/person → infinity*
    - $O(100M)$  PCs and  $O(8B)$  embedded processors/year
- **Electronic tags and intelligent objects**
  - creating the ubiquitous infosphere
  - SC02 smart tag experience



Silicon Valley

<http://www.toaster.net/wireless/aplist.php>




National Center for Supercomputing Applications



National Computational Science ALLIANCE

# Open Source Software

- **Open source development**
  - not all software has an economic market
    - market cannot sustain development costs
    - HPC software is the canonical example
  - but some does
    - Linux, Apache, ...
- **U.S. PITAC report on open source**
  - <http://www.ccic.gov/ac/pres-oss-11sep00.pdf>
  - urges investment in open source projects
- **NCSA open source license**
  - <http://www.opensource.org/licenses/UoI-NCSA.html>



“The Linux issue is whether this is a fundamentally disruptive technology, like the microprocessor and the Internet. We’re betting that it is.”

Irving-Wladawsky-Berger, IBM




National Center for Supercomputing Applications



National Computational Science ALLIANCE

# Open Source Community Building



“In a Box” Initiative

- **Packaged software and tools**
  - cluster software (aka OSCAR)
  - Grid infrastructure
  - Access Grid for collaboration
  - display wall for visualization
- **CDs distributed at SC01**
  - Alliance-wide “In a Box” effort
  - packaged for community building
  - also available via the web
    - [www.ncsa.uiuc.edu](http://www.ncsa.uiuc.edu)
- **Gelato consortium**
  - Linux development
  - current partners
    - ESIEE, NCSA, UIUC, Waterloo
    - HP, UNSW, and Tsinghua




National Center for Supercomputing Applications




National Computational Science ALLIANCE

# NCSA Terascale Linux Clusters

- **1 TF IA-32 Pentium III cluster (Platinum)**
  - 512 1 GHz dual processor nodes
  - Myrinet 2000 interconnect
  - 5 TB of RAID storage
  - #41 on November 2001 Top500 list, 594 GF
- **1 TF IA-64 Itanium cluster (Titan)**
  - 164 800 MHz dual processor nodes
  - Myrinet 2000 interconnect
  - #34 on November 2001 Top500 list, 678 GF
- **Breakthrough calculations on both**
  - molecular dynamics (Schulten)
    - first nanosecond/day calculations
  - gas dynamics (Woodward)
  - others underway



IBM intel Myricom



## NCSA machine room expansion

- capacity to 100 TF and expandable
- dedicated September 5, 2001




National Center for Supercomputing Applications

National Computational Science ALLIANCE

# The U.S. NSF PACI TeraGrid



NCSA, SDSC, Argonne, Caltech plus \$7.5M Illinois I-WIRE Initiative



"Internet" circa 1969



Internet circa 1999

Source: Bill Cheswick

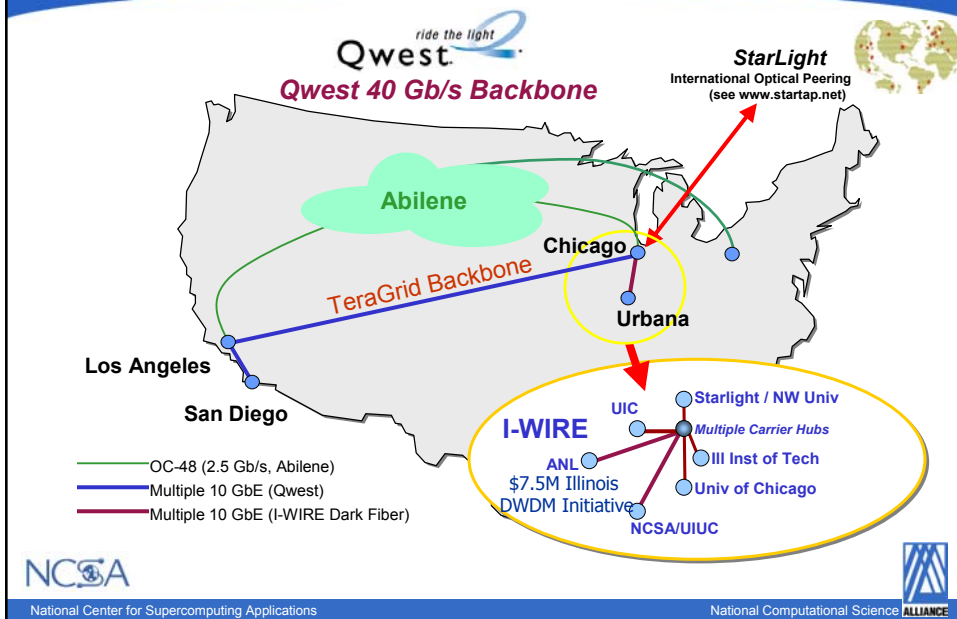


National Center for Supercomputing Applications

National Computational Science ALLIANCE



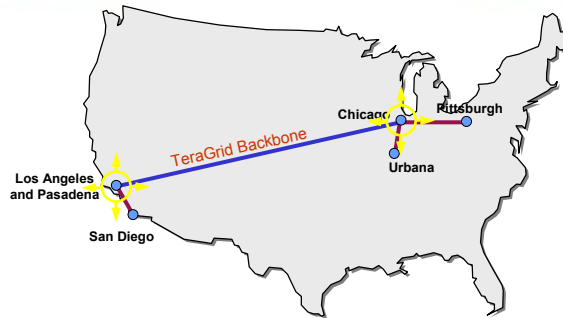
# TeraGrid Network Backbone



# Measuring TeraGrid Success

- **An extensible Grid**
  - design principles assume heterogeneity and more than four sites
    - scalable, replicable, and interoperable standards
    - documented design and infrastructure
    - encourage, support, and leverage open source software
- **Breakthrough science via new capabilities**
  - site capabilities more powerful than existing PACI resources
  - integrated site resources as a coordinated system
  - extant and new user base
    - current users and new communities requiring Grids
- **A pathway for current users**
  - evolutionary paths from current practice
    - provide examples, tools, and training to exploit Grid capabilities
  - *user support, user support, and user support*

# TeraGrid, ETF, and Beyond



- **Extensible Terascale Facility (ETF)**
  - TeraGrid generalization and extension
  - preparation for additional sites in FY03
- **NCSA, SDSC, Argonne, Caltech, and PSC**
  - plan due to NSF June 3
- **Connection points for interface to additional sites**

NCSA • Los Angeles and Chicago

National Center for Supercomputing Applications

National Computational Science ALLIANCE



# A Tool and Market for Every Task



- **Multiple development models**
  - standard market economics
  - long-term government investment
- **Each targets different applications**

NCSA • understand application needs




National Center for Supercomputing Applications

National Computational Science ALLIANCE



# The Really Big Questions

- **Life and nature**
  - structures, processes, and interactions
- **Matter and universe**
  - origins, structure, manipulation, and futures
  - interactions, systems, and context
- **Humanity**
  - creativity, socialization, and community
- **Answering big questions requires**
  - boldness to engage opportunities
  - new approaches and infrastructure
  - new collaborations
  - interdisciplinary partnerships




ATP: Klaus Schulten et al

NCSA

National Center for Supercomputing Applications

# Cosmological Challenges

- **Much more mass than we can see**
  - more than 90% of believed mass is not visible
  - supporting evidence
    - rotation curves in spiral galaxies
    - globular clusters and gas by elliptical galaxies
    - gravitational lensing by clusters
- **Possible missing mass candidates**
  - baryonic
    - Massive Compact Halo Objects (MACHOS)
  - non-baryonic
    - Weakly Interacting Massive Particles (WIMPS)
- **Massive new simulations required**
  - *multi-particle and MHD*



NCSA

National Center for Supercomputing Applications

Naoki Yoshida and Lars Hernquist

National Computational Science ALLIANCE

# Molecular Cloud Collapse

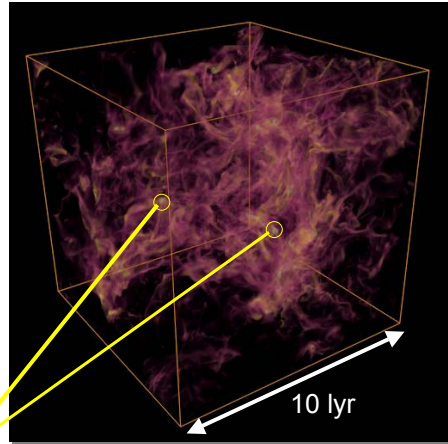
- **Molecular cloud simulation**

- ZEUS-MP code
- turbulent and magnetized
- self-gravitating

- **Current state of the art**

- 512<sup>3</sup> grid
- NCSA Origin
  - 256 processors
- 100,000 processor hours

Log of surface density



Collapsing dense cores

Source: Li and Norman

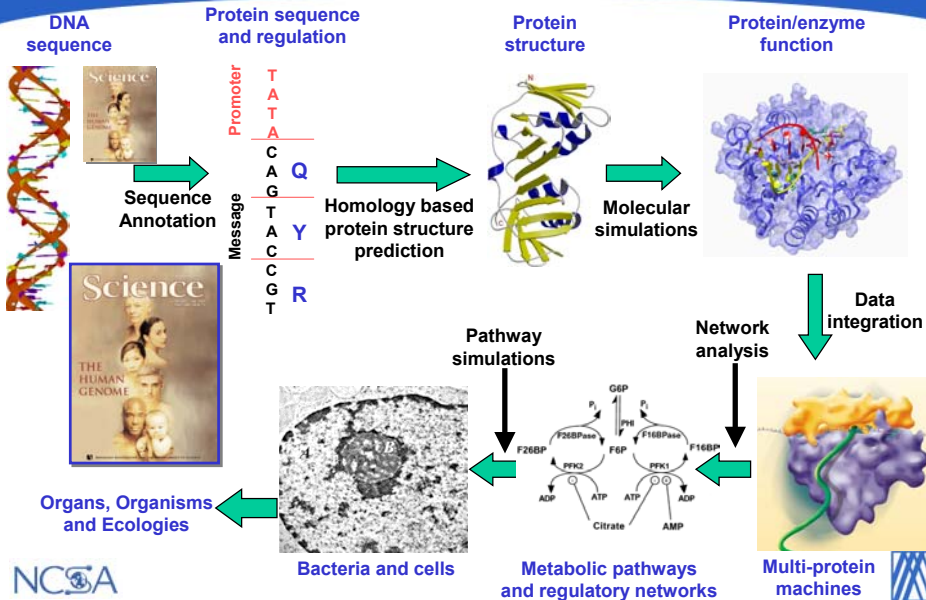


National Center for Supercomputing Applications



National Computational Science

# Systems Biology: Exaflop Challenges



National Center for Supercomputing Applications



National Computational Science

# Grid Examples


- **NSF Network for Earthquake Engineering Simulation (NEES)**

- seamless testing and simulation
- earthquake hazard mitigation
  - structural, geotechnical & tsunami
- national IT infrastructure
- leadership
  - NCSA and UIUC civil engineering



- **Distributed weather prediction**

- distributed radar and satellite data
- multiple application domains
- near real-time response needed



National Center for Supercomputing Applications

National Computational Science

# Storm Scale Prediction

- **Sample four hour forecast**

- Center for Analysis and Prediction of Storms
- Advanced Regional Prediction System
- full-physics mesoscale prediction system



- **Execution environment**

- NCSA Itanium Linux Cluster



- **Fort Worth forecast**

- four hour prediction, 3 km grid
- initial state includes assimilation of
  - WSR-88D reflectivity and radial velocity data
  - surface and upper air data, satellite, and wind



National Center for Supercomputing Applications

Source: Kelvin Droegemeier

National Computational Science

# Distributed Virtual Astronomy


- Futures**

- homogeneous, multi-wavelength data
- observations of millions of objects
  - mega-sky surveys (2MASS, SLOAN, ...)



- U.S. National Virtual Observatory (NVO)**

- Caltech, JHU, ALMA, HST, ...
- enable discovery in archives via new tools



- Data mining and archives**

- discovering significant patterns
  - analysis of rich image/catalog databases
- understanding complex astrophysical systems
  - integrated data/large numerical simulations

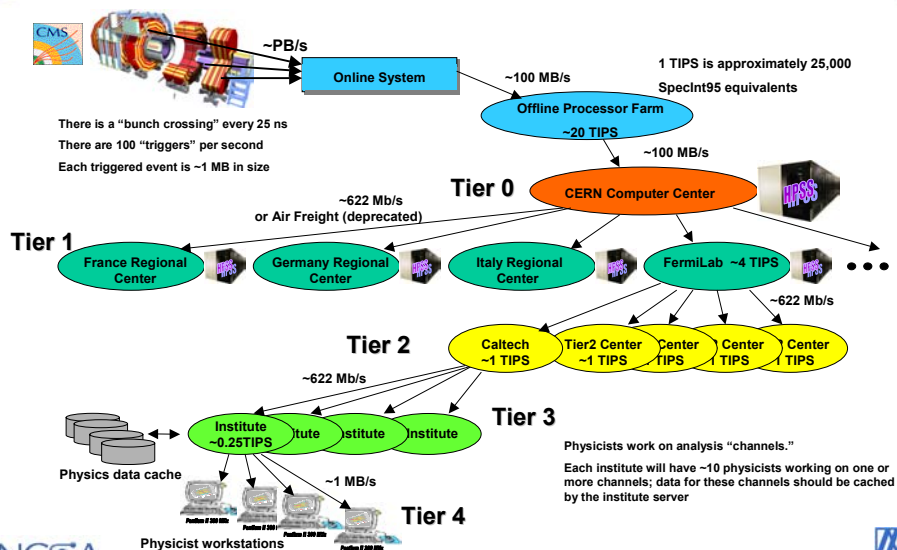


National Center for Supercomputing Applications



National Computational Science ALLIANCE

# Data Grids for High Energy Physics



National Center for Supercomputing Applications

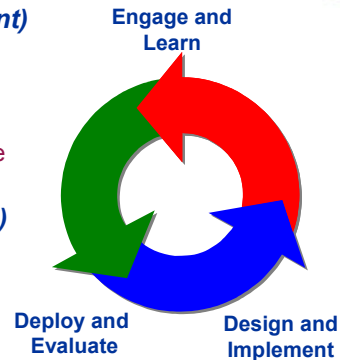
Source: Harvey Newman/Ian Foster



National Computational Science ALLIANCE


# The Virtuous Cycle

- **Live in the future (*research and development*)**
  - track evolving infrastructure trends
  - prototype advanced infrastructure
- **Ride the exponentials**
  - see qualitative change from qualitative change
  - recognize the inflection points
- **Bring insights to production (*infrastructure*)**
  - translate prototypes into production
  - empower applications and learn
  - expand the user community
- **Empower new science breakthroughs**
  - big questions and big answers
- **Learn from experience (good and bad)**
  - enhance the good
  - fix the bad and explore alternatives
- **Recognize that the cycles continue *ad infinitum***
  - commit to continual investment



# HPF: I Feel Your Pain


- **Lessons**
  - **irregular data structures**
    - language support needed
  - **data distributions**
    - best not part of the language
  - **compilation and tuning**
    - major research challenges
    - inverse mappings for tuning
- **Observations**
  - HPF locality model is semi-implicit
  - we expected too much too soon
    - 5-10 years needed for mature compilers and idiom identification
    - languages and idioms must evolve



# MPI: It Hurts So Good

## • Observations

- “assembly language” of parallel computing
  - lowest common denominator
- upfront effort repaid by
  - system portability and explicit locality management
- *remember what Churchill said about democracy*
  - *it applies to MPI as well*



## • Costs and implications

- human productivity
  - low-level programming model
- software innovation
  - limited development of alternatives



National Center for Supercomputing Applications



National Computational Science ALLIANCE

# Infrastructure Futures

## • Petaflops and beyond

- petaops for breakthrough discovery/advantage
  - rethinking technology leverage
- desktop teraflops at O(\$25K)
  - revolutionizing “day to day” access
- terabit and ubiquitous sensor networks
  - reducing the tyranny of distance



## • Data explosion

- multi-petabyte data archives and personal petabytes
  - data to knowledge and insight
- cross-lingual indexing and management
  - multidisciplinary coordination and integration



## • Ubiquitous infosphere/pervasive Grid environment

- thousands of lambdas/fiber and terabits/second
  - the Grid writ large (building up)
- wireless, sub-millimeter sensor/actuators
- adaptive, responsive, contextual environments
  - the Grid writ small (building up)



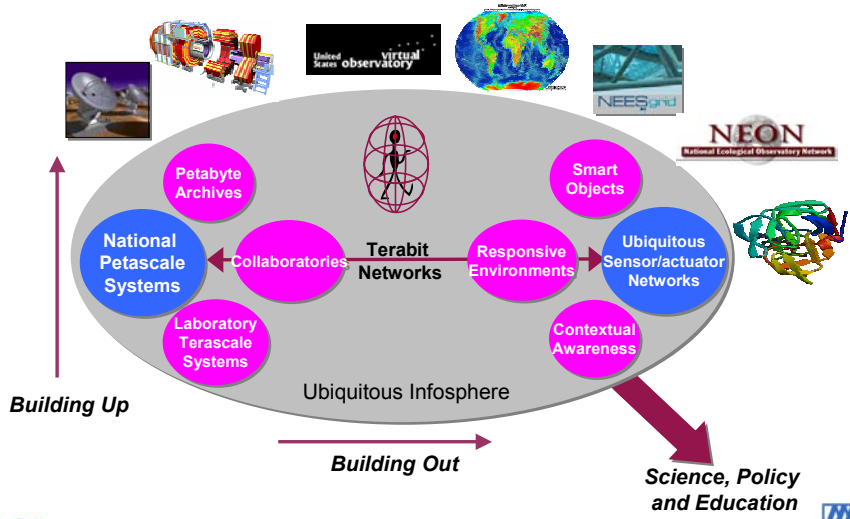
National Center for Supercomputing Applications



National Computational Science ALLIANCE



# Responding With Breakthrough Science




NCSA

National Center for Supercomputing Applications

National Computational Science ALLIANCE

# It's About the Science




Gravitational Waves: Ed Seidel



Turbulence: Paul Woodward



Cholesterol: Eric Jakobsson et al



ATP: Klaus Schulten et al

NCSA

National Center for Supercomputing Applications

National Computational Science ALLIANCE

# Questions?



NCSA

National Center for Supercomputing Applications



National Computational Science ALLIANCE