The Future: Release 2016

Dan Reed
reed@renci.org

Chancellor's Eminent Professor
Vice Chancellor for Information Technology
University of North Carolina at Chapel Hill

Director, Renaissance Computing Institute (RENCI)

renci
renaissance computing institute

Presentation Outline

"Where there is no vision, the people perish."
Proverbs 29:18

- **A geobiology primer**
  - nature has a few lessons to share
- **A bit of computing history**
  - heed the words of Santayana
- **A snapshot of current reality**
  - power, scaling and reliability
- **Musings on the future**
  - technical, political and scientific

renci
renaissance computing institute

A Geobiology Primer

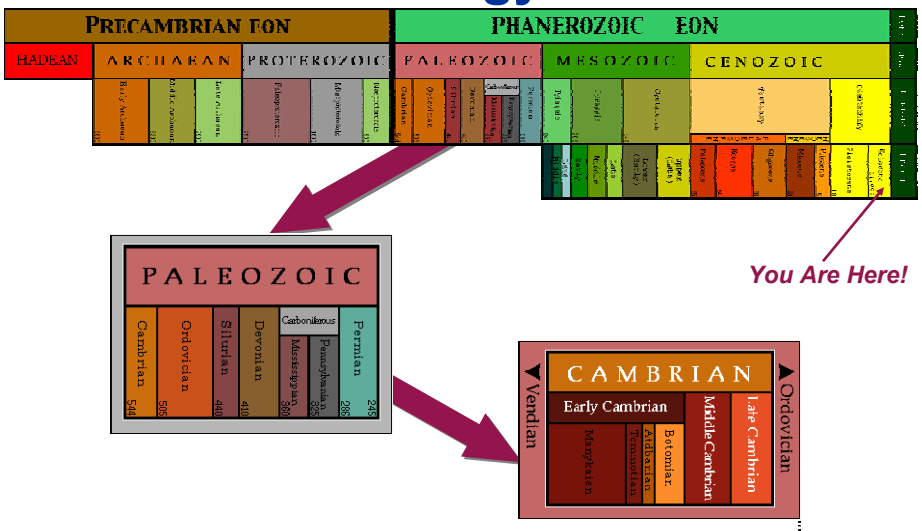

Source: UC-Berkeley Museum of Paleontology

renci

A Few Key Points

- ***Cenozoic/Quaternary (2M years)***
  - *humans dominate near the end (100K years)*
- **Cenozoic/Tertiary (65M years)**
  - mammals dominate as large animals
  - *K/T mass extinction due to meteor impact*
- **Mesozoic/Cretaceous (130M years)**
  - first flowers and primates appear
- **Mesozoic/Jurassic (180M years)**
  - dinosaurs dominate; birds appear
- **Mesozoic/Triassic (230M years)**
  - dinosaurs appear
  - Atlantic Ocean forms
- **Paleozoic/Permian (270M years)**
  - reptiles dominate land; seas contract
  - *Permian mass extinction (95% of all life)*

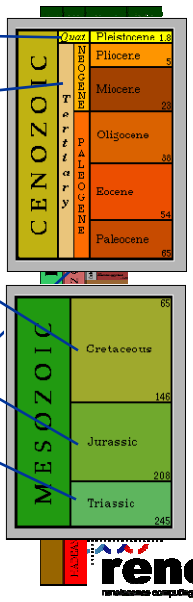

renci

High Performance Computing

- **IBM Stretch**
  - design goal: 100-200X IBM 704
    - world's fastest machine until 1964
    - parallelism as an enabler
  - design timeline
    - 1961 LASL delivery; retired 1971
    - 1962 Harvest NSA delivery; retired 1976
    - $13.5M list price ($95M in current $)
  - architectural features
    - interleaving [illegible]
    - [illegible] forwarding
- **Illinois/Burroughs ILLIAC IV**
  - world's fastest machine as design goal
    - launched 1974, retired 1982
    - $30M circa 1972 ($130M in current $)
  - 64 processor SIMD (1/4th design target)
  - array language support (Glypnr and IVTRAN)
  - thin film memory (2K words/processor)
  - ARPANET for remote access

Failing Gloriously

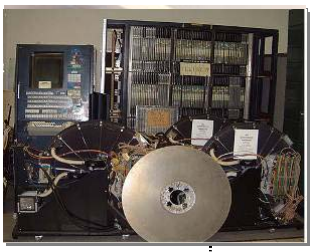

renci

Technology Gambles

By sacrificing a factor of roughly three in circuit speed, it's possible that we could have built a more reliable multi-quadrant system in less time, for no more money, and with a comparable overall performance. The same concern for the last drop of performance hurt us as well in the secondary (parallel disk) and tertiary (laser) stores.

Dan Slotnick

renci

Hunan-Computer Symbiosis

- **PLATO (Programmed Logic for Automated Teaching Operations)**
  - begun in 1960, led by Illinois' Don Bitzer
  - several spinoffs via CDC, NovaNET,
  - Illinois classroom use until 1985
    - 10 million hours 1978-1985
    - over 3 million hours in Notes
  - early online community
    - computer music and plasma touch panel displays
  - lessons later gave us Lotus Notes and Mosaic™

Failing Gloriously

- **Project MAC**
  - Man and Computer or Multiple Access Computer
  - $25M ARPA funding from 1963-1970
    - >$100M in current $
  - J.C.R. Licklider suggestion, Robert Fano leadership
  - Multiplexed Information and Computing Service (MULTICS)
    - virtual memory, hierarchical file systems, time sharing, …
  - a host of innovative ideas and collaborations

renci

MULTICS and UNIX

"... the problem was the increasing obviousness of the failure of MULTICS to deliver promptly any sort of usable system, let alone the panacea envisioned earlier."

Dennis Ritchie

renci

Human-Computer Symbiosis

It seems reasonable to envision, for a time 10 or 15 years hence, a 'thinking center' that will incorporate the functions of present-day libraries together with anticipated advances in information storage and retrieval.

The picture readily enlarges itself into a network of such centers, connected to one another by wide-band communication lines and to individual users by leased-wire services. In such a system, the speed of the computers would be balanced, and the cost of the gigantic memories and the sophisticated programs would be divided by the number of users.

J.C.R. Licklider, 1960

renci
renaissance computing institute

ARPANET

Len Kleinrock

BBN IMP Team

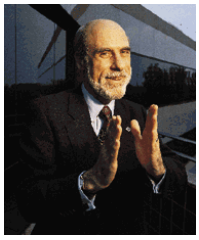

Vint Cerf

Larry Roberts

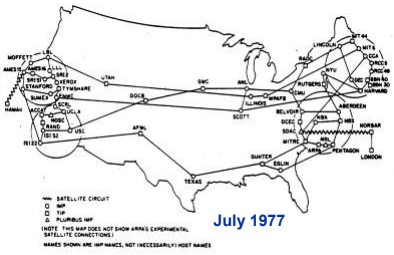

Note the timescale!

Bob Kahn

renci
renaissance computing institute

Web and Social Processes

- **Google**
  - it's a search engine, it's a verb, …
- **Blogs**
  - published self-expression
- **Instant Messenger**
  - social networks
- **Wireless messaging**
  - semi-synchronous
- **Internet commerce**
  - the dot.com boom/bust
  - EBay, Amazon
- **Spam, phishing, …**
  - anti-social behavior

"On the Internet, nobody knows you're a dog."

Digital Book Scanning

- **Features**
  - pneumatic page turning
    - 1500-3000 pages/minute
    - 200-800 dpi
  - Google book project

www.4digitalbooks.com

Google Book Search BETA

books.google.com

What Can Geobiology Teach Us?

- **Nature**
  - explores alternatives and fills ecological niches
  - responds to both gradual and rapid change
  - prunes failed options without remorse
  - supports homeostasis and immunity
  - develops social structures and specialization
    - ants, termites, …

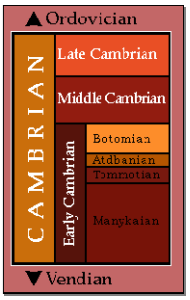

- **Computing**
  - limited exploration of alternatives
  - encourages homogeneity not diversity
    - attack prone, security breaches, …
  - *exponential change and punctuated equilibrium*

renci

Ten Years: Past and Future

- **Looking back, in the public mind**
  - there were few or no experiences with …
    - web sites, email, spam, phishing, computer viruses
    - e-commerce, digital photography or telephony
    - digital job offshoring
  - cell phones were rare and expensive
  - a Sony Walkman was state of the art
  - CDs were still pretty cool
  - WiFi was almost unknown

- **Looking forward ten years …**

renci

Imagine a Future Where …

- **Your car finds a parking space for you**
  - … and also parks the car (a 2008 Lexus feature)
- **The radio only plays music you love**
  - … because it knows every song you've ever heard
- **Your phone only rings when you want to answer**
  - … because it knows your emotional state

- **Your body calls an ambulance when you are ill**
  - … via implanted, biologically powered diagnostic sensors
- **Your DNA sample determines personalized treatment**
  - … because genotype-phenotype models are specific
- **Your office adjusts its behavior to your needs**
  - … because it knows (semantically) what you are doing

renci

Imagine a Future Where …

- **Your every physical movement is tracked/logged**
  - … by embedded sensors on all human artifacts

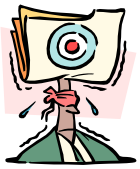

- **Your neighbors know all the books you read**
  - … because your electronic financial identity was stolen
- **Your every call is monitored for content**
  - … by deep semantic analysis and logging

- **Your utilities fail due to a virus attack**
  - … because security was penetrated by a 10 year old
- **Your DNA sample/lifestyle determine health cost**
  - … because you are targeted as a high risk genotype/lifestyle
- **Cyberwar destroys U.S. financial institutions**
  - …because U.S. lacks ability to construct IT infrastructure

renci

Understanding the Future

- **Some rules of thumb**
  - in the *near term*, we *overestimate* change
  - in the *long term*, we *underestimate* changes
- **Outside their field of expertise**
  - experts are often better at predictions
    - the contra-Delphi effect
- ***Inventing the future* is more successful**
  - recognize exponentials
    - quantitative change brings qualitative change
  - recognize multidisciplinary coupling
- **Technological and social change**
  - different rates with differing consequences
- **Packaging the story does matter**

renci

The Coming of Consumer Parallelism

- **Technology trends**
  - slower rise in clock rates
  - multicore processors
    - IBM Power5/6 and SUN UltraSPARC IV
    - Intel Core and AMD Opteron
    - quad core and beyond are coming
  - reduced power consumption
    - laptop and mobile market drivers
  - greater I/O and memory integration
    - PCI Express, Infiniband, …
  - Justin Ratter (Intel)
    - "100's of cores on a chip in 2015"
- **Moore's law isn't a birthright**
  - CMOS scaling issues are now a challenge
    - power, junction size, fabrication line costs, …

AMD 64 Opteron

intel Core Duo inside

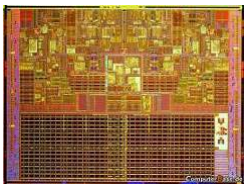

renci

We're Still Trying to Get There …

February 1994

PETAFLOPS II

2nd Conference on Enabling Technologies for Peta(fl)ops Computing

February 15-19, 1999

Doubletree Hotel — Santa Barbara, CA

Conference Chair — Paul Messina, Caltech
Program Chair — Thomas Sterling, Caltech/JPL
Steering Committee Chair — Paul H. Smith, DOE

Sequential **Terascale** **Petascale**

renci
renaissance computing institute

Los Alamos Roadrunner

- **RFP let on May 10**
  - baseline 60 TF Linux cluster
  - accelerator augmentation option
    - ~1 petaflop peak target via board/node
- **Multiple accelerator options**
  - IBM Cell, Clearspeed, …
- **Clearspeed example**
  - 50 GF at 25 watts
  - 1 GB onboard memory
  - PCI-X interface

Los Alamos NATIONAL LABORATORY EST.1943

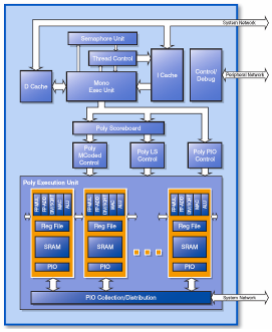

ClearSpeed

renci
renaissance computing institute

Petascale Reliability

- **Facing the issues**
  - ASCI Q boot time is ~8 hours
    - not far from the system MTTF
  - application checkpoint frequency
  - MTTF $1/\lambda = 1\text{-}r$

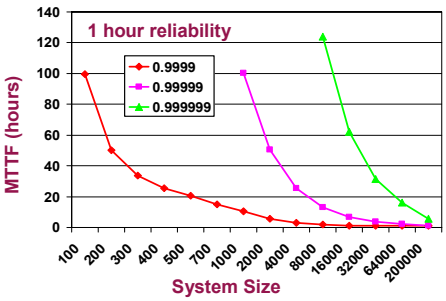

- **A few assumptions**
  - assume independent component failures
    - an optimistic and not realistic assumption
  - *N* is the number of processors
  - *r* is probability a component operates for 1 hour
  - *R* is probability the system operates for 1 hour
- **Then** $R = r^N$ **or** $R \approx \frac{1}{e^{\lambda N}}$ **for large *N***

Power Consumption and Reliability

- **Power has many implications**
  - cost, for sure
  - but also physical system size and reliability
  - Blue Gene/L uses low power processors
    - that's no accident

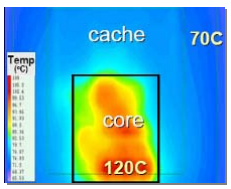

David Barkai, Intel

- **Scaling also affects power and RAS**
  - at ~50 nm feature size (without new approaches)
    - static power (leakage) is comparable to dynamic (switching) power
  - leakage *increases dramatically* with operating temperature
  - SRAM soft error rate (SER) increased 30X (Intel)
    - when moving from 0.25 to 0.18 micron geometry and from 2 to 1.6V
  - ECC does not catch all errors (Compaq)
    - perhaps 10 percent uncaught
  - worse, cheap systems have no ECC

Dell Cluster: Top500 Benchmarking

- **Configuration**
  - 512 3.6 GHz Xeon dual processor nodes
  - Infiniband interconnect
- **UC Berkeley/Crossbow motes**
  - temperature measurements
    - air inlets and outlets
- **Multiple benchmarks (primarily HPL)**

Mote Sensor Locations

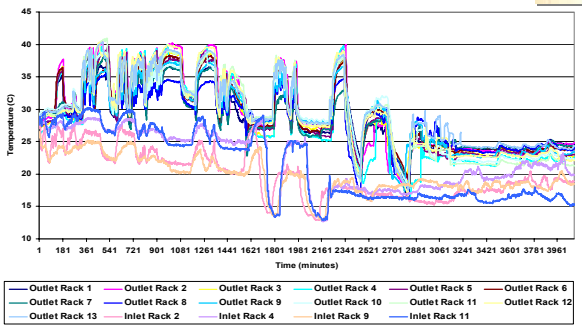

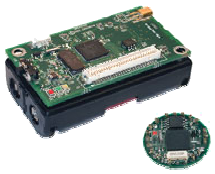

Shobana Ravi, Duke

renci

Petascale and Beyond

- **CMOS will support ~10 years of device progress**
  - Moore's law will likely hold at some level through 2016
- **Current designs for petascale systems**
  - IBM's BG/P and PERCS, Cray's Cascade
  - SUN's Hero and SGI's Ultraviolet
  - among the last large-scale extrapolations of current parallel designs
- **Architecture challenges are rising**
  - available investment is declining
- **Petascale and beyond**
  - power and cooling, packaging density
  - exposed concurrency in programming models
  - scalability of system software
  - application tailoring and optimization
  - programmability and ease of use

renci

Changing Design Challenges

- **Dramatic growth in computing base**
  - then (1960s-1970s)
    - thousands of users
    - field repair of design errors
      - extra revenue ☺
  - and now
    - hundreds of millions of users
    - massive product recall of design errors
      - life and death corporate issue ☹

- **Where are today's intellectual giants?**
  - Eckert, Mauchly, Cocke, Tomasulo, Slotnick, …
- **Designing user systems, not components**

Ideas from Bob Colwell/Intel

renci

Telematics, Biomedicine and Picosatellites

- **Smart cars**
  - Bluetooth, navigation, digital audio, active cruise control
  - rain sensitive wipers, light sensitive mirrors, tire pressure
  - parking assistance, EFI and variable timing
  - road tracking and drowsy warning
  - on-board vehicle diagnostics
    - OBD II standard
    - Controller Area Network (CAN)
  - inter-vehicle communication
    - *ad hoc* and interest-based
- **Biomedicine**
  - technology drivers
    - wireless communications and circuit density
  - capsule endoscopy, diabetes, pacemakers, drug infusers
- **Picosatellite**
  - ~$10K-$40K construction cost
  - 10 cm cube (one liter) to 10x10x30 cm
    - industry standard PC-104 boards
  - ~$50K launch cost
    - secondary payload on commercial launcher

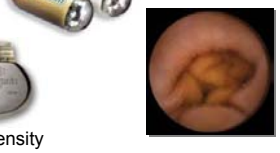

renci

We Are Becoming A "Mature Industry"

- **Examples**
  - airplanes, cars, railroads, …
  - toothpaste, breakfast cereal, …
- **Common characteristics**
  - economics limit radical change
  - installed base shapes evolution
  - primary competition on price
  - infrastructure compatibility
- **Or in technological terms**
  - declining returns on investment
  - complex component interactions
  - large volume commodity shipments
  - targeted niche applications

renci

Influence: The Computing Food Chain

- **Our late term influence has waned**
  - overall market growth/technology maturation
  - incrementalism generally flourishes
- **Only big events can kill dinosaurs**
  - diversified, evolutionarily successful ecosystem
- **We still have influence**
  - new foundational technologies
  - revolutionary approaches
  - disruptive product niches

  but in different ways …

- **Hence, the big question is …**
  - evolution or revolution

renci

Disruptive Technology

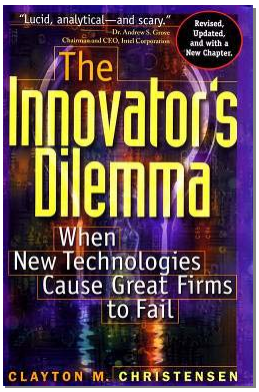

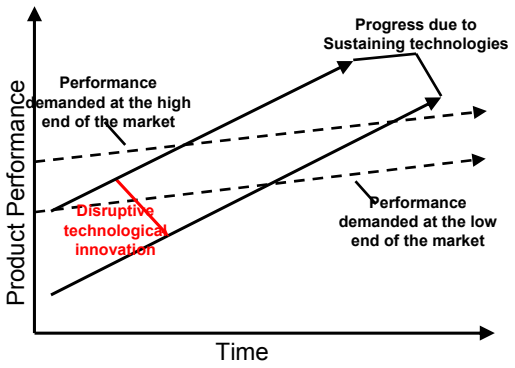

- **Initially less profitable**
- **Mass market acceptance for full value**
- **Cheaper, smaller, simpler, …**

renci

The Big Questions

- **Life and nature**
  - structures, processes and interactions

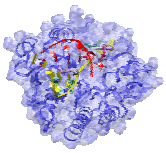

- **Matter and universe**
  - origins, structure, manipulation and futures
  - interactions, systems and context
- **Humanity**
  - creativity, socialization and community

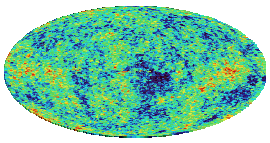

- **Answering big questions requires**
  - boldness to engage opportunities
  - expandable approaches
  - world-leading infrastructure
  - (often) interdisciplinary partnerships

renci

Big Science Visions Are Common

- **Multilevel biological modeling**
  - from molecules and structures to organisms and ecologies
    - petascale systems and beyond
- **Distributed, virtual astronomy**
  - real-time data analysis and multi-modal data fusion from distributed archives
- **Personalized, *in situ* medicine**
  - drug design to [illegible] micro-transfusers
- **High-energy physics/cosmology fusion**
  - dark matter, the standard model, and the theory of everything
- **Integrated climate change and urban/social planning**
  - multidisciplinary data fusion, modeling, and analysis

Big Questions to Get Big Answers

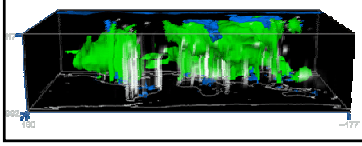

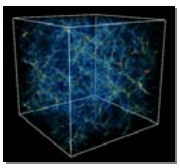

renci
renaissance computing institute

Integrated, Predictive Biology

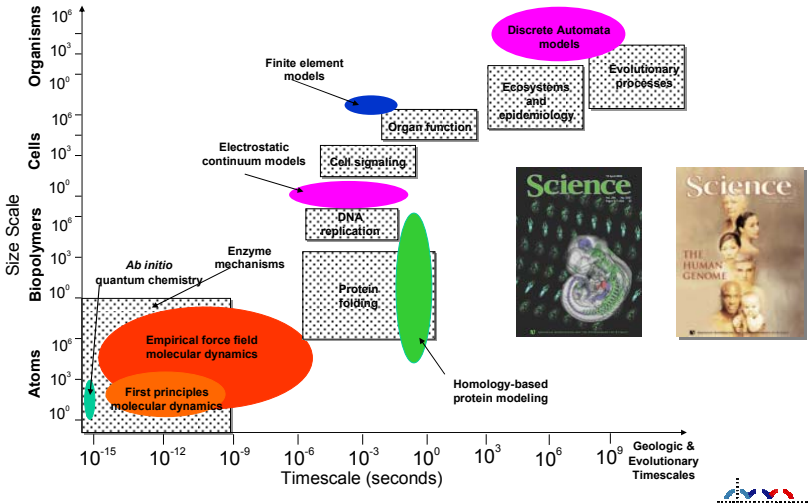

renci
renaissance computing institute

Source: DOE Genomes to Life

Unanswered Physical Questions

- **The Standard Model**
  - quantum theory
    - electroweak and strong forces
- **Gravity and relativity**
  - no "Grand Unified Theory"
    - gravity integration and rationale for mass
      - search for the Higgs boson
- **Dark matter and dark energy**
  - most of the universe's mass is "invisible"
    - mass candidates: MACHOS (baryonic), WIMPS, …
  - expansion seems to be accelerating
    - vacuum energy (cosmological constant)
    - dynamic field (quintessence)
- **Recent experimental data**
  - Wilkinson Microwave Anisotropy Probe (WMAP)
    - universe is 13,400 ± 300 million years old and flat

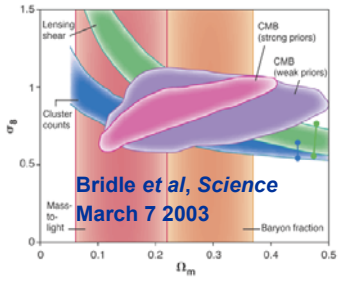

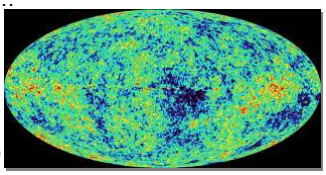

renci

Big Means What?

- **Big projects are getting smaller!**
  - remember the effects of inflation $P = C(1+R)^N$
- **We need to think bigger!**
  - what is a >$100M project?

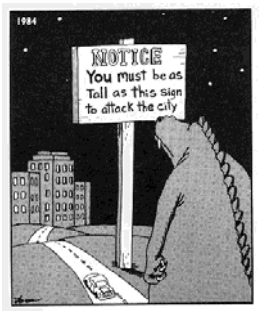

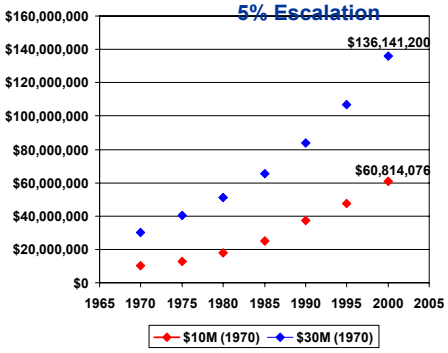

renci

The Public Perception …

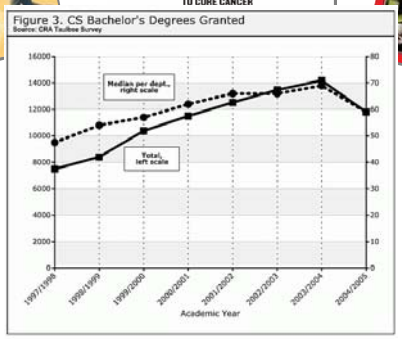

U.S. Trade {Im}Balance

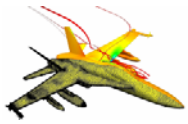

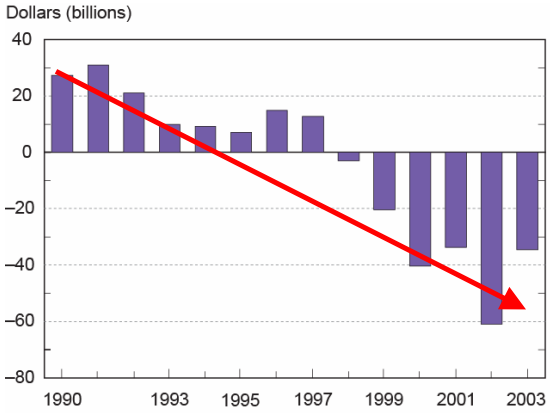

NOTE: Includes aerospace, pharmaceuticals, office and computing equipment, communications equipment, and scientific instruments.

Source: NSF Science and Engineering Indicators, 2006

The Ecosystem Really Matters

- **Local change at a single agency can**
  - upset national agendas
  - have long-term repercussions
- **Universities and Federal R&D agencies must *make coordinated, fundamental, and structural changes* that affirm the integral role of computational science**
  - the most important problems are multidisciplinary, multi-agency, multi-sector, and collaborative

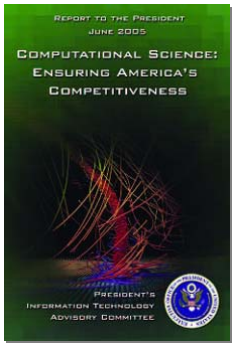

Wetland Restoration

renci
renaissance computing institute

American Competitiveness Initiative (ACI)

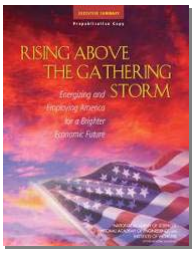

- **Innovation**
  - double NSF, DOE SC and NIST
  - make permanent the R&E tax credit
- **Education**
  - 70,000 new teachers
  - alternate teacher certification
- **Workforce/immigration**
  - expand worker training programs
  - flexible H-1B caps, reform visa issues

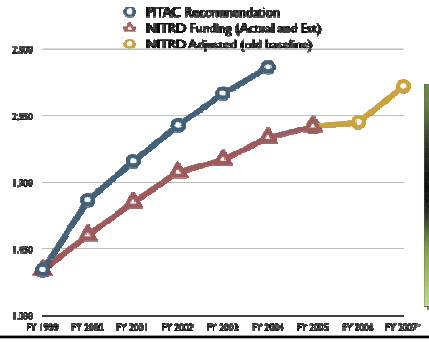

renci
renaissance computing institute

Community Computing Consortium (CCC)

- **CCC vision**
  - community building and vision setting
  - research agenda and funding advocacy
- **Other disciplinary examples**
  - astronomy decadal survey
  - physics priority list
- **CCC leadership**
  - Computing Research Association (CRA)
- **Proposal leaders**
  - Dan Reed, CRA Board Chair
  - Randy Bryant, CMU
  - Susan Graham, UC Berkeley
  - Ed Lazowska, Washington
  - Dick Karp, UC Berkeley
  - Ken Kennedy, Rice
  - Peter Lee, CMU
  - Wim Sweldens, Lucent
  - Jeff Vitter, Purdue

COMPUTING RESEARCH ASSOCIATION
CRA

renci
renaissance computing institute

What's the Moral?

- **Set some priorities**
  - no priorities means no vision
  - no vision means no intellectual commitment
- **Choose some directions**
  - technology and applications
    - identify a driving problem
- **Think at appropriate scales**
  - financial and temporal
    - you must be tall enough to attack the city ☺
- **Take some bigger risks**
  - technical and political
  - most innovative projects fail
    - at least by narrow technical measures
    - and that's just fine!

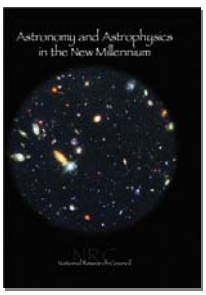

renci
renaissance computing institute