Computing at a Cross-Roads: Big Data, Big Compute, and the Long Tail

William Gropp

www.cs.illinois.edu/~wgropp

What this talk is…

- A request for help, in keeping with the topics of this meeting
- The US NSF has asked for a "study on Future Directions for NSF Advanced Computing Infrastructure to Support US Science in 2017-2022"
- The study is being conducted by the National Academies
- See Project Website: http://www8.nationalacademies.org/cp/projectview.aspx?key=49628
- Responsible Staff Officer: Jon Eisenberg

NATIONAL ACADEMY OF SCIENCES

1863–2013 • Celebrating 150 Years of Service to the Nation

Chartered by Congress to:

> Whenever called upon by any department of the Government, investigate, examine, experiment, and report upon any subject of science or art.

Project Scope

- Examine **anticipated priorities and associated tradeoffs** for advanced computing in support of NSF-sponsored science and engineering research considering:
- The **contribution of high end computing to U.S. leadership and competiveness** in basic science and engineering and **the role that NSF should play in sustaining this leadership**
- **Expected future national-scale computing needs**: high-end requirements, those arising from the full range of basic science and engineering research supported by NSF, as well as the computing infrastructure needed to support advances in both modeling, simulation and data analysis
- **Complementarities and tradeoffs** that arise among investments in supporting advanced computing ecosystems; software, data, communications
- The **range of operational models** for delivering computational infrastructure, for basic science and engineering research, and the role of NSF support in these various models
- **Expected technical challenges** to affordably delivering the capabilities needed for world-leading scientific and engineering research

Study Committee

William Gropp, UIUC (co-chair)
Robert Harrison, Stony Brook/Brookhaven (co-chair)
Mark R. Abbott, Oregon State
David Arnett, Univ. of Arizona
Robert Grossman, Univ. of Chicago/Open Data Group
Peter Kogge, Notre Dame
Padma Raghavan, Penn. State
Daniel A. Reed, Univ. of Iowa
Valerie Taylor, Texas A&M
Katherine Yelick, UC Berkeley/LBNL
Jon Eisenberg, National Academies study director

Timeline and Reports

Interim report (Late Summer 2014) to identify key issues and discuss potential options. 2017-2020 time frame

Final report (2015) to include a framework for future decision-making about NSF's advanced computing strategy and programs. 2017-2022 time frame.

- how to prioritize needs and investments and how to balance competing demands for cyberinfrastructure investments
- approach: identifying issues, explicating options, and articulating tradeoffs and general recommendations

No recommendations concerning the level of federal funding for computing infrastructure

International Perspective

- We are very interested in how others are addressing these questions, including:
  - Relative investment in hardware, software, algorithms
  - Planning horizons
  - Facility operations and management
  - Are Data intensive facilities fundamentally different from Compute intensive? Is one a subset of the other?

How to Provide Comments

- Please send suggestions of information sources (people, reports, etc.) to the study director, jeisenbe@nas.edu
- Submit your comments on study topic via National Academies Current Projects system, http://tiny.cc/inputacistudy

- For more on project, follow www.cstb.org

Disclaimer for the Rest of the Talk

- The interim report is not yet available, so I can't discuss any details or content
- The remainder of this talk is my own, not the study committee's, and reflect my own interests and biases
- However, I encourage you to consider all issues and questions presented and provide comments and feedback to the study committee, including the three members at this conference

Some Things that I'd Like To Know

- How to create advanced computing infrastructure that enables integrated discovery involving experiments, observations, analysis, theory, and simulation?
- How would you define the challenges in providing future advanced computing, including:
  - The growing demand for both capacity and capability computing
  - Growth of data intensive computing, supporting scientific workflows and multidisciplinary research
- What are the challenges in software and algorithms for next generation hardware?
  - Especially ones other than the usual suspects, or with quantitative, rather than qualitative, data

Some Things that I'd Like To Know

- Comprehensive data on capacity needs by research area
- Comprehensive data on capability needs and impact by research area
- Comments on the role of private industry in providing cyberinfrastructure, including support for different operational models
- Definition of computing needs by research area, integrated over NSF
- What are the most important operational issues for computing centers, particularly in the next ten years?

Planning in Astronomy

- Something that the CSE Community is not ready for but should be aiming for: The "Decadal Survey"
- Massive, community-wide effort
  - 450 white papers
  - 17 town halls
  - 27 panel meetings
  - Study committee had 6 meetings + >100 telecons
- Does recommend funding levels
  - Balanced program between large and small scale programs

Current Community and Investments

- The following two slides provided by NSF; provide historical trends on use of facilities
- Note the long tail – many users, groups, domains
- Note still compute-centric
- Note large Blue Waters users are > 1M XD SUs/year
  - Blue Waters is not part of XSEDE
  - Performance over 13 PF/peak; over 1PF sustained on wide variety of applications including all I/O

Some Data on the Use of NSF Computing

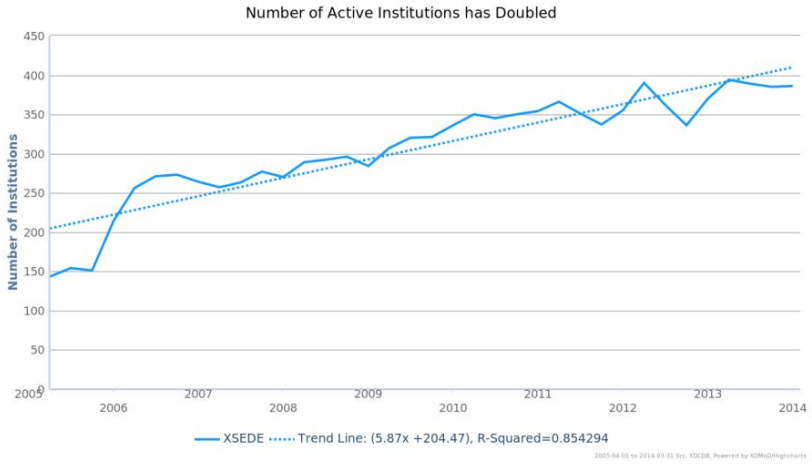

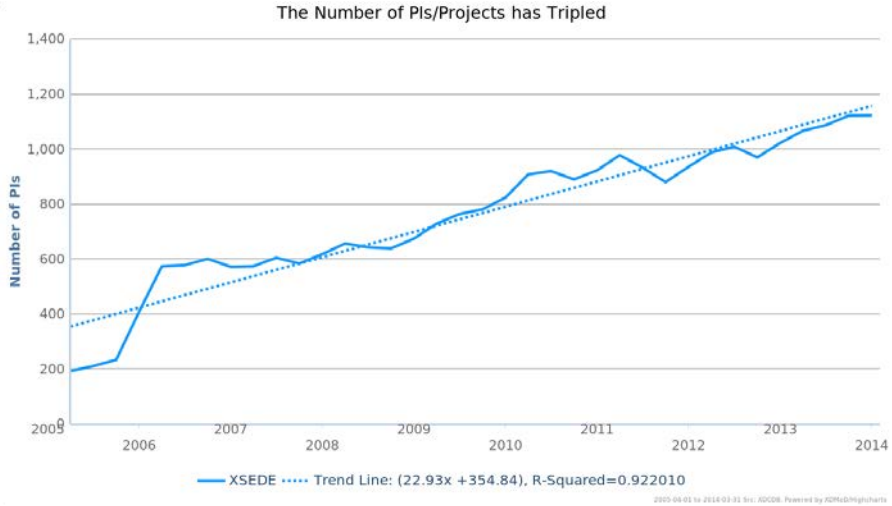

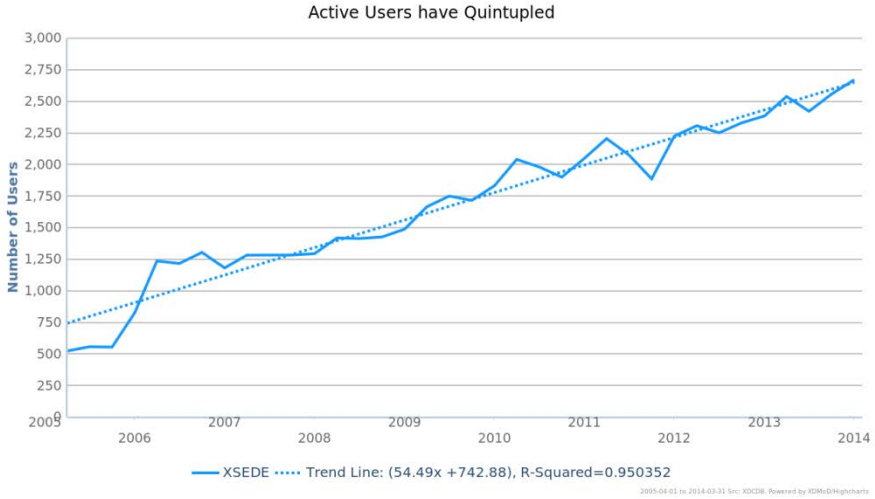

PARALLEL@ILLINOIS

Supply Side View

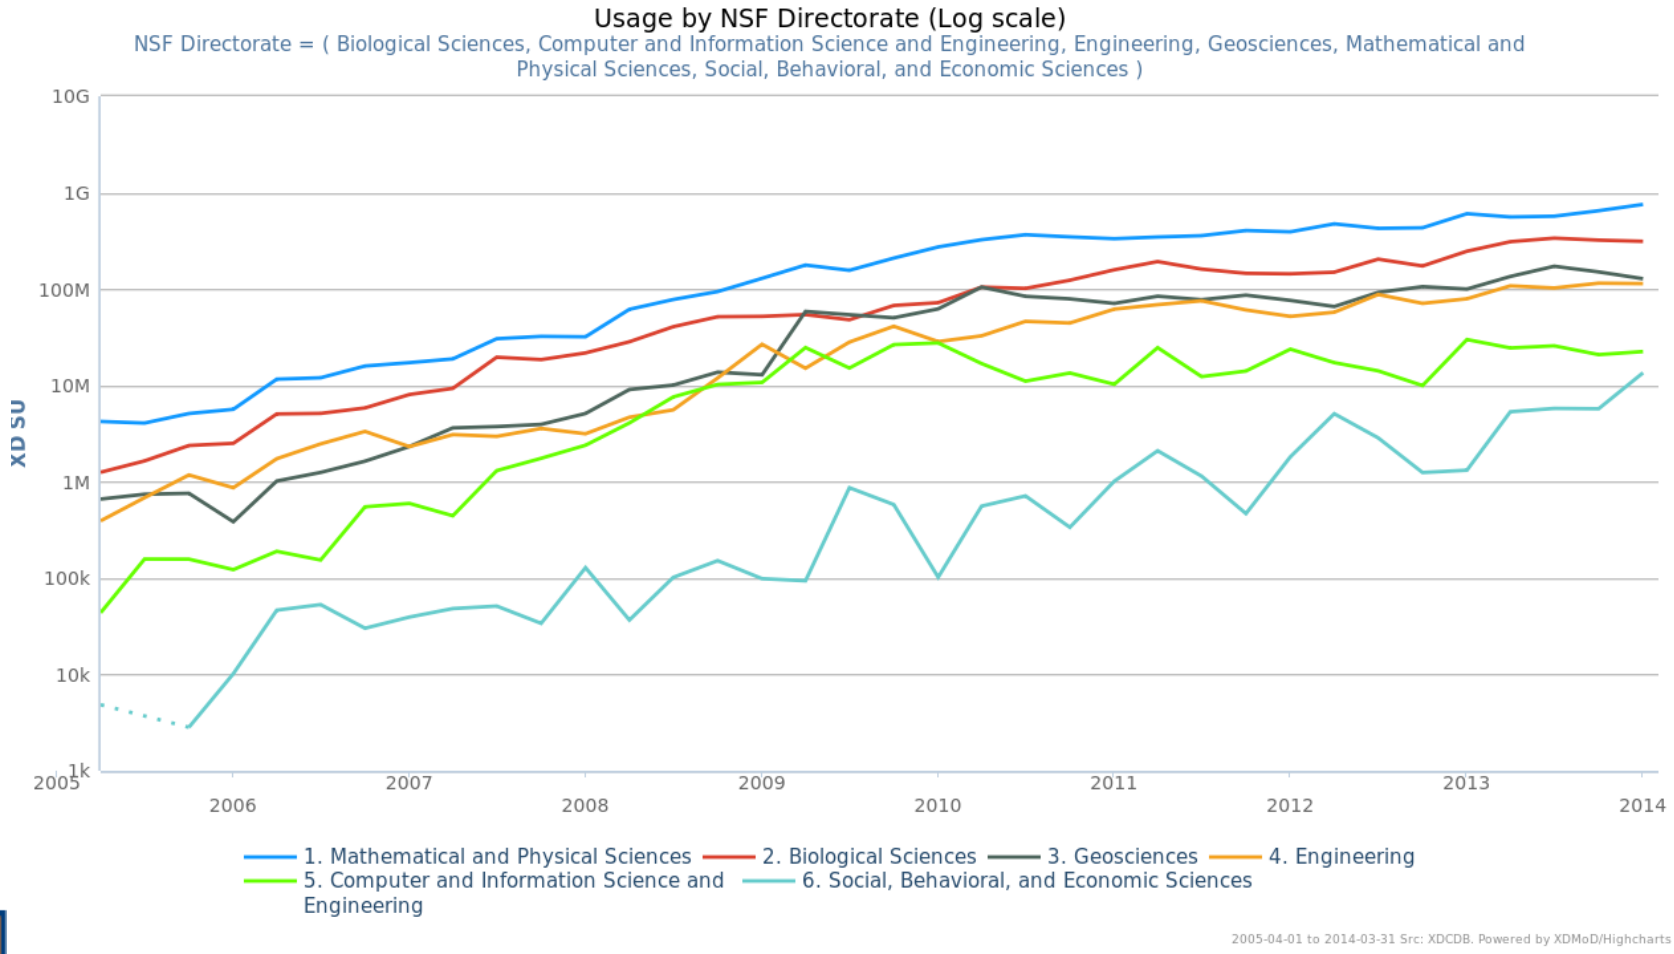

PARALLEL@ILLINOIS

Demand Side View

- Surprisingly hard to get data on the demand
  - Next slide: XSEDE as of 2012
  - DOE INCITE has provided information on the total size of requests
- Obvious caveat: Demand expands to exceed available resources
  - How much is essential for progress?
- Less obvious caveat: Requests are not centralized. Not every NSF funded researcher even knows about XSEDE. Allocation process and overhead selects some users out before submission. Potentially particularly strong for the "long tail"

Available-Requested-Allocated SUs (Data to 2012)

- (slide pending release)

Thanks to John Towns

US NITRD

- Networking and Information Technology Research and Development Program
- https://www.nitrd.gov/Index.aspx
- Cross-government coordination of NIT R&D
- Caveats: Budget data is self-reported; allocation by category imprecise; DOE split in inconsistent ways; no explicit "big data" (but a "Big Data Senior Steering Group")

US NITRD Areas

- CSIA – Cyber Security and Information Assurance
- HCI&IM – Human Computer Interaction and Information Management
- HCSS – High Confidence Software and Systems
- HEC I&A – High End Computing Infrastructure and Applications
- HEC R&D – HEC Research and Development
- LSN – Large Scale Networking
- SDP – Software Design and Productivity
- SEW – Social, Economic, and Workforce

Why Look at Budget Numbers?

- Information on agency priorities, relative levels of investment within and between agencies
- Information on absolute investments
- Caveats: Categories imprecise, often trailing indicators, definitions vary (e.g., for NIH, any cluster apparently qualifies as HEC)

US Government Wide

- See http://itdashboard.nitrd.gov/Dashboard.aspx

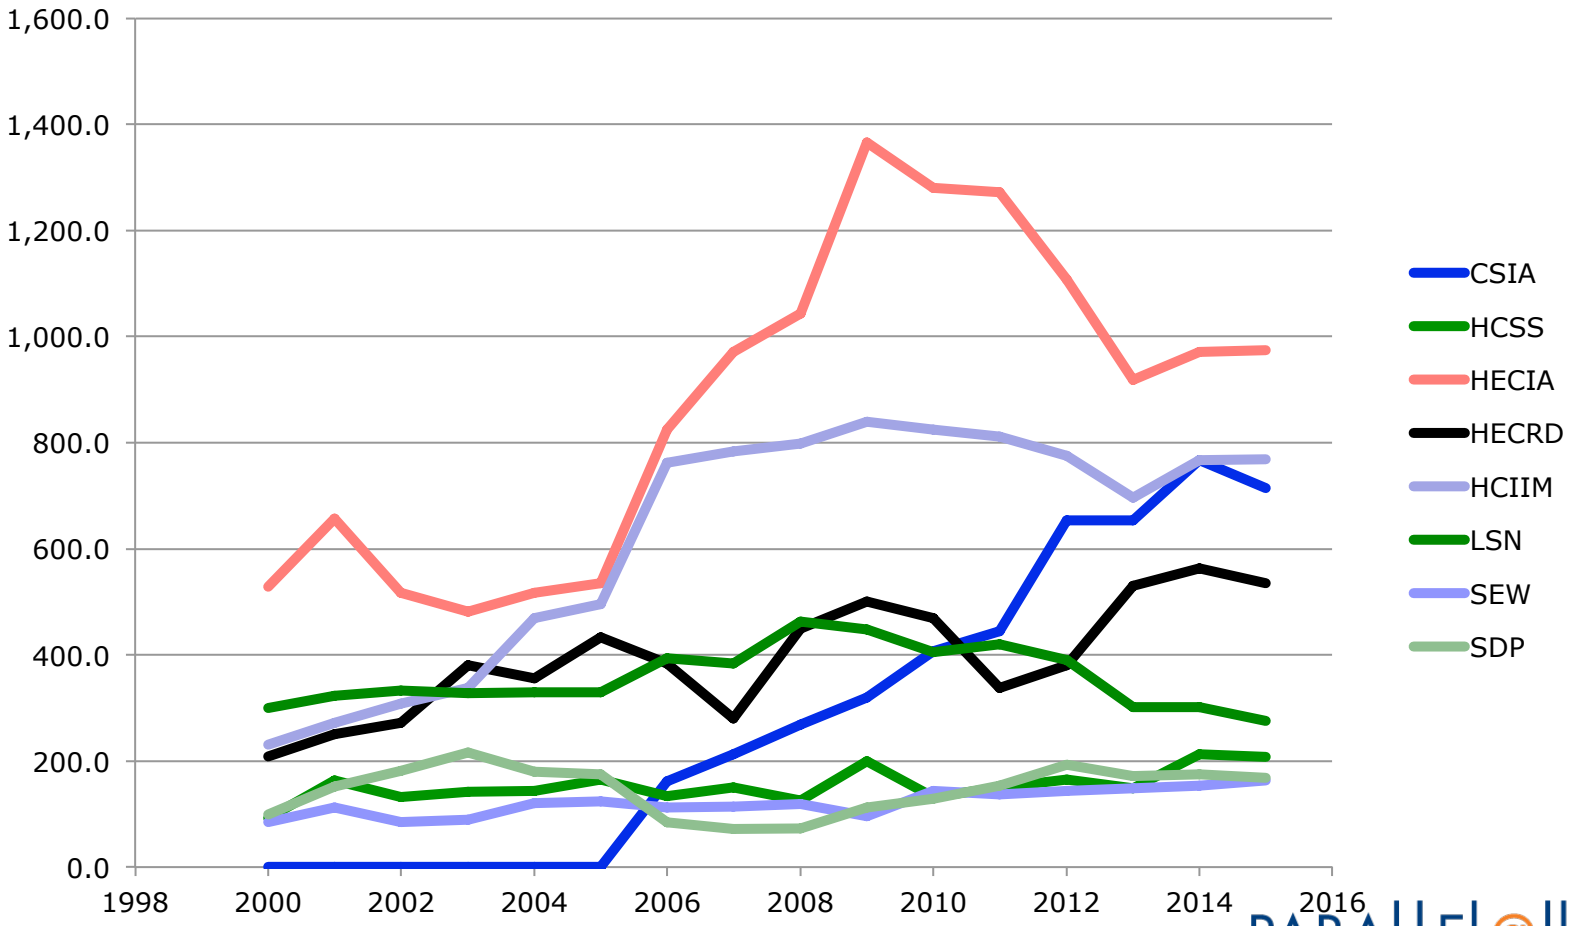

PARALLEL@ILLINOIS

US (all agencies) NITRD (% of total)

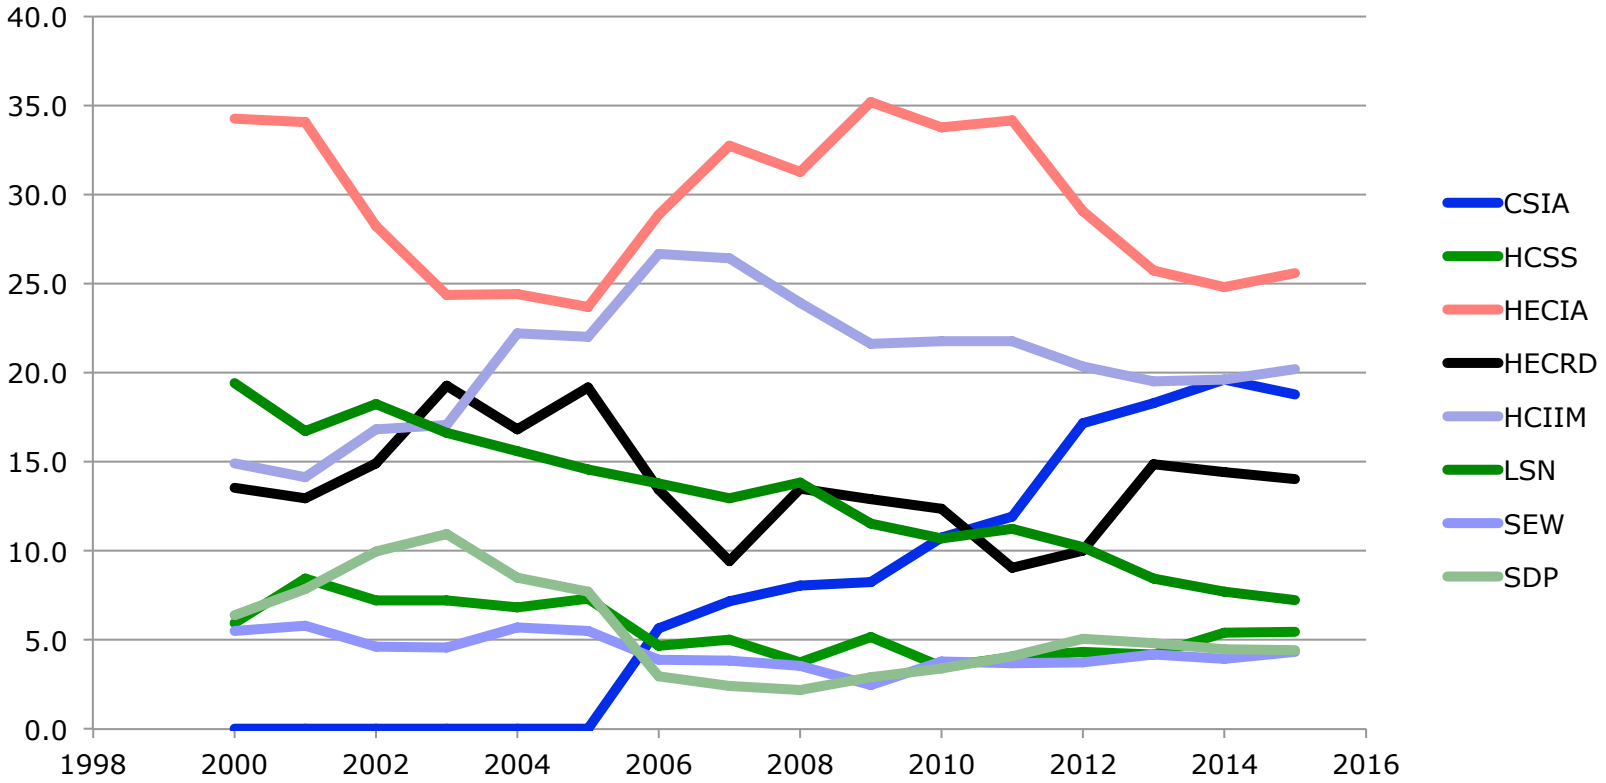

US NSF NITRD

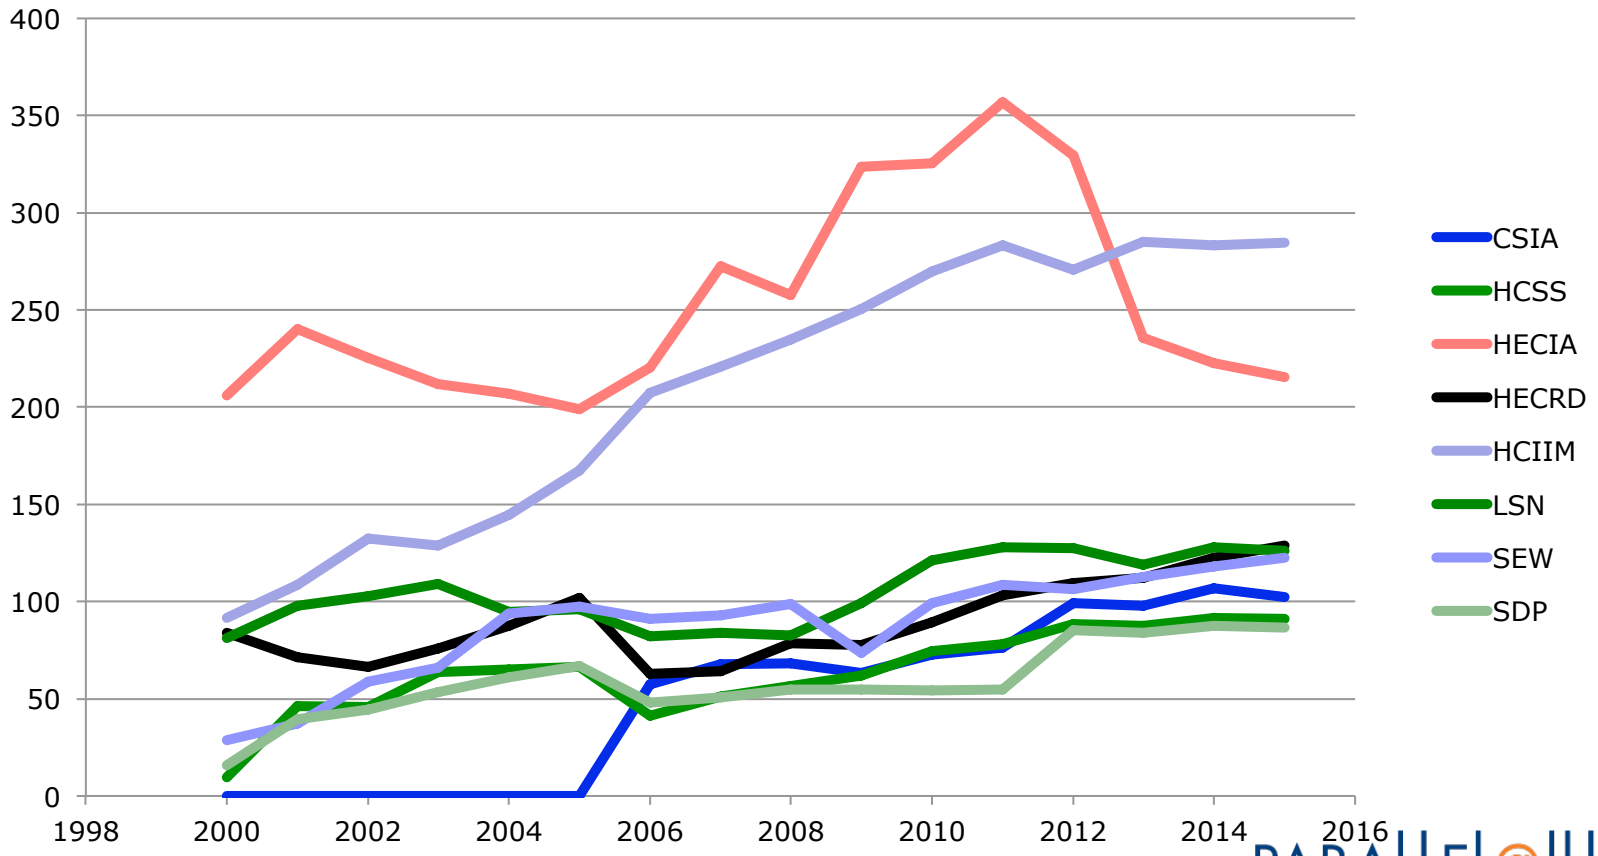

PARALLEL@ILLINOIS

US NSF NITRD (% of total)

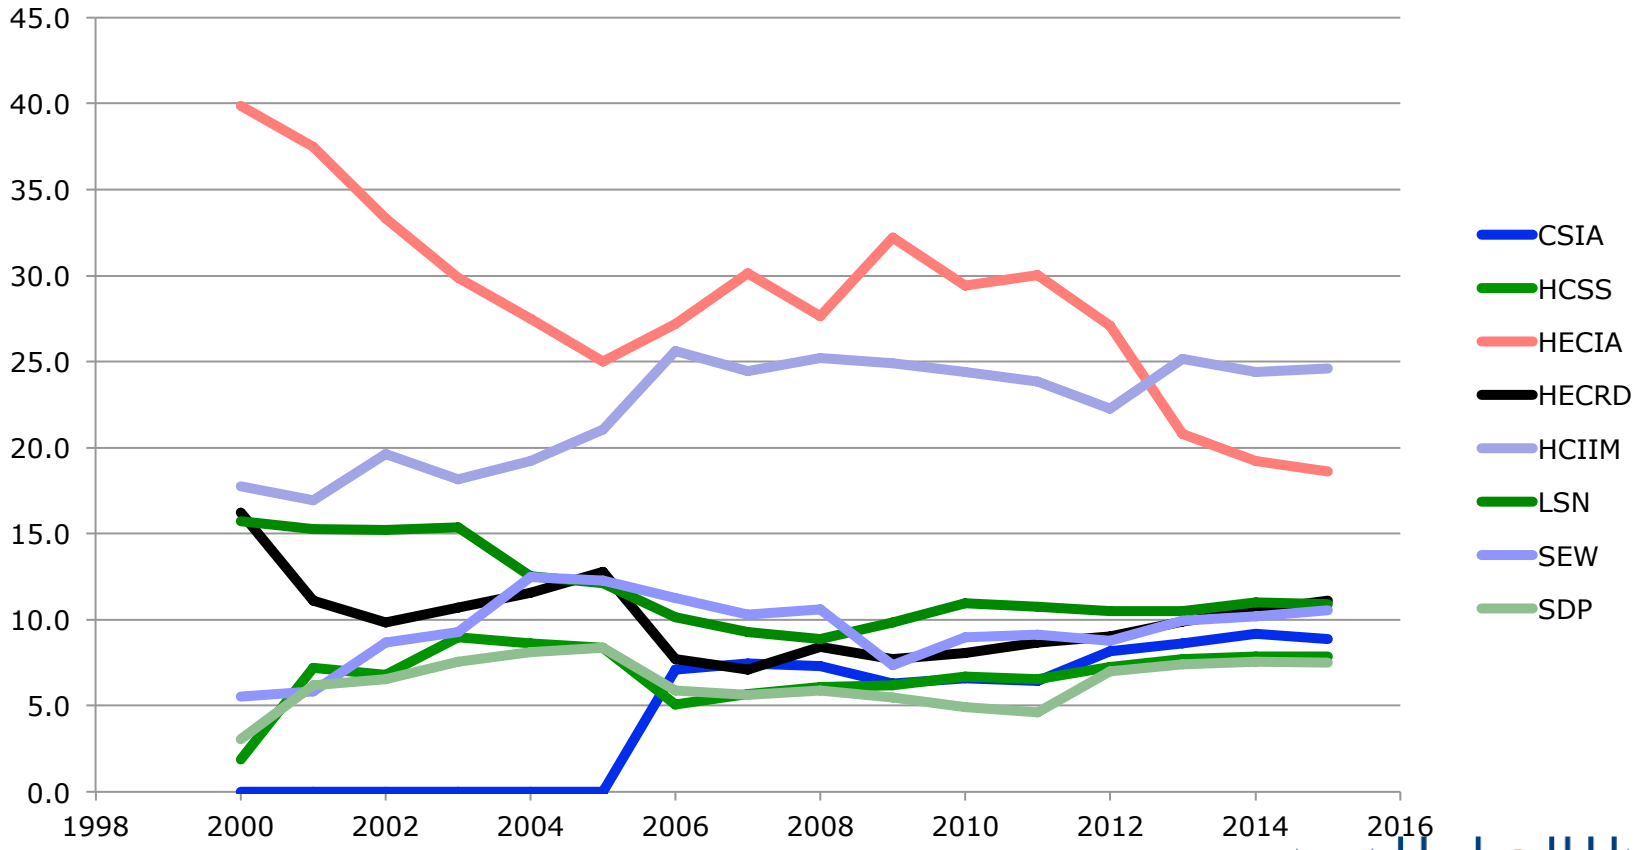

NSF Budget as a % of Total

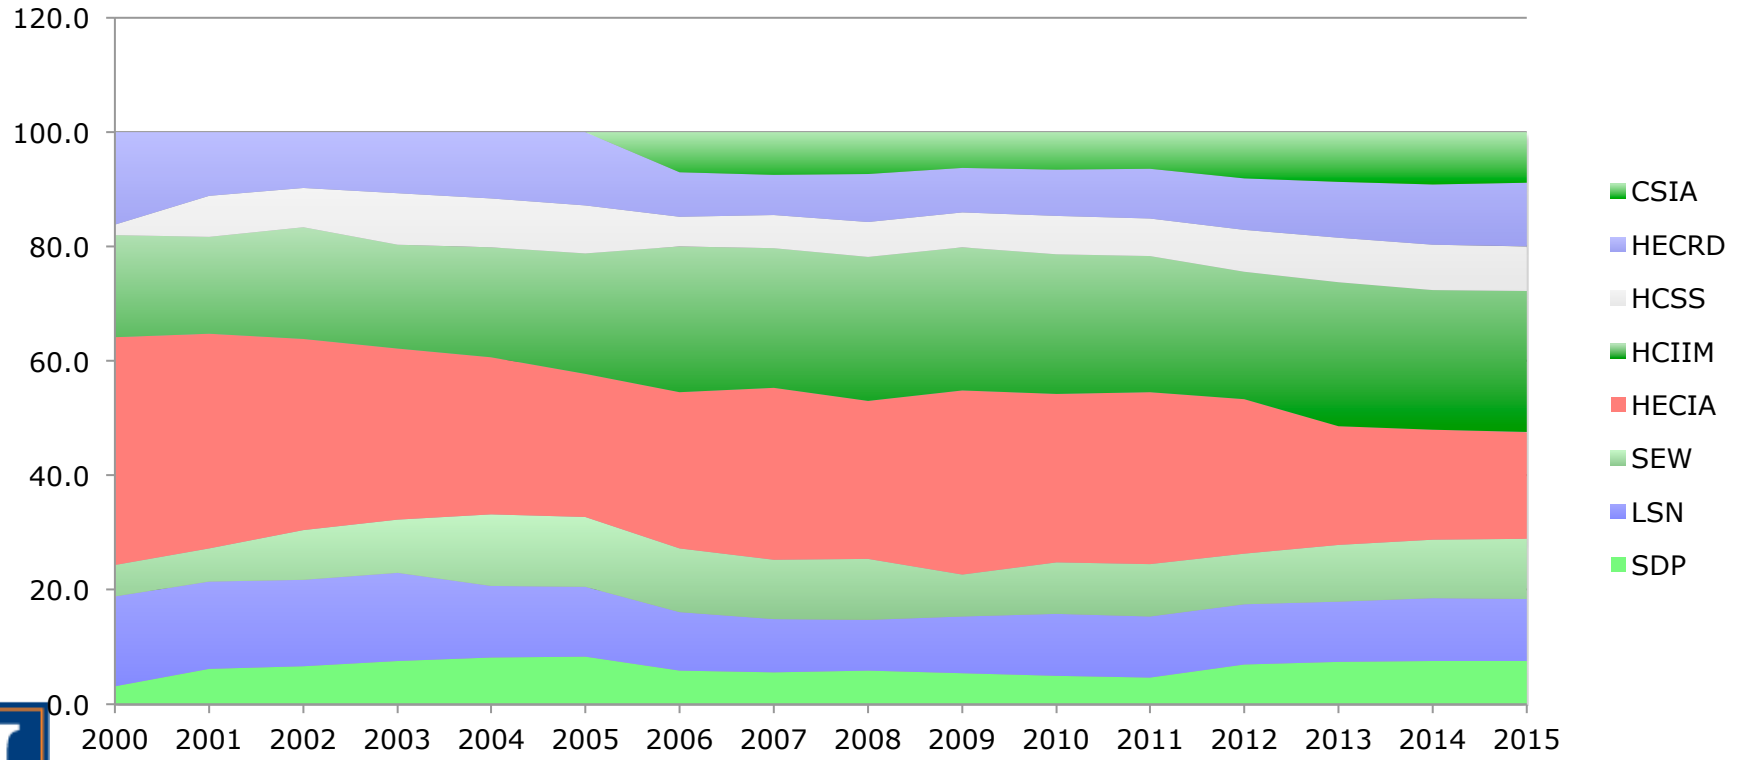

Some Observations

- CSIA (cybersecurity) apparently carved out of HEC R&D, HCSS
- HEC Infrastructure has declined by more than 50% as % of total
- HCIIM (HCI and information management) largest budget category since 2012

Is it Time to Change the Model for How Resources are Provided?

- Computing (and data) as major equipment
  - ♦ Acquire and operate as other major NSF facilities
- Computing and data as a commodity service
  - ♦ Contract out to private sector providers
- Computing and data as major equipment, operated as a commodity
- Other?

What about Clouds?

- What is it about a "cloud" that is important? Or, what do you mean when you talk about "cloud"?
  - Flexible service model: time available on demand with low administrative overhead
    - Is there any guarantee of availability?
      - Guarantees have cost – typically extra resources to guarantee that peak demands can be met
    - This is "On-demand self-service" in the NIST definition, plus "Rapid elasticity"
  - Shared resource with cost advantage
    - It has to cost less?
    - This is "Resource pooling" and "Measured service"
  - Goes without saying: accessible
    - This is "Broad network access" in the NIST definition
  - Virtualized resources
    - A critical part of the definition for some, irrelevant for others

What about Clouds?

- How should costs for different resource models be compared?
  - ♦ Common to charge per virtual CPU
  - ♦ Data and networking charged separately
  - ♦ Support
    - Not everyone needs support
    - Some do and don't know it
    - Would it make more sense to unbundle these costs?
      - – Would investigators include those costs in their budget? (Ha!)

Background Calculation: Cost of a Leadership-sized system on the Cloud

- Amazon EC2 instance:
  - 22,000 nodes with 60 GiBytes and 32 cores/node:
    - $27,939,694/month on demand
    - $6,028154.76/month 3 year heavy reserved
  - 4224 nodes with GPU
    - $2,083,614.43/month on demand
    - $203,464.81/month 3 year heavy
  - Not really comparable to center facility
    - No data in/out of system
    - Far less data support than, e.g., Blue Waters (26 PB disk; 320 PB tape, >1TB/sec to disk)
    - No high-speed, low latency interconnect
    - No expert user support
    - Basic Linux (SUSE Enterprise adds about $200k/month)
  - About $373,897,080 over 5 years (heavy reserved)
    - $1.8 Billion on demand
  - Numbers from online tool, August 28, 2014

How to Provide Comments

- Please send suggestions of information sources (people, reports, etc.) to the study director, jeisenbe@nas.edu
- Submit your comments on study topic via National Academies Current Projects system, http://tiny.cc/inputacistudy

- For more on project, follow www.cstb.org