Return on Investment from Academic Supercomputing: SC14 Panel

Panelists:

- Greg Newby (KAUST), Moderator
- Amy Apon, Clemson University
- Nick Berente, University of Georgia
- Rudolph Eigenmann, National Science Foundation
- Susan Fratkin, Coalition for Academic Scientific Computation
- David Lifka, Cornell University
- Craig Stewart, Indiana University

New Orleans Convention Center

Friday, November 21 2015

- **ROI** == Return on Investment. An assessment of outcomes, relative to expenses.
- *Challenge*: academic outcomes are not primarily financial
- *Challenge*: outcomes of research computing are often not measured well.
- *Challenge*: supercomputing is expensive, so scrutiny is high.

Let's talk about ROI

Returns on Investment

Effects of Investment in High Performance Computing

Amy Apon
Clemson University

Two Research Studies

- Regression Analysis
- Frontier Efficiency Estimation

Study #1: Regression Analysis

- Investment in high performance computing, as measured by entries on the Top 500 list, is a predictive factor in the research competitiveness of U.S. academic institutions.

We study Carnegie Foundation institutions with "Very High" and "High" research activity – about 200 institutions

Data Acquisition

<u>Dependent variables</u>

- NSF and other federal funding summary and award information
- Publication counts
- U.S. News and World Report rankings

<u>Independent variables</u>

- Top 500 List count and rank of entries
  - Mapped from "supercomputer site" to "institution"
  - We note that entries are voluntary – the absence of an entry does not mean that an institution does not have HPC

Data from the Top 500 List

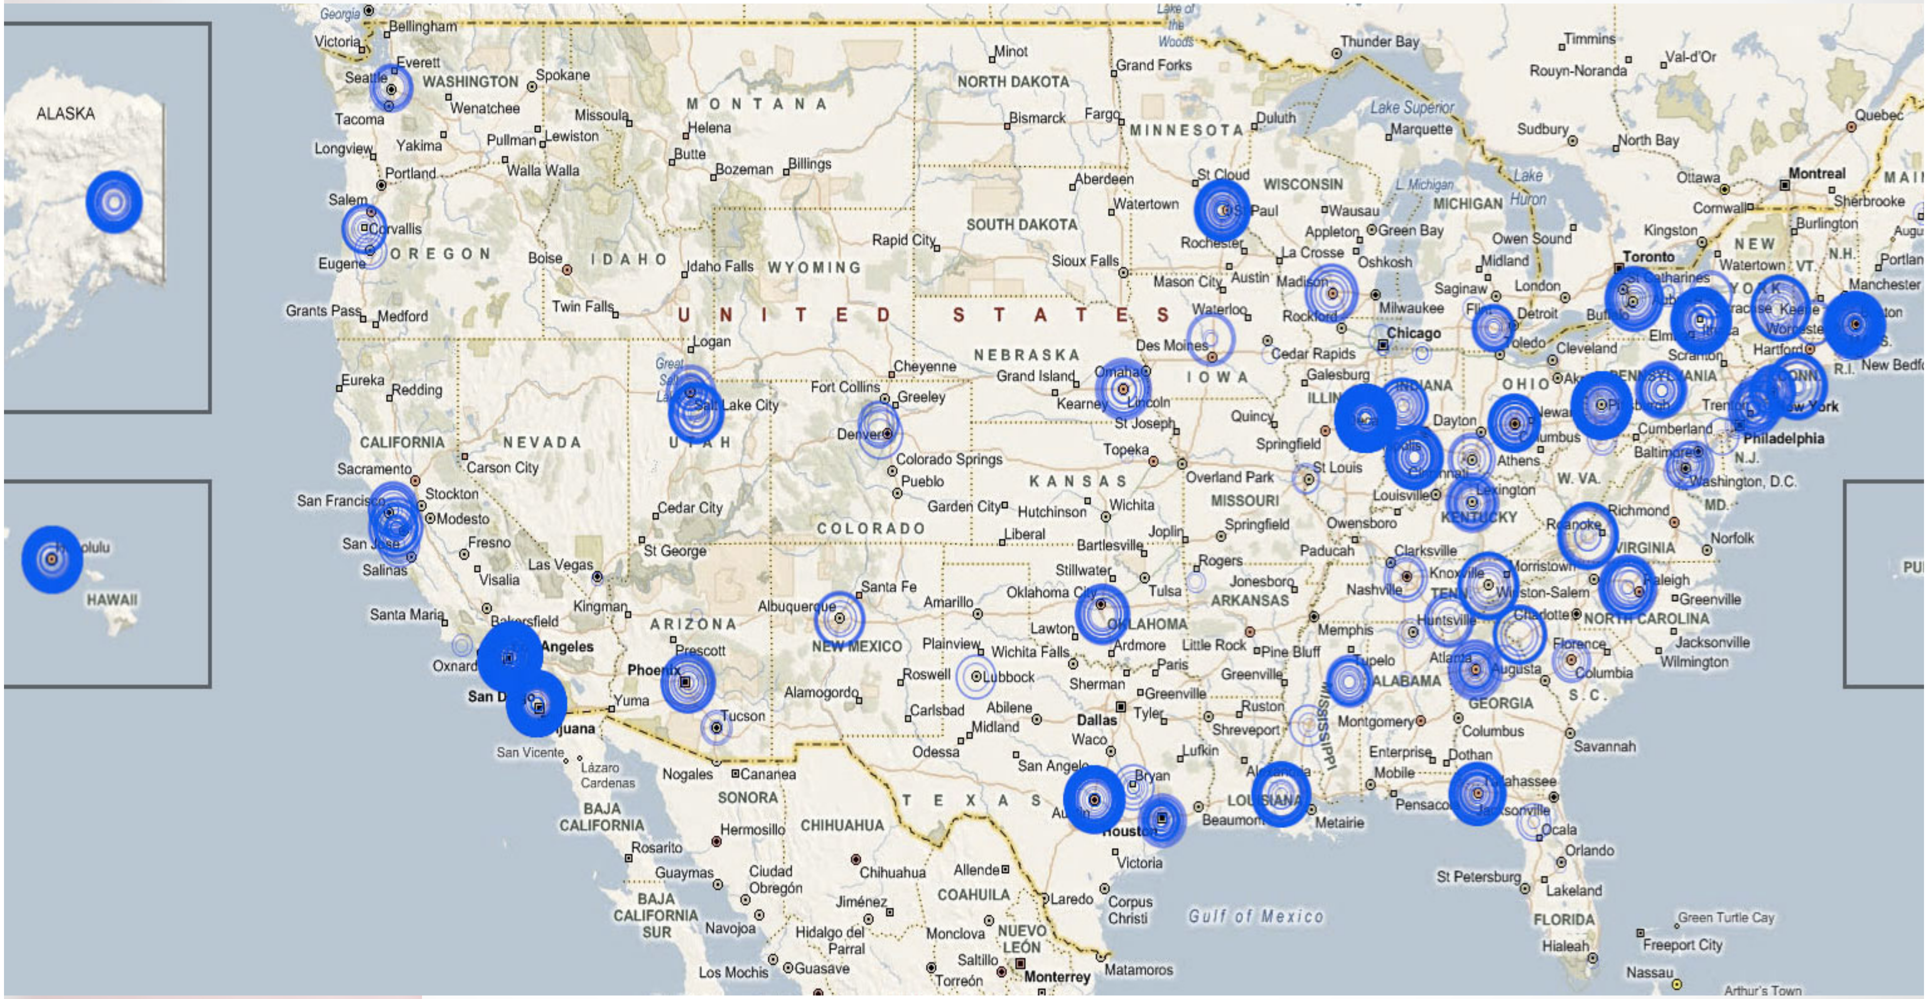

An historical record without comparison of supercomputers

Study #1: Regression Analysis

- Two Stage Least Squares (2SLS) regression is used to analyze the research-related returns to investment in HPC
- We model two relationships
  - **Model 1**: NSF Funding as a function of contemporaneous and lagged Appearance (APP) on the Top 500 List Count and Publication Count (PuC), and
  - **Model 2**: Publication Count (PuC) as a function of contemporaneous and lagged Appearance on the Top 500 List Count (APP) and NSF Funding

Results

- HPC investment yields economically and statistically significant immediate returns in terms of new NSF funding available
  - An entry on a list (two lists a year) is associated with increased yearly NSF funding of $2.4M
    - The midpoint of estimated range of $769K-$4M, with 95% confidence level

Results

- HPC investment yields economically and statistically significant immediate returns in terms of increased academic publications
  - An entry on a list is associated with average increased publications of about 60 in a year
    - The midpoint of estimated range of 19-100, with a 95% confidence level

Results

- HPC investments suffer from fast depreciation over a 2 year horizon
- Analysis on the rank of the system shows that capability has a less strong but still positive impact to competitiveness

Consistent investments in HPC, even at modest levels, are strongly correlated to research competitiveness.

Apon, et. al., "High Performance Computing Instrumentation and Research Productivity in U.S. Universities," Journal of Information Technology Impact, vol. 10, No. 2, pp 87-98, 2010.

Discussion

- Regression approach suffers from problems with endogeneity of Publication Counts (PuC) and NSF Funding.
- To correct for this, we deployed a 2SLS estimation method, with number of undergraduate Student Enrollments (SN) acting as an instrumental variable in the first stage regression for PuC (Model 1) and NSF (Model 2).
- In both cases, SN was found to be a suitable instrument for endogenous regressors.
- This also motivated a new research approach using frontier efficiency analysis, and more fine grained data

Study #2: Frontier Efficiency Estimation

- We study whether investment in high performance computing, as measured by entries on the Top 500 list, is consistent with more efficient (i.e., higher) research productivity
- We study eight different academic departments using data from the National Research Council

Frontier Analysis

We define P as the set of feasible combinations of *p* inputs and *q* outputs, also called the production set.

$$\mathcal{P} = \{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}_+^{p+q} \mid \boldsymbol{x} \text{ can produce } \boldsymbol{y}\}$$

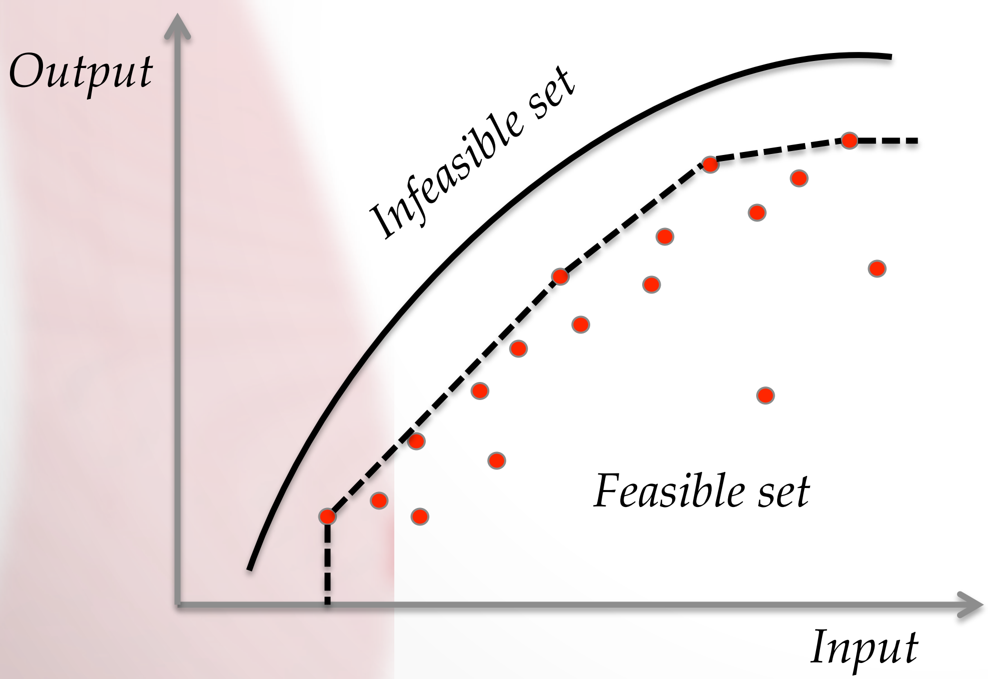

- There exists a maximum level of output on a given input (the concept of efficiency)
- The efficiency score is an estimation with regard to the true efficiency frontier
- Range: [0,1]

Results

- The availability of Top500-scale HPC resources enhances the technical efficiency of research output in Chemistry, Civil Engineering, Physics, and History,
- but not in Computer Science, Economics, nor English.
- We find mixed results in Biology.
- Our results provide a critical first step in a quantitative economic model for investments in HPC.

References

Apon, Amy W., Stanley Ahalt, Vijay Dantuluri, Constantin Gurdgiev, Moez Limayem, Linh B. Ngo, and Michael Stealey. "**High performance computing instrumentation and research productivity in US universities**." Journal of Information Technology Impact 10, no. 2 (2010): 87-98.

Apon, Amy W., Linh B. Ngo, Michael E. Payne, and Paul W. Wilson. "**Assessing the effect of high performance computing capabilities on academic research output**." Empirical Economics (2014): 1-30.

Core Hours delivered by HPC systems

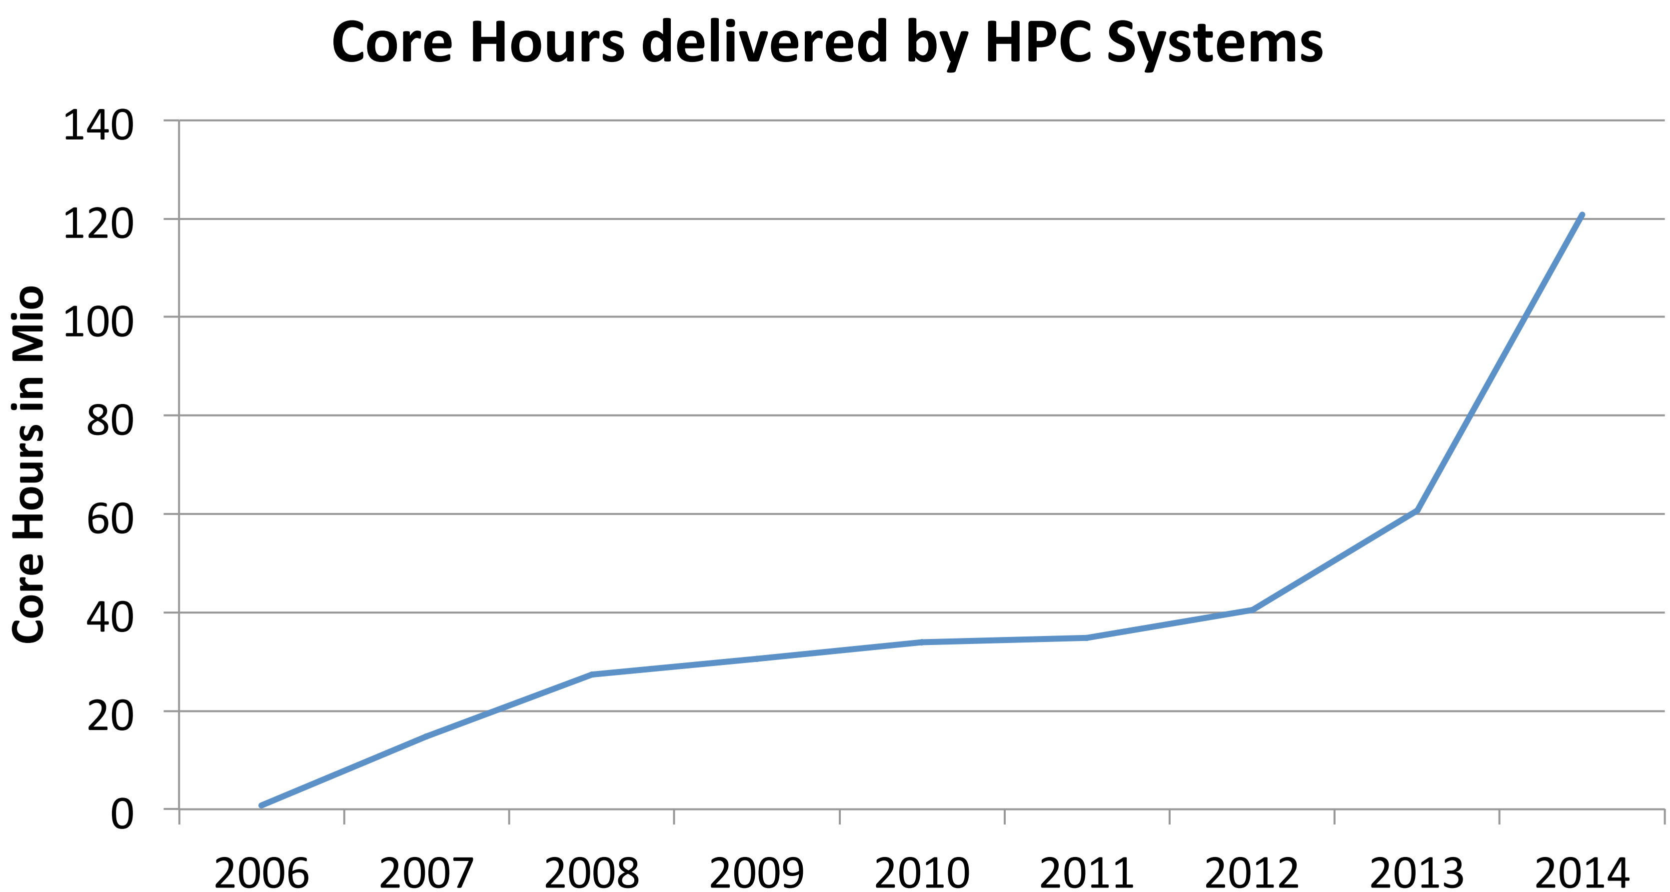

Support grant dollars vs. HPC investments

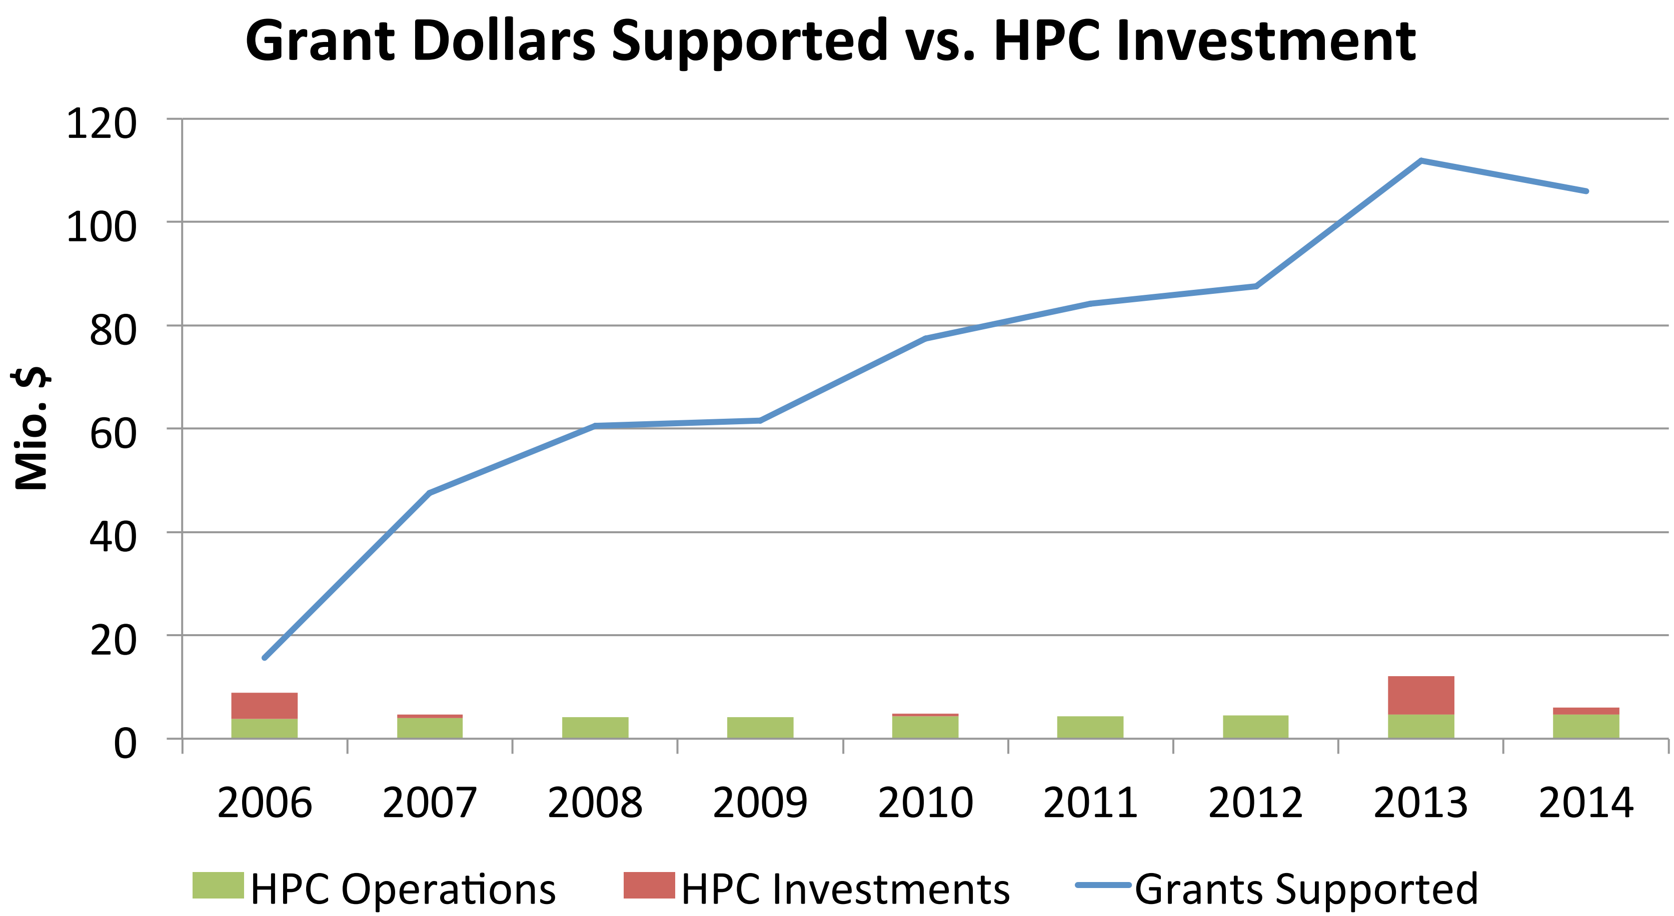

Statistics Gathering Context

- Tap into operational statistics of every group
  - Batch system, file system, tape archive,…
- Tie into enterprise information systems
  - Identity management, financial systems
- Only store derived information
- Relate scientific value and IU grant successes to investment in HPC operations (people and software) as well as HPC Hardware

Technology

- MySQL database
- Ruby on Rails scripting
- CAS authentication
- Connection to data sources
  - Cron jobs
  - SQL connections
  - File based importing
- Plan: Leverage XDMOD

ROI of Academic HPC

Rudi Eigenmann
NSF/CISE/ACI

Disclaimer: I'm presenting my own view, not an official NSF position

Academic vs Business ROI

- ROI question is very important for NSF, ACI
  - Large investments
  - Review panels of large awards are asked about
    - is taxpayer money spent wisely?

- Academic ROI is very different from the business definition.
  - Business ROI is usually short-term return. Is measurable.
  - Good Academic research has a long-term return. Is hard to measure.

  => the key discussion: how can we assess long-term impact of research – here: HPC-enabled research?

Anecdotal Evidence of Difficulty

I asked a recent NSF panel the ROI question.

The result:

- Long, heated discussion
- The panel report would say only a few lines

=> the topic is hot, but solution is elusive

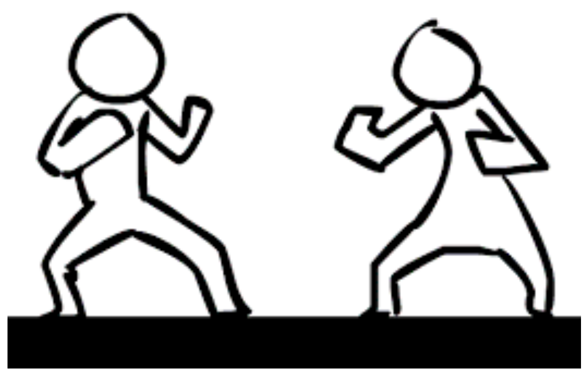

Approaches to Assessing long-term impact of HPC

with a focus on HPC infrastructure

- Scientific impact
- Enhanced productivity
- Growth, diversity of community

As a result of investments in

- Hardware (cycles, storage) and operation
- Software
- Shared services

Specific Elements of ROI

Greg mentioned in Panel proposal

- Papers, presentations, posters
- Theses and dissertations
- Graduating students who utilized supercomputing or other services
- Workshops, training
- High-value jobs for graduates
- High-value jobs within the center, including student jobs

Others

- Nobel prices
- Breakthroughs – what could not have been done without HPC – today the "third pillar of science" is arguably the biggest
  - 3D supernova simulation leads to new understanding of the universe
  - Full HIV virus simulation
  - Detailed shake maps could save lives and damage in the billions
- Breakthroughs are especially relevant when investing in leadership class machine
- The difference made by a particular investment in all the above

From Greg's Questions

- What is the basis for your ROI calculations? Does this seem to work well, at your institution?
  - Shared v distributed center model
- Do your bosses "get it?"
  - Science goals are most important for NSF; HPC infrastructure seems secondary
- What do you use to accompany ROI reporting, such as success stories, key facts, or the competitive advantage of your institution?
  - Success stories of scientific advances are key!

1. Bigger impact, due to coordinated effort of resource providers and thus broader geographic reach and engagement of a larger, more diverse community

* defining formats for processes, documentation, APIs, organizational interfaces, user interfaces

* negotiating agreements and standards for digital services in US and internationally

* negotiate MOUs with other organization in the US and in and countries. (Note, this is not just a matter of a better negotiation position; some agreements would unlikely happen if MOUs with multiple SPs were necessary.)

* Training and education authority. Even industry is looking at XSEDE for training.

2. Much improved user experience through uniformity across sites and a one-stop shop:

* Users see a more uniform experience when using multiple sites. This is important both for multi-site projects and for researchers using multiple sites over time.

* Users can apply for small and large resource on all site via one application process and can transfer machine allocations between sites

* Data at different sites appear as part of one file system, with drag-and-drop tools available.

* XSEDE.org offers single portal to CDS&E information and advanced digital services.

4. Better balancing of human and digital resources.

* technical expertise (e.g., ECSS). This includes expertise sharing and complementing. Not all sites need all expertise all the time. Users benefit from the combined expertise.

* Staff exchange (examples of staff exchange when needed?)

* Resource allocation (machine cycles, storage)

5.Better decision making based on information from across the entire VO:

* Ticket handling can draw from full history of prior incidents at all sites of the VO.

* Better projection of overall computing needs for the entire VO

* Better mapping of computational problems to the most suitable resources.

* Better security forensics

* Better implementations, leading to higher productivity

6. Coordinated Training and Education with broader geographic reach and engagement of a larger, more diverse community

(both higher impact and cost savings)

* drawing from expertise across VO

* multi-site webinar offerings with multi-site Q&A sessions to key experts

* reaching audiences across and beyond VO

* campus champion program coordinates CDS&E activities across XY universities

7. Reduced duplication of effort

* creation of educational material and lecture offerings

* definition of formats for documentation, APIs, organizational interfaces, user interfaces

* definition and execution of processes for development, operation and security, user services, communication, and external interactions.

* software implementations

8. Success stories that may not have happened without a VO (make clear what added value of the VO has enabled the success)

* Nobel prize

* HPCwire Readers' and Editor's Choice Award at SC13

* Gates Foundation (Malaria eradication) CI is using XSEDE components

9. Softens impact of resource competition by retaining staff

Four Dimensions of ROI for HPC

Nick Berente
University of Georgia

NSF CI TEAM #1240160
NSF RCN # 1148996

Four Dimensions of ROI for HPC

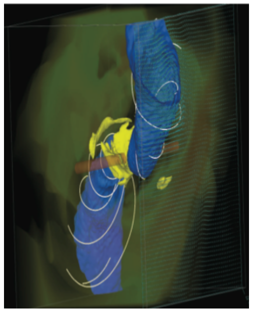

Science

- Enablement
- Findings
  - Quantity
  - Novelty
  - Quality & Importance

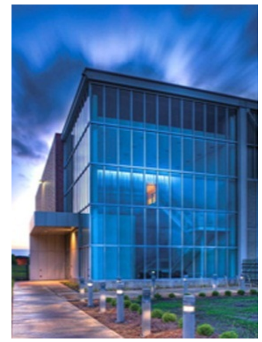

Economic

- Dollars / matching dollars
- Multipliers local / national
- Industry creation – potential & analogous
- Cost savings

Workforce

- Up-skilling
- Next-generation capabilities
- Community of practice (capacity) – risk mitigation
- Interdisciplinarity
- Diaspora & user graduation
- Alumni

Innovations

- Commercialization
- Industrial collaboration & product development
- Translation
- Technologies / patents
- New ventures

1. Identify key stakeholders (vertical axis)

	Science	Economic	Workforce	Innovation
Stakeholder 1				
Stakeholder 2				
Stakeholder 3				

2. Value Propositions for those Stakeholders

	Science	Economic	Workforce	Innovation
Nation (NSF)	✓			✓
State (GA)		✓		✓
University (UGA)		✓	✓	

2. What Brings those Stakeholders Value? – specifically!

	Science	Economic	Workforce	Innovation
Nation (NSF)	*Impact Quantity*			*Translation*
State (GA)		*Generate funds to region*	*STEM workforce*	*Commerce*
University (UGA)		*Bring funds to University*	*Impact students*	

Stakeholder Segmentation
Identify key stakeholders and specific value propositions

	Operational	Science	Economic	Workforce	Innovation
University 1					
University 2					
User 1					
User 2					
NSF 1					
NSF 2					
State 1					
State 2					
Employee 1					
Employee 2					


"All the good business leaders I know are maniacal about measuring things. They know their sales data and customer-satisfaction numbers, which divisions of their company are beating expectations and which are lagging behind…
Measurement is a big part of mobilizing for impact. You set a goal, and then you use data to make sure you're making progress toward it."
-Bill Gates

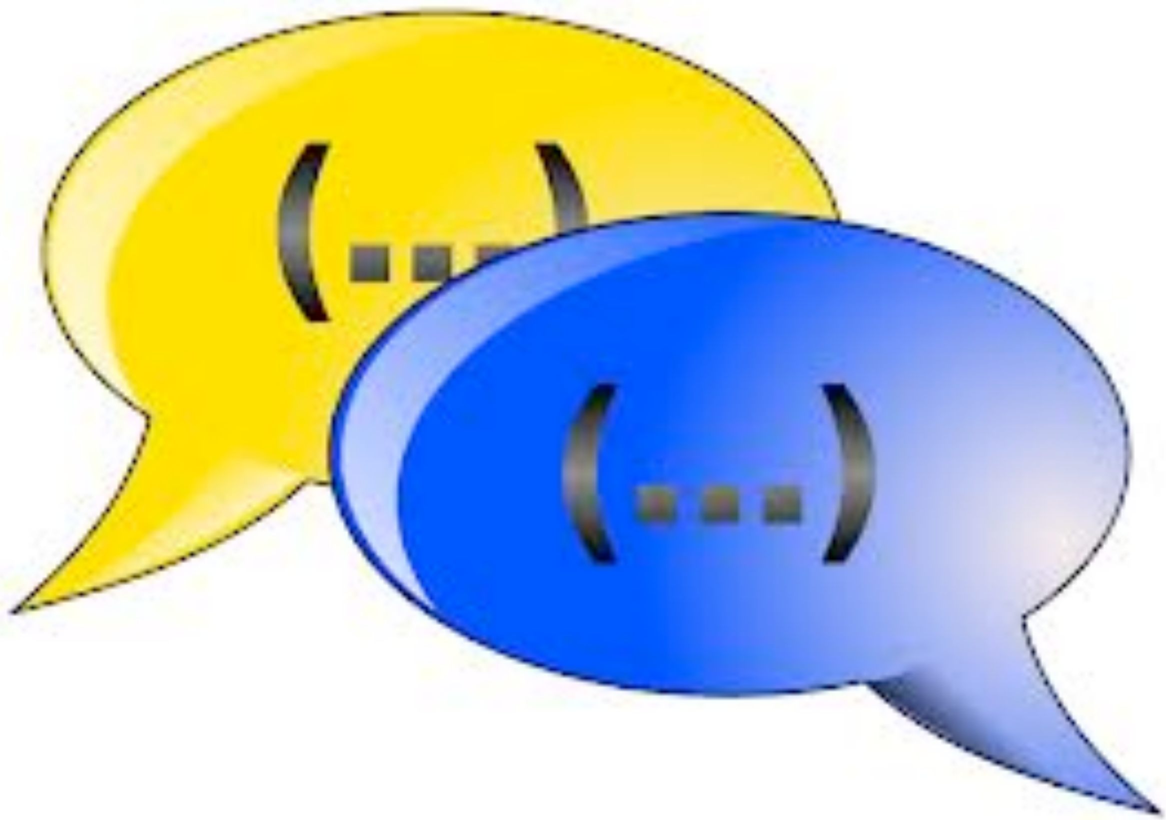

ROI as a discussion: Measure what matters!

The major products of The Academy:

Publications

Graduations

Utilization

How many faculty & student users?

Institutional outcomes

Enhancements to reputation

High-value jobs (student workers)

Intellectual ROI
Reputation ROI

Extramural funding

Research grants & contracts

Campus fees, including recharge

Innovations

Patents; royalties

Financial ROI

Equipment & related

- Costs of equipment - amortized
- Annual maintenance contracts
- Software

Personnel

- Dedicated personnel FTEs
- Shared personnel
- Breakdown of types of support provided (i.e., systems, user support, scientific programming, etc.)

Investments: Capital & Support

Energy and energy impacts

Lifecycle analysis of pollution, materials used, and other factors of operation

Data center infrastructure

HVAC

UPS

Networking

Investments: Operations

1 dozen sections on a variety of topics

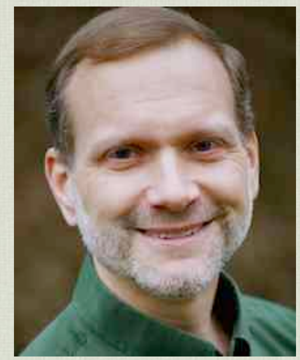

http://hpcbiz.readingroo.ms

Newby’s Monograph

Return on Investment from Academic Supercomputing: SC14 Panel

- **Feedback Links:**
  - http://bit.ly/sc14-eval
  - https://submissions.supercomputing.org/?page=SessionEval&id=stype147

SC14 Panel

Panel title:

ROI for Academic Supercomputing

Moderator:

Greg Newby, KAUST

Panelists:

Amy Apon, Clemson University
Nicholas Berente, University of Georgia
Rudolf Eigenmann, National Science Foundation
Susan Fratkin, Coalition for Academic Scientific Computation
David Lifka, Cornell University
Craig Stewart, Indiana University

Description:

Return on Investment or ROI is a fundamental measure of effectiveness in business. It has been applied broadly across industries, including information technology and supercomputing. In this panel, we will share approaches to assessing ROI for academic supercomputing.

There are good reasons why administrators at colleges and universities ask for ROI reports from center managers. For example, supercomputers are major capital expenses, with high operational costs and specialized staff. They are relatively self-contained, and exist at clearly defined physical locations. They provide built-in mechanisms for accounting for utilization of resources, such as CPU or node hours and storage blocks. The "investment" side of ROI is, therefore, fairly well defined. It includes capital expenses and ongoing operational expenses. Utilization can be tied to particular users within the enterprise.

The panel will address the challenge that "returns" from supercomputing and other computationally based research activities are often not financial. This is major distinction from other industrial sectors, where product sales, inventions, and patents might form the basis of ROI calculations. How should ROI be assessed for high performance computing in academic environments? What inroads to ROI calculations are underway by the panelists? What are challenges of ROI calculations?

Financial outcomes do occur at colleges and universities, notably the receipt of research funding or contracts. However, even when there are such financial outcomes, it can be difficult to apportion the revenues due to supercomputing versus other things. While campus funding is generally tracked diligently, centralized records are usually not kept of the different resources (such as supercomputing) needed to complete the project. It is therefore incumbent on supercomputing centers to interact with faculty members and other constituents to assert the fiscal returns that are due to the center.

Non-financial outcomes are major intellectual products of colleges and universities, and this is therefore a departure from typical ROI calculations in industry. These outcomes include:

- Papers, presentations, posters
- Theses and dissertations
- Graduating students who utilized supercomputing or other services
- Workshops, training
- High-value jobs for graduates
- High-value jobs within the center, including student jobs

Most of these outcomes can be counted, although determining whether a particular outcome is partially or entirely due to supercomputing or related resources can be challenging. Turning them into ROI ratios such as “papers per CPU cycle” or “high value jobs per terabyte” is unlikely to be intuitively appealing to campus administrators.

The panel will discuss how the financial and non-financial outcomes of academic high performance computing can be enumerated, quantified and presented to stakeholders as components of ROI discussions. They will consider how tracking such items over time can be instructive for demonstrating the growing importance of large-scale computation and storage.

Expect a lively exchange, with themes of common interest across academic supercomputing centers. Audience members will be invited to ask questions and (briefly) share their own experiences. Some of the thought-provoking questions panelists will be asked to address may include:

- What is the basis for your ROI calculations? Does this seem to work well, at your institution?
- What’s the wackiest question you have been asked by administrators to justify supercomputing's ROI?
- How do ROI considerations for supercomputing compare to other facilities and resources at your University?
- What units or offices are your biggest detractors? Your biggest supporters?
- Do your bosses "get it?"
- How do you inspire your users to actively participate in providing ROI measures?
- How do you present ROI to your stakeholders?
- Are your ROI analyses different for compute, storage, big data, support, or other components of your operations?
- What are some of the most important business concepts you utilize in discussing ROI?
- What do you use to accompany ROI reporting, such as success stories, key facts, or the competitive advantage of your institution?
- Can we work together, as an industry, on standardized reporting rubrics and data gathering strategies for calculating ROI?