

XDMoD Value Analytics: A Tool for Measuring the Financial and Intellectual ROI of Your Campus Cyberinfrastructure Facilities

Ben Fulton

Pervasive Technology Institute
Indiana University
Bloomington, Indiana 43017-6221
befulton@iu.edu

Steven Gallo

Center for Computational Research
State University of New York at
Buffalo
Buffalo, New York
smgallo@buffalo.edu

Robert Henschel

Pervasive Technology Institute
Indiana University
Bloomington, Indiana
henschel@iu.edu

Tom Yearke

Center for Computational Research
State University of New York at
Buffalo
Buffalo, New York
tyearke@buffalo.edu

Katy Börner

Information Visualization Laboratory
Indiana University
Bloomington, Indiana
katy@indiana.edu

Robert L. DeLeon

Center for Computational Research
State University of New York at
Buffalo
Buffalo, New York
rldeleon@buffalo.edu

Thomas Furlani

Center for Computational Research
State University of New York at
Buffalo
Buffalo, New York
furlani@ccr.buffalo.edu

Craig A. Stewart

Pervasive Technology Institute
Indiana University
Bloomington, Indiana
stewart@iu.edu

Matt Link

Pervasive Technology Institute
Indiana University
Bloomington, Indiana
mrlink@iu.edu

ABSTRACT

Understanding the financial and intellectual value of campus-based cyberinfrastructure (CI) to the institutions that invest in such CI is intrinsically difficult. Given today's financial pressures, there is often administrative pressure questioning the value of campus-based and campus-funded CI resources. In this paper we describe new financial analytics capabilities being added to the widely used system analysis tool Open XDMoD (XSEDE Metrics on Demand) to create a new realm of metrics that will allow us to correlate usage of high performance computing with funding and publications. The capabilities to be added will eventually allow CI centers to view metrics relevant to both scientific output in terms of publications, and financial data in terms of awarded grants.

The creation of Open XDMoD Value Analytics was funded by the National Science Foundation as a two year project. We are now nearing the end of the first year of this award, during which we focused on financial analytics. During the second year of this project we will focus on analytics of intellectual output. This module will allow the same sorts of analyses about systems and users as the financial analytics module, but in terms of intellectual outputs such as number of publications, citations to publications, and H indices. This module will also have capabilities to visualize such

data, integrated with financial data. We plan to present these tools at PEARC '18.

CCS CONCEPTS

•General and reference → Metrics; Evaluation; •Human-centered computing → Visualization; •Social and professional topics → Management of computing and information systems; Funding; •Applied computing → Business intelligence; Decision analysis; •Information systems → Data mining;

KEYWORDS

XDMoD, HPC, HPC monitoring, ROI, funding

ACM Reference format:

Ben Fulton, Steven Gallo, Robert Henschel, Tom Yearke, Katy Börner, Robert L. DeLeon, Thomas Furlani, Craig A. Stewart, and Matt Link. 2017. XDMoD Value Analytics: A Tool for Measuring the Financial and Intellectual ROI of Your Campus Cyberinfrastructure Facilities. In *Proceedings of PEARC17, New Orleans, LA, USA, July 09-13, 2017*, 7 pages.
DOI: 10.1145/3093338.3093358

1 INTRODUCTION

Universities and colleges throughout the US are under increasing financial pressures, which in many cases, leads to a comprehensive questioning of how a given institution allocates its funds. At institutions of higher education that invest heavily in local research cyberinfrastructure, there is very typically an ongoing discussion of "should we pay this much for local research computing systems?" This discussion is often instigated by campus financial leaders and decision makers. At institutions that invest less heavily there is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PEARC17, New Orleans, LA, USA

© 2017 ACM. 978-1-4503-5272-7/17/07...\$15.00
DOI: 10.1145/3093338.3093358

often an ongoing discussion of "shouldn't we invest more in local research cyberinfrastructure facilities?" This discussion is often led by faculty members and other researchers. Both questions are fair: local investment in research cyberinfrastructure (CI), particularly high performance computing (HPC) systems, involves significant and ongoing investment in equipment and staff. And these questions may also focus on different metrics: a budget discussion may focus on the financial return on investment in such local facilities, while academic researchers may be more concerned with the intellectual return on investment - particularly their ability to perform innovative research and publish their results. How does a research institution understand the value it receives from investment in research cyberinfrastructure in general and local high performance computing systems in particular? (NB: in this paper we will simply use "cyberinfrastructure systems" which might locally mean computing clusters, supercomputers, supercomputers, special purpose computational systems, storage systems, visualization systems, or all local research cyberinfrastructure facilities combined.)

Existing studies of return on investment (ROI) in campus cyberinfrastructure show that significant investment in high performance computing - particularly sustained investment - leads to increases in publications and grant income [8]. However, the work of Apon et al. [7] is correlative: increases in investment lead to increases in grant success and increased rates of peer-reviewed technical publications; the question of whether the investment was worth it to begin with is left open. The work of IDC Research, Inc. (a market intelligence provider for the information technology industry) presents a very rosy picture of the financial ROI in high performance computing [4]. However, having participated in the surveys that led to IDC's conclusions our impression is that this analysis focuses on success stories, and does not measure fully the total cost of investments in operating high performance computing facilities. There are two analyses of the return on investment in federal high performance computing facilities. One measured the financial ROI on the federal government's investment in XSEDE (the eXtreme Science and Engineering Discovery Environment), and found that the financial ROI for the federal government was greater than 1 (that is, more than \$1 in return for every dollar invested) [14]. However, this was for a unique national facility. Another study looked at intellectual impact of investment in advanced cyberinfrastructure facilities. Laszewski and collaborators [12] found that publications that utilized XSEDE and NCAR (National Center for Atmospheric Research) resources were cited more often than papers in peer journals that did not use these resources. Hence, they concluded that utilization of XSEDE and NCAR resources exerted a strong positive impact on scientific research.

Only one paper that we are aware of looks specifically at the financial ROI of investment of high performance computing: a study that looked at the financial ROI on Indiana University's investment in its one PetaFLOPS Big Red II supercomputer [15]. This paper examined grant income to Indiana University going to members of teams (primary investigators/PIs or co-investigators/Co-PIs) that had at least one person who had an account on Big Red II. The amortized total dollars were \$11.7M, \$24.5M, and \$39.8M from FY14 through FY16 respectively (see also Figure 2), and Facilities and Administration funds ("indirect costs") from these grants would, on average, be slightly more than \$8M per year, while the amortized

cost of purchasing and operating Big Red II is in the neighborhood of \$3M per year. So, there is evidence that the financial ROI for IU's investment in Big Red II is positive (> 1).

1.1 Open XDMoD

The most widely used software package in the US for understanding how computer clusters and supercomputers are used is Open XDMoD. More than 200 institutions are already using Open XDMoD to analyze the usage of their cyberinfrastructure resources. The Center for Computational Research (CCR) of the University at Buffalo began developing an open source tool to provide metrics, basic accounting and visualization of CPU (central processing units) and storage usage at that institution [13]. Initially called "UBMoD", this tool was developed to answer questions about the value of the CCR by documenting the use of CCR high performance computing systems by researchers - particularly leading researchers - of the University of Buffalo. Project leader Thomas Furlani and his colleagues were awarded a grant from the National Science Foundation (NSF) to extend the functionality of this tool for all XSEDE-supported high performance computing systems [6]. XSEDE is the largest single cyberinfrastructure project funded by the NSF, and it integrates and supports several high performance computing (HPC) systems [16]. With NSF funding, CCR staff expanded the application to include job-level performance data and quality of service metrics, along with a high quality graphical user interface. Importantly, the NSF recommended that all major CI facilities supported by XSEDE operate XDMoD as a way to assess system utilization and performance; this will expand usage quickly. XDMoD proved so valuable that the CCR created an open source version of that tool, called Open XDMoD [3], which can be installed on any computing cluster or HPC system to provide usage, performance, and quality of service metrics. Figure 1 shows an example: a line chart displaying the CPU hours utilized per job for a variety of fields of science in 2016. Open XDMoD includes a variety of methods and drill-down mechanisms to measure resource utilization and performance and facilitate monitoring of the jobs that are running on the system to determine their efficiency and resource consumption. This knowledge is beneficial in planning for future upgrades and acquisitions.

1.2 Goals of this Report

Indiana University (IU) and the University at Buffalo aimed to add two categories of functionality to Open XDMoD, creating a variant of this program that will be called Open XDMoD VA (Value Analytics). Open XDMoD VA will, when completed, expand existing capabilities so that the software can be used for:

- Financial analytics: Analyze grant income to researchers that use local cyberinfrastructure (CI) and relate that income to use of local CI systems.
- Intellectual output analytics: Analyze publications by researchers that use local CI systems and citations to those publications

The creation of Open XDMoD VA is planned to be a two year project, with the first year focusing on adaptation and release of the financial analytics tools, and the second year focusing on release

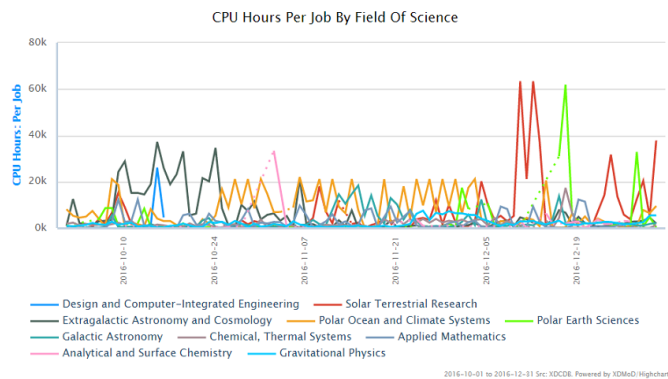


Figure 1: XSEDE CPU hours per job by field of science for calendar year 2016. This is an example of a chart that is already made available by XDMoD.

of the intellectual output analytics tools. In this paper, we describe the current status of the financial analytics tools.

2 FINANCIAL ANALYTICS

Financial leaders in any organization want to talk in terms of finances, in addition to other topics that might be relevant to an organization. Universities aim to persist over time, and any organization that so aims has to make sure that its finances balance in the long run. Indiana University has in the past 20 years made significant investments of its own funds in local cyberinfrastructure. To address local questions about the financial value of these investments, Indiana University's Pervasive Technology Institute (IU PTI) [5] began development of a tool to better understand usage of IU's local cyberinfrastructure and explain the value of that cyberinfrastructure to researchers and institutional leaders in financial terms.

The tool developed by IU PTI included some of the functionality included in Open XDMoD, but were distinguished from the functionality previously offered by Open XDMoD in that IU PTI's tool included a great deal of financial information. Indiana University uses the Quali Financial Services package [1] to track grants awarded to university researchers from the NSF, the National Institutes of Health (NIH) and other sources. IU PTI was able to get permission to draw data from this financial data system to analyze the relationship between use of IU HPC systems and grant income to the institution. Figure 2 shows an example of the financial analysis that this tool enables.

With the NSF having established the University at Buffalo's Center for Computational Research XDMoD tool as the *de facto* standard for assessment of advanced cyberinfrastructure use, IU PTI contacted CCR about integrating some of the unique capabilities of the financial and Return on Investment (ROI) analysis tools IU PTI had been developing into Open XDMoD. These modules are intended to provide a starting point for assessing the value of investment in campus-based CI in financial terms (grant income from researchers who use campus CI). By presenting a view of the institutional financial, collaboration, and publication data alongside the current HPC usage analysis in XDMoD, the module provides

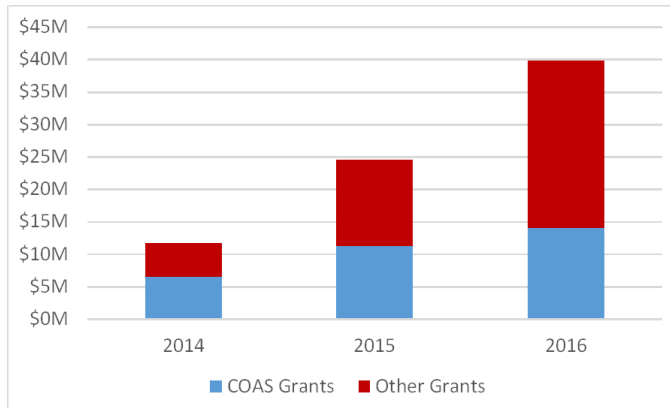


Figure 2: Example of the sort of data analyses that will be available via OpenXDMoD Value Analytics. This figure shows grant income to Principal Investigators with accounts on Indiana University's Big Red II cluster. A particular point of local concern has been value to users in the College of Arts and Sciences, hence the breakdown in this graphic.

users with invaluable insight into the value of investing in cyberinfrastructure. The financial analysis module added to Open XDMoD allows users to analyze grant income to a university and relate it to use of local cyberinfrastructure systems. The module currently has early support for generating datasets and charts linking PIs' grant income to their usage of HPC resources, but further data and usability testing is necessary before release of this feature. In the future, we would like to extend the linking capabilities of the system to involve Co-PIs, key personnel, and departments or other academic units. We currently provide data export tools for Quali Coeus, but we plan to extend our support to include Peoplesoft, Ellucian, and the NSF and NIH web sites. Taken in sum, the anticipated and existing key capabilities provided by the module are as follows:

- Analyze grant income to a university and relate it to use of local cyberinfrastructure systems, relating grants (and grant income) to use of local CI systems by Principal Investigators or Co-PIs;
- Correlate amount of usage of local CI systems with funding by PIs and / or Co-PIs;
- Analyze grant income to a university by department or other academic unit and correlate that with use of local CI systems;
- Analyze grant income using just current awards, or current and past (expired) awards;
- Ingest financial data from the following sources:
 - Quali Coeus
 - Peoplesoft
 - Ellucian
 - NSF web site
 - NIH web site

To facilitate the loading of data into Open XDMoD, we have defined schemas for JSON (JavaScript Object Notation) [2] documents that specify the structure of the data that XDMoD is designed to




Figure 3: Scatter plot of the Indiana University award funding in millions of dollars vs HPC usage in millions of CPU-hours covering the time period from January 2015 to May 2017. The red line is a linear fit. The fit p-value indicates that the fit is statistically significant.

handle: one schema for grant information and one schema for information about people associated with a grant. Information from supported financial systems will be transformed into these standardized file formats, and from there loaded into the Open XDMoD database. The point in the case of both documents is to get as much local financial information as possible regarding grant awards and people set in a format that Open XDMoD can read. For each grant award, the schema includes fields for PIs, Co-PIs, and Key Personnel, starting and ending dates, dollar amounts, the funding agency, and grant identifiers both for the organization and the agency. For lists of people, the schema includes basic information including some local identifier, so that the people list and grant award list can be properly linked and analyzed and to assist in cross-referencing with HPC usage data. Providing as full a listing of grant awards including award Principal Investigators and Co-PIs as possible enables an analysis of the role of cyberinfrastructure systems in grant successes. People are the key link between systems usage and grant award success. Because XDMoD VA is open source, users of course have the option of making modifications to the ingestion process if a local financial system is not currently supported. The use of JSON documents means that direct interfaces to financial systems are not needed; all that is required is a way to extract data from a financial system and transform it into a properly formatted JSON

document. We remain open to expanding the list of supported financial systems based on community input.

XDMoD VA is being developed so that it does not require a direct connection to a university or college's local financial systems, providing flexibility for data acquisition. Where permitted, XDMoD VA will ingest data directly from the list of financial management systems listed above. These financial systems are in use by 56% of United States higher education institutions [11]. If it is not possible for a college or university's IT organization to have direct read access from the institution's financial systems, XDMoD VA will provide the capability to download award data directly from NSF and NIH web services. This is a tremendously important point. The fullest source of financial award data will always be the office of the chief financial officer of any institution - such data sets will include for example grant awards from private foundations, which are generally not well publicized. So if a cyberinfrastructure facility can work with its local institutional financial organization, great. If not, a cyberinfrastructure organization staff can go directly to the NSF and NIH web sites to obtain complete information about grant awards to their own institution. This provides a good look at the most important sources of grant funding to users of very many (probably most) campus computing centers. While this is simply not as complete as the data one might get with full access

to campus financial systems, ingestion from NIH and NSF web sites provides a way for computing centers to have access to financial data independent of cooperation from local campus financial organizations.

Using this grant data, administrators of cyberinfrastructure systems may analyze cross-referenced data about system usage with funding data. For example, the Indiana University award funding correlates linearly with the HPC usage over the 2.5 year time frame for which both are available, as shown in Figure 3. The linear fit (as shown by the red line) is highly statistically significant with a p -value < 0.01 . Note that the usual caution should be stated: correlation does not necessarily imply causality. However, we know from interviews with well-funded researchers who make extensive use of IU cyberinfrastructure facilities that they could not do their current research projects in the absence of these resources [9]. Indeed we hear from many such researchers that local availability of supercomputers, on demand, is critical to their research activities and competitiveness for grant awards. So, while correlation does not necessarily prove causation, we have a correlation with independent sources of data that suggest competitiveness for grant awards is aided by IU cyberinfrastructure resources. Figure 3 also reflects one element of IU strategy: we intentionally invest significant cyberinfrastructure resources in disciplines that typically do not generate much grant funding, because IU views these areas (including digital humanities, interior design, and apparel merchandising) as important to the overall mission of IU.

As another example, a computing center may want to view all grants brought in by users of a particular cluster for the previous financial year. In Figure 4 below, we show many of the data already assembled and summarized by Open XDMoD (from the start of the list at "Jobs Summary" to "Jobs by User") and several of the new financial reports that are available in Open XDMoD VA.

2.0.1 Availability. As of this writing, the XDMoD Value Analytics financial modules were expected to be available in beta form on July 1st, 2017 from the XDMoD-VA Github repository (<https://github.com/ubccr/xdmod-value-analytics>). Our current work plan is to have these modules available in final production form on November 1st, 2017.

2.0.2 Return on Investment. Finances are important to any organization that hopes to persist over many years. But institutions of higher education - research institutions in particular - exist to create and transmit knowledge. A critical aspect of ROI is the return in terms of intellectual accomplishment on financial investment in HPC systems. With the capabilities now available in the closed beta version of Open XDMoD VA one can ask and answer questions such as:

- What is the total amount of grant money brought into the university by users of a particular cyberinfrastructure system, or by groups of cyberinfrastructure systems? (The tool is thus useful for both departmentally managed CI systems and CI systems managed by the central IT department)
- What is the total amount of facilities and administration funds brought into the university by users of a particular cyberinfrastructure system, or by groups of cyberinfrastructure systems?

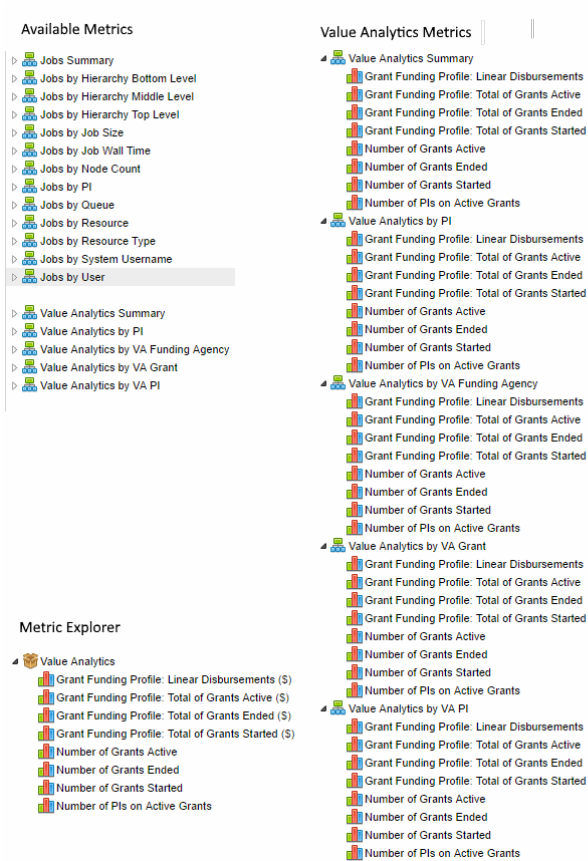


Figure 4: Examples of metrics available via Open XDMoD with the Value Analytics module enabled.

- What is the correlation between the amount of use of an institution's CI systems and the total amount of grant funding brought in by a particular researcher, group of researchers, department, or school?

These analyses are still imperfect. For example, we have yet to create a tool that will assess how important use of a particular CI system is to the success of any particular grant proposal that is funded by some funding agency. That is a future step and requires among other things the surveying of researchers to obtain their opinions on this matter, and precisely calculating a financial return on investment in cyberinfrastructure facilities will be possible only with this information. While Open XDMoD is by itself not quite sufficient to calculate ROI (Return on Investment) in cyberinfrastructure facilities, in some circumstances it can be used to make qualitative yet convincing arguments that ROI for a given institution is greater than 1 (that is, more than one dollar of financial benefit to the institution for every dollar invested). For example, Thota et al. [15] looked at the grants received by researchers at Indiana University who used IU's Big Red II supercomputer, and compared the annual cost of operating that system (amortized system cost, maintenance, and staff support) to the facilities and

administration funds brought into the university by the researchers that used Big Red II. Thota and his colleagues found that

The facilities and administration funds income to [Indiana University] that comes along with grant awards to people who use Big Red II is twice what Big Red II costs to operate. [15]

In a case like this, it is fair to make the assessment that the CI system is a good investment. (And it's hard to make the counter arguments such as 76% percent of the university's grant income goes to users of Big Red II and other CI systems, the facilities and administration income to the university is two times the cost of operating Big Red II, and somehow having this supercomputer is a bad deal?). This line of argument isn't perfect, but it is credible and better than can be done in the absence of the tools we are implementing as part of XDMoD VA.

3 INTELLECTUAL OUTPUT ANALYTICS

Measuring intellectual outcomes is difficult, particularly since the results of intellectual accomplishments may take years or decades to be fully realized. Measures of intellectual output commonly include:

- Total publications by a researcher (or group of researchers)
- Citations to a researcher (or group of researchers)
- H index [10]
- The ability to compare total publications, citations, and H-indices between researchers (or groups of researchers) who do and who do not use a particular CI resource
- The ability to create maps of "influence" and "collaboration" among researchers, research groups, and academic units
- Multivariate diagrams that relate use of CI systems, funding, and publication metrics

4 CONCLUSIONS

The Open XDMoD- Value Analytics module (Open XDMoD VA) will allow users to analyze financial metrics related to use of local cyberinfrastructure facilities by researchers, groups of researchers, and academic units. The capabilities of Open XDMoD VA will allow ingestion of data from an institution's financial system if local policy permits this. If not, then the staff of a local cyberinfrastructure facility can download grant award data directly from NIH and NSF web pages. With this information, it will be possible to answer questions such as

- What is the total amount of grant award funding brought into an institution by all users of a particular cyberinfrastructure facility (or set of facilities)?
- How do researchers (or groups of researchers) who use local cyberinfrastructure facilities compare to researchers who do not, in terms of grant award success?

The financial analytics module does not quite yet allow one to calculate Return on Investment in a financial sense. For example, there is no mechanism yet to assess how important a particular cyberinfrastructure facility (or groups of facilities) are in terms of obtaining a particular grant. But used carefully, it can allow in certain circumstances the drawing of conclusions regarding ROI. Open XDMoD VA in general should be used carefully and with

discretion. On the one hand, it could be easily possible to double (or more) count the value of cyberinfrastructure systems by analyzing them each individually and summing the results to arrive at a misleading value, when in fact there is often much overlap of users of cyberinfrastructure systems. Those who use CI resources tend to use many of those available to them. On the other hand one must be careful in discussing finances. While it is essential when having discussions with institutional financial leaders to have those discussions in financial terms, it is all too easy to cause offense when talking dollars and ROI with faculty members.

In this report we have focused on the financial analytics capabilities being developed as capabilities to be added to the already well known and widely used XDMoD, creating a tool called XDMoD VA. The financial analytics capabilities will be available in a closed beta on July 1st, 2017 and will be production-ready on November 1st, 2017. During the second year of this project we will focus on analytics of intellectual output and visualization of results; we plan to report on this at PEARC '18.

5 ACKNOWLEDGMENTS

This material is based, in part, upon work supported by the National Science Foundation under Grant No. 1053575 and 1566393 and also by the Indiana University Pervasive Technology Institute (founded with funding from Indiana University and the Lilly Endowment, Inc.). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, Indiana University, or the Lilly Endowment. We would like to thank our vendor partners Cray and DDN for their collaboration in the implementation of Big Red II, which drove much of the early financial analytics development at IU. We thank Dr. Winona Snapp-Childs for her editorial efforts; any errors remain the responsibility of the authors.

REFERENCES

- [1] 2005. Kuali Financial Services. <http://www.kuali.org/>. (2005). Accessed: 2017-03-02.
- [2] 2009. JSON Schema: A Media Type for Describing JSON Documents. <http://json-schema.org/latest/json-schema-core.html>. (2009). Accessed: 2017-02-24.
- [3] 2010. XDMoD (XD Metrics on Demand). <https://github.com/ubccr/xdmod>. (2010). Accessed: 2017-03-02.
- [4] 2015. IDC Presents an Update on ROI with HPC. <http://insidehpc.com/2015/09/idc-presents-an-update-on-roi-with-hpc/>. (2015). Accessed: 2017-05-24.
- [5] 2015. Indiana University Pervasive Technology Institute. <http://www.PTI.IU.edu>. (2015). Accessed: 2017-05-24.
- [6] 2015. XD Metrics Service. https://www.nsf.gov/awardsearch/showAward?AWD_ID=1445806. (2015). Accessed: 2017-05-24.
- [7] Amy Apon, Stanley Ahalt, Vijay Dantuluri, Constantin Gurdgiev, Moez Limayem, Linh Ngo, and Michael Stealey. 2010. High Performance Computing Instrumentation and Research Productivity in U.S. Universities. *Journal of Information Technology Impact* 10, 2 (2010), 87–98. <https://ssrn.com/abstract=1679248>
- [8] A. W. Apon, L. B. Ngo, M. E. Payne, and P. W. Wilson. 2015. Assessing the effect of high performance computing capabilities on academic research output. *Empirical Economics* 48, 1 (2015), 283–312.
- [9] Lizanne DeStefano and Lorna Rivera. 2015. *Cyberinfrastructure Value Assessment Report*. Technical Report.
- [10] Hirsch JE. 2005. An index to quantify an individual's scientific research output. In *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 102. National Academy of Sciences, 16569–16572. Issue 46. DOI: <http://dx.doi.org/10.1073/pnas.0507655102>
- [11] L. Lang and J.A. Pirani. 2014. "The Financial Management System: A Pivotal Tool for Fiscal Viability.". Research bulletin. ECAR, Louisville, CO.
- [12] G. v. Laszewski, F. Wang, G. C. Fox, D. L. Hart, T. R. Furlani, R. L. DeLeon, and S. M. Gallo. 2015. Peer comparison of xsede and ncar publication data (*IEEE*

- International Conference on Cluster Computing*). 531fi?!532. DOI : <http://dx.doi.org/10.1109/CLUSTER.2015.98>
- [13] JT Palmer, SM Gallo, TR Furlani, MD Jones, DeLeon RL, JP White, N Simakov, AK Patra, J Sperhac, T Yearke, R Rathsam, M Innus, CD Cornelius, JC Browne, WL Barth, and RT Evans. 2015. Open XDMoD: A Tool for the Comprehensive Management of High-Performance Computing Resources. *Computing in Science & Engineering* 17 (July/August 2015). DOI : <http://dx.doi.org/10.1109/MCSE.2015.68>
- [14] Craig A. Stewart, Beth Plale, Von Welch, Matthew R. Link, Therese Miller, Eric A. Wernert, Michael J. Boyles, Ben Fulton, David Y. Hancock, Robert Henschel, Scott A. Michael, Marlon Pierce, Robert J. Ping, Tassie Gniady, Geoffrey C. Fox, and Gary Miksik. 2015. Pervasive Technology Institute Annual Report: Research Innovations and Advanced Cyberinfrastructure Services in Support of IU Strategic Goals During FY 2015. (2015). <http://hdl.handle.net/2022/20566>
- [15] Abhinav Thota, Ben Fulton, Le Mai Weakley, Robert Henschel, David Y. Hancock, Matthew Allen, Jenett Tillotson, Matthew Link, and Craig A. Stewart. 2016. A PetaFLOPS Supercomputer as a Campus Resource: Innovation, Impact, and Models for Locally-Owned High Performance Computing at Research Colleges and Universities. In *SIGUCCS 16*. ACM.
- [16] John Towns, Timothy Cockerill, Maytal Dahan, Ian Foster, Kelly Gaither, Andrew Grimshaw, Victor Hazlewood, Scott Lathrop, Dave Lifka, Gregory D. Peterson, Ralph Roskies, J. Ray Scott, and Wilkins-Diehr Nancy. 2014. *Computing in Science & Engineering* 16, 5 (Sept.-Oct. 2014), 62–74. DOI : <http://dx.doi.org/10.1109/MCSE.2014.80>