Assessing the Impact of a CyberTraining Project: Expanding the Metrics

Sandra B. Nite
Texas A&M University
s-nite@tamu.edu

Joshua Winchell
Texas A&M University
jwinchell@tamu.edu

Marinus 'Maikel' Pennings
Texas A&M University
pennings@tamu.edu

Dhruva K. Chakravorty
Texas A&M University
chakravorty@tamu.edu

Keith Jackson
Texas A&M University
kjackson@tamu.edu

ABSTRACT

As training on cyberinfrastructure resources becomes more common, we show the progression of metrics used to measure the effectiveness and impact of informal computational training courses that are provided by the Texas A&M University High Performance Research Computing facility. These courses were built to support researchers from research groups that have a background in computing practices. As such, the courses were structured as information-sharing sessions with the primary method to measure course success being frequency of participation. While these metrics inform about the interest in these courses, they relied on researchers continuing the learning process in their laboratories. As computing becomes ubiquitous in research programs, researchers who have no peer-learning mechanisms participate in these courses. Researchers are now participating in a continuum of courses that cover introductory to advanced topics and rely on them to build proficiency in research computing technologies.

We report on a pilot program that pivots along the way to support these researchers. We collected additional metrics to learn about the impact of the training materials for individual researchers. These metrics include participation in course activities, time spent logged on compute clusters, and views of course recordings and other asynchronous training materials. Surveys are now structured to identify the needs of individual researchers. Some of these metrics require additional processing time but will assist in understanding how researchers learn in these environments.

KEYWORDS

cybertraining, short courses, high performance computing

1 INTRODUCTION

Texas A&M University (TAMU) High Performance Research Computing (HPRC) was established in 1989 [13] and serves the research computing needs of all Texas A&M universities and state agencies and provides researchers with exceptional High Performance Computing (HPC) resources. Besides hardware, users have access to several commercial, free, and open-source software packages. HPRC staff includes system administrators and research scientists as well as graduate and undergraduate student assistants.

2 THE HPRC TRAINING PROGRAM

In 2017 TAMU HPRC was awarded a CiSE ProS Cybertraining award that was funded by the National Science Foundation (NSF). This funding, followed by other state and federal funding agencies helped expand training opportunities to the research community. As part of this effort, development of the short course program accelerated. A list of topics in computing were identified as key areas of importance for researchers interested in computing. For each topic area, complementary learning objectives and outcomes were identified to guide the design of courses. Delivery was offered remotely and in-person, and hybrid instruction was introduced in 2018. Initially, each course had a lecture component followed by a series of hands-on exercises. Exercises were incorporated throughout the courses to keep attendees engaged throughout the training [1, 3, 8, 10, 11]. After considerable trial and error, it was determined that two and a half hours was the appropriate duration for a short course. Each course was followed by a brief survey asking researchers whether the course content suited their learning needs. The short courses are modular, and a series of short courses can be stacked on top of each other to develop a semester-long series of classes. While this model was refined over the years, the underlying structure largely remained the same. Researchers can learn from a series of topics from introductory to advanced levels.

Each semester TAMU HPRC offers 60+ CI-training focused camps, hackathons, workshops, one-on-one consulting sessions, and train-the-trainers programs for scientific applications. The guiding principle for these training programs is to help researchers use our CI resources for science. In this spirit, the focus remains on teaching researchers how to use science and engineering applications software and/or workflows on CI. Underlying technologies (e.g., containerization, AI/ML frameworks, composability) are taught under the auspices of these applications. We have developed pedagogical approaches to CI training and were arguably the first group to offer micro credentials in research computing. Our informal training program offers 30+ CI-specific courses in 3 pedagogical formats to 4,000+ researchers every semester. These training courses are modular and emphasize hands-on activities. Asynchronous self-paced training for research workflows and CI are offered on the Canvas and Google Classroom Learning Management Systems. These asynchronous classes give researchers the opportunity to learn at

https://doi.org/10.22369/issn.2153-4136/15/2/2

their own pace, while earning micro credentials. In a parallel thrust, the TAMU HPRC YouTube channel offers ADA-compliant training videos in 5-, 50- and 160-minute increments. The channel has 1,100+ subscribers and typically garners 1,100 views every month. Finally, HPRC is home to CI communities like the NSF SWEETER CyberTeam and BRICCs consortia. These efforts have extended to tutorials and workshops at the annual SC, PEARC, and IEEE Frontiers in Education conferences as well. These training programs dovetail into one-on-one deep consulting with researchers via the "Bring Your Own" Science, Code, and Data series. Research focused training programs for skilled CI researchers have been supported by previous NSF CyberTraining awards (OAC-1829799). HPRC scientists are members of the MATCH CSSN (Computational Science Support Network) and CCEP (CSSN Community Engagement Program) awardees. They also lead into our strong K-12 programs that extend from teacher preparation workshops with IEEE and ACM at (SuperComputing22 (SC22) and SuperComputing23), and week-long summer camps that have introduced over 400 school students to computing. The strength of these programs is best represented by our ability to recruit students from various backgrounds. Local news outlets have featured the camps as well, broadening the reach throughout the community.

The focus for this study is the short courses and primers that are taught each semester. This is the largest piece of our training program, and we are working to improve the metrics for formative assessment. Summative assessment will occur at the end of the first year of the ACES testbed implementation. HPRC provides a series of short courses and primers each fall and spring semester to teach researchers how to enhance their use of research computing clusters. Primers are hour-long events and are often taught by graduate students. Short courses are 2.5-hour-long events taught by HPRC staff and industry partners. These courses build from introductory topics to advanced courses. The content of HPRC's short courses includes the basics needed to use the clusters, coding specifications for certain types of processing, and programming languages for specific applications. They are provided over a period of about 10 weeks, with a morning and afternoon training event on Tuesdays and Fridays. There is a mixture of face-to-face, online through Zoom, and hybrid venues. In particular, the ACES courses are delivered on Zoom or as hybrid courses so that they are available to attendees across the U.S. Most of the training courses incorporate practice exercises for the participants to work through to ensure their understanding of the concepts and procedures being presented. As the series of courses were developed, the most important metric was a count of participants registered for the courses. This metric showed that the courses were being attended and gave information about which courses enjoyed the highest attendance. This has continued to be the primary metric of interest. As the series of short courses expanded, the metrics used are changed to support a better understanding of the effectiveness and impact of the training program. The knowledge gained about the effectiveness of the training will be useful in the effort to increase the impact on the university campus and impact a broader community of researchers across the United States.

3 PURPOSE OF THE STUDY

In this study we show the process we have used to gather better data about our training program. We will examine what is missing from a thorough understanding of the impact of the short course training program at our institution and beyond as we reach out across the United States to provide high performance compute resources for research. Our training program will expand, and we need better metrics to fully understand the impact and how to improve our effectiveness in providing the resources researchers need. Our research questions are: 1) How well are our current short course and primer offerings serving the needs of the research community? 2) What metrics can deepen our understanding of what researchers need and how to fulfill those needs?

4 METHODOLOGY

For the first few years, registration numbers and numbers of courses offered were the primary metrics considered. However, the courses are continually evolving and offerings expanded as compute cluster components become more sophisticated. For example, with the composability available on FASTER and ACES, some short courses needed to address this. Similarly, with the novel accelerators offered on these clusters, the porting of code is not a straightforward process. Thus, the training model for these courses needed to evolve [7]. As we considered the metrics of numbers of short courses and numbers of registrations for each, it became clear that we needed to improve our data collection to include additional information about the impact of our training program. The number of registrations was not sufficient to determine the level of interest. We needed to take attendance to see how many actually followed through with their intent to take the course. Further, in order for researchers to use the clusters as taught in the short courses, attendees needed to do more than watch a presenter. Thus, we have worked to ensure that the short courses are not lecture presentations but instruction with brief exercises for participants to try out during the training. To this end, we added a new metric in the Fall 2023 semester, tracking the number of attendees that logged in to the cluster. Besides gathering additional information to examine this new metric, we realized we needed a different teaching style from lecture and demonstration. We needed to know how many attendees participated by logging on and trying the exercises given in the training. From experiences teaching credit courses, we know that attendees who watch the presenter but do not log on to the computing cluster and complete the exercises are much less likely to use the cluster later. We shifted our focus beyond the metrics of frequency of registrants and attendees. We began checking the cluster during the training to ensure attendees are logged into the clusters. If we see they are not, we offer help at the beginning of the training to get them logged on to the cluster so that they can complete the tasks and exercises throughout the course

5 RESULTS

Although the HPRC training program is multi-pronged, we limited ourselves here to the most homogeneous events that are most often in advertisement and delivery across semesters. Thus, we focused on the short courses and primers offered regularly during school semesters, with the exception of Summer 2020, when a

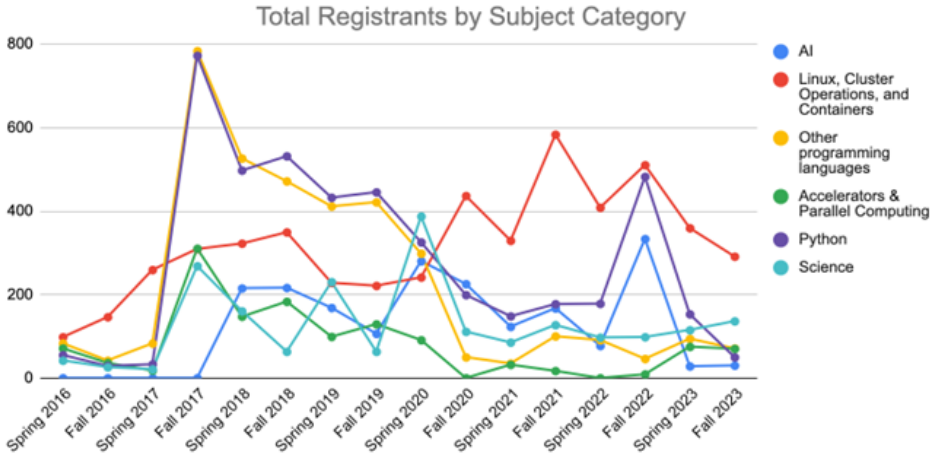

Figure 1: Registration counts for groups of classes from Spring of 2016 to Fall of 2023.

series was offered online during the COVID-19 pandemic. Short courses are generally in Fall and Spring because the Summer is busy with K12 outreach and conferences. There is some difficulty in analyzing even the short course and primer data because our specific course offerings varied by semester, meaning a course-by-course representation would be both too bulky and too difficult to use to draw cross-semester conclusions. Thus, we grouped our classes into six broad categories that we tracked across semesters.

Figure 1 shows the registrations for six categories of short courses taught during school semesters since Spring of 2016. From these we can see the areas in which interest has grown or waned over time. The number of registrations gave us information about the level of interest in the courses, but the number of attendees gave us a greater understanding of the strength of that interest. These metrics for the short courses delivered from 2016-2018 and for 2020 were previously reported [2, 4]. We make two observations from Figure 1. Firstly, there is a jump in registrations in 2017, especially for courses about "Python" and "Other programming languages," as new funding allowed for remote classes and expanded offerings of in-person classes [2]. Secondly, there was a drop in overall registration in the pandemic years, which is now beginning to recover.

In Figure 2, we look at more detailed metrics for the Fall 2023 semester, checking attendance and cluster logins in addition to registrations. These courses were offered at no cost; thus, researchers who registered were not fully committed if something else arose. The metrics that were especially useful are the number who attended and the number who logged into the cluster.

In addition to our live courses, we process videos from recordings of short courses and provide YouTube videos on the TAMU HPRC channel, providing greater access for those who cannot attend or want to review the material. There have been almost 10,000 views to date. Some courses (e.g., use of containers, programming, cybersecurity) are very popular with researchers [5, 12].

Our first research question was, "How well are our current short course and primer offerings serving the needs of the research community?" The metrics shared thus far have several indicators that we are meeting those needs well. We can see this in the numbers that attended the new offerings. One of the newest clusters, ACES, just reached the testbed state. This cluster has a number of specialty resources for researchers. During the summer of 2023, we held a training conference and invited researchers to attend to be introduced to these resources. The same was true of the newest cluster, Launch, which went online in December 2023. A quick look shows that attendees were extremely happy with the training. We identify needed improvements through surveys.

The second research question was, "What metrics can deepen our understanding of what researchers need and how to fulfill those needs?" We recognized through this study that there are additional metrics that could be useful as we continue to improve our training program. First, we would like to know how many attendees completed the exercises during the courses. We have discussed possibly giving them an assignment to complete and providing a certificate for attending and completing the assignment. However, there is work to do to prepare the assignment as well as preparing and sending certificates after checking the assignment. We are considering customizing the survey for different courses to gain optimal information.

6 DISCUSSION

We believe collecting compute times on the clusters for short course attendees in subsequent semesters could add to the useful metrics. The average compute time of short course attendees will help us see how much they use the HPRC resources in their classes and/or research, if we can obtain these data. A study on e-resources usage revealed that among the five factors that had an impact on usage, "influencers" were the most powerful in affecting intention to use the resources. E-resources need to be organized, be easy to access, and meet the researcher's needs. With a good first experience, users will try again and spread the information to friends and colleagues [14]. A survey question about how likely an attendee is to recommend our training courses might prompt the attendee to encourage fellow researchers to avail themselves of the training sessions. Finally, we plan to construct a survey for PIs to ascertain how much they believe the short courses help their students in their

Figure 2: Frequency of registrations, attendance, and engagement for Fall 2023 Short Courses.

research projects. A similar survey could be sent to all attendees at the end of each semester's series of primers and short courses. We may create a survey to ask if there are other training courses that would further enhance the work of the researchers that use our clusters.

Understanding the impact of a training program on the intended audience is not simple or straightforward. Counts of registration and attendance are the bare minimum metrics to inform the training provider. Actual engagement in the training goes a step forward. Before that metric can be used, the design of the training may need to be changed as well. HPRC's training program began more as a lecture series where attendees were encouraged to follow along on their computers, but little help was available during the session to ensure that level of engagement. The courses have evolved in a way that provides helpers in the session to assist attendees who have difficulty accessing the computing clusters. In addition, exercises for the attendees are embedded. The presenter pauses for response from the attendees, answers questions, and then shows the process and answer for the exercises. We are noticing that the engagement assessed through percentage of those logged onto the cluster is approaching 100% of the attendees.

6.1 Expanding Beyond the Institution

The metrics discussed have been primarily considered in view of the impact of the short courses at our institution, TAMU. However, our new ACES cluster is available to researchers across the United States and just moved into the testbed phase. We plan to use many of the metrics discussed previously on the short courses for which attendees use the ACES cluster. Accounts are controlled through ACCESS, (the successor of XSEDE), but we will be able to collect information about logins and compute times on this cluster. These metrics and others we may design for the future will provide information about how our short courses impact the greater research community that uses HPC resources. It is currently and will always be necessary to develop new training for the latest hardware and software used in HPC [6]. We need a variety of metrics to help us understand better where to focus our efforts moving forward.

7 CONCLUSIONS

In discussing the metrics involved in the cyberinfrastructure training program, we approach the topic from the standpoint that the results we obtained gave us ideas about how various factors may be related; in the relatively open ecosystem of a university–plus external outreach–there is a limit to how specific our conclusions may be regardless of how much data we are able to collect, but we can at least identify trends to explore. We can suggest theories and sometimes gather data to test that theory, but we often must simply operate on our conjectures [9].

> The key concept is not that the metrics themselves had a direct causal effect on eventual outcomes, but rather that the metrics were chosen so that actions and decisions which moved the metrics in the desired direction also moved the organization in the direction of the desired outcomes and goals (p. viii) [14].

This quote defines the purpose of metrics in cyberinfrastructure training programs. In general, the goals are to serve the research community by providing various training resources to promote their research activities on HPC resources for the good of society.

Thus, adaptation, innovation, and assessment are in constant movement as technology quickly develops continually. We do our best to use the metrics to understand and conjecture about what will be needed next and work to develop the resources needed to support that work.

ACKNOWLEDGEMENTS

We gratefully acknowledge support from the National Science Foundation CyberTraining award CIP-1730695.

REFERENCES

[1] D. Christopher Brooks, Jearl Walker, and Paul Baepler. 2016. *A Guide to Teaching in the Active Learning Classroom: History, Research, and Practice.* Stylus Publishing.

[2] Dhruva K. Chakravorty, Marinus 'Maikel' Pennings, Honggao Liu, Zengyu 'Sheldon' Wei, Dylan M. Rodriguez, Levi T. Jordan, Donald 'Rick' McMullen, Noushin Ghaffari, and Shaina D. Le. 2019. Effectively Extending Computational Training Using Informal Means at Larger Institutions. *Journal of Computational Science Education* 10, 1 (2019), 40–47. https://doi.org/10.22369/issn.2153-4136/10/1/7

[3] Dhruva K. Chakravorty, Marinus 'Maikel' Pennings, Honggao Liu, Zengyu 'Sheldon' Wei, Dylan M. Rodriguez, Levi T. Jordan, Donald 'Rick' McMullen, Noushin Ghaffari, Shaina D. Le, Derek Rodriquez, Crystal Buchanan, and Nathan Gober. 2019. Evaluating Active Learning Approaches for Teaching Intermediate Programming at an Early Undergraduate Level. *Journal of Computational Science Education* 10, 1 (2019), 61–66. https://doi.org/10.22369/issn.2153-4136/10/1/10

[4] Dhruva K. Chakravorty, Lisa M. Perez, Honggao Liu, Braden Yosko, Keith Jackson, Dylan Rodriguez, Stuti H. Trivedi, Levi Jordan, and Shaina Le. 2021. Exploring Remote Learning Methods for User Training in Research Computing. *Journal of Computational Science Education* 12, 2 (2021), 11–17. https://doi.org/10.22369/issn.2153-4136/12/2/2

[5] Dhruva K. Chakravorty and Minh Tri Pham. 2022. Evaluating the Effectiveness of an Online Learning Platform in Transitioning Users from a High Performance Computing to a Commercial Cloud Computing Environment. *Journal of Computational Science Education* 11, 1 (2022), 93–99. https://doi.org/10.22369/issn.2153-4136/11/1/15

[6] Eric Coulter, Jodie Sprouse, Resa Reynolds, and Richard Knepper. 2019. Extending XSEDE Innovations to Campus Cyberinfrastructure - The XSEDE National Integration Toolkit. *Journal of Computational Science Education* 10, 1 (2019), 16–20. https://doi.org/10.22369/issn.2153-4136/10/1/3

[7] Zhenhua He, Sandra Nite, Joshua Winchell, Abhinand Nasari, Hieu Le, Jiao Tao, Dhruva Chakravorty, Lisa M. Pereze, and Honggao Liu. 2023. Development of a training framework for novel accelerators. In *Proceedings of the 2023 IEEE Frontiers in Education (FIE) Conference: Engineering Education in a Diverse, Global World.* College Station TX. https://doi.org/10.1109/FIE58773.2023.10343498

[8] Richard Lawrence, Tri M. Pham, Phi T. Au, Xin Yang, Kyle Hsu, Stuti H. Trivedi, Lisa M. Perez, and Dhruva K. Chakravorty. 2020. Expanding Interactive Computing to Facilitate Informal Instruction in Research Computing. *Journal of Computational Science Education* 13, 1 (2020), 44–49. https://doi.org/10.22369/issn.2153-4136/13/1/8

[9] Miriam Madsen. 2021. The Configurative agency of metrics in education: A research agenda involving a different engagement with data. *Journal of Education Policy* 36, 1 (2021), 64–83. https://doi.org/10.1080/02680939.2019.1682679

[10] Michael Prince. 2004. Does active learning work? A review of the research. *Journal of Engineering Education* 93, 3 (2004), 64–83. https://doi.org/10.1080/02680939.2019.1682679

[11] Kem Saichaie, D. Christopher Brooks, Phil Long, Bob Smith, Richard Holeton, Carole Meyers, Adam Finkelstein, and Shirley Dugdale. 2017. 7 Things You Should Know About Research on Active Learning Classrooms. Educause.

[12] Jinsil Hwaryoung Seo, Michael Bruner, Austin Payne, Nathan Gober, Donald 'Rick' McMullen, and Dhruva K. Chakravorty. 2019. Using Virtual Reality to Enforce Principles of Cybersecurity. *Journal of Computational Science Education* 10, 1 (2019), 81–87. https://doi.org/10.22369/issn.2153-4136/10/1/13

[13] Texas A&M High Performance Research Computing. 2024. https://hprc.tamu.edu

[14] Vandana, Ruchi Jain Garg, Vinod Kumar, and Alok Kumar Singh. 2023. Investigating the Impact of Usage Factors on Satisfaction and Intention to use e-Resource. *Performance Measurement and Metrics* 24, 1 (2023), 1–11. https://doi.org/10.1108/PMM-08-2021-0042