The Use of High-Performance Computing Services in University Settings: A Usability Case Study of the University of Cincinnati's High-Performance Computing Clusters

Mahmoud Junior Suleman, MSc.
University of Cincinnati
sulemamr@mail.uc.edu

Amy Latessa, PhD.
University of Cincinnati
latessak@mail.uc.edu

Shane Halse, PhD.
University of Cincinnati
halsese@mail.uc.edu

Jess Kropczynski, PhD.
University of Cincinnati
kropczjn@mail.uc.edu

ABSTRACT

High-performance computing (HPC) clusters are powerful tools that can be used to support a wide range of research projects across all disciplines. However, HPC clusters can be complex and difficult to use, limiting their accessibility to researchers without a strong technical background. This study used a mixed method to investigate ways to make HPC clusters more accessible to researchers from all disciplines on a university campus. A usability study of 19 university researchers was conducted to understand the needs of HPC users and identify areas where user experience could be improved. Our findings reveal the need to build a customized graphical user interface HPC management portal to serve users' needs and invest in workforce development by introducing an academic credit-based High-Performance Computing Course for students and partnering with other faculties to introduce special programs, e.g., Student Cluster Competitions which would draw more student interest.

KEYWORDS

High Performance Computing, Computing Resources for Academic Researchers, Usability, Think Aloud, Accessibility

1 INTRODUCTION

High-performance computing (HPC) has become an essential tool to advance scientific discovery over the last two decades, and it's an area where researchers actively create larger systems to accommodate new modes of scientific discovery with complex workflows. Advocating for large-scale scientific programming and HPC have become more essential to achieving national goals considering the discoveries made by academic researchers and industry professionals who contribute significantly to national development and further increase the importance of adequately educating the next generation [14]. National labs, academic institutions, and industries need scientists and staff who understand high-performance computing (HPC) and the complex interconnections across individual topics in HPC. However, domain science and computer science undergraduate programs need to provide more educational resources and are far from conveying the interdisciplinary and collaborative nature of the HPC environment [10].

Academic institutions that are research (R1) based invest resources to acquire and build centers to manage HPC Systems on campus to enable scientists to improve and foster their research capabilities. Over the past few years, some HPC centers have encountered challenges in encouraging scientists to use HPC resources effectively for their research. There is an HPC expertise and knowledge gap because very few educators have the skill set to use HPC Systems made available on campus, making it difficult for non-stem students to learn and access available HPC resources [17].

HPCs are used across various fields, making it challenging for students to understand their applications. Launching HPC applications is complex, requiring multiple components and specialized skills [30]. Using HPCs effectively demands a broad knowledge base and significant practice. Typical HPC curricula often focus on STEM students, covering topics like computer literacy, programming, parallel computing, version control, and debugging [31]. Students from non-computing departments need extra training to use HPCs effectively, hindering their research progress. To improve the accessibility and usability of local HPC resources, we conducted a usability study to determine the factors hindering their adoption by students and faculty. Toward this goal, this paper aims to address the following overarching research question: *How can HPCs be made more accessible for use across disciplines in institutes of higher education?*

2 BACKGROUND

2.1 Evolution of High-Performance Computers

Over the years, there has been rapid growth in computing and communications technology; the past decade has witnessed a proliferation of robust parallel and distributed systems and an ever-increasing demand for the practice of high-performance computing (HPC). HPCs have moved into the mainstream of computing. They have become a key technology in determining future research and development activities in many academic and industrial branches. They must cope with very tight timing schedules when solving large and complex problems [11].

https://doi.org/10.22369/issn.2153-4136/16/1/5

In 1985, the National Science Foundation established a partnership between five research centers: the San Diego Supercomputer Center (SDSC) at the University of California San Diego, the Pittsburgh Supercomputer Center (PSC) at the University of Pittsburgh, the National Center for Supercomputing Applications (NCSA) at the University of Illinois Champagne-Urbana, the Cornell Theory Center at Cornell University [29]. In the last decade, there has been significant advancement in the reliability and performance of computing elements such as processors, disks, and network devices. Computational power has increased in desktops and laptops by the availability of processors; these reliable and robust off-the-shelf computational elements have also spurred a new generation of high-performance computing systems [4].

2.2 The Value of High-Performance Computing Centres on University Campuses

Cyberinfrastructure includes computing systems, data storage, instruments, repositories, visualization, and people, connected by high-speed networks for scholarly innovation as defined by Stewart et al. [20]. Research shows a campus supercomputer positively impacts research output; Apon's work analyzed the ROI of cyberinfrastructure, finding that HPC investment benefits research productivity in fields like Chemistry, Physics, and Civil Engineering [1]. Apon's studies show statistical analyses of cyberinfrastructure's impact [2]. Indiana University calculated the ROI of 3 cyberinfrastructures, finding cost savings compared to commercial cloud alternatives [19]. Scrivner developed XDMOD-VA, a visualization plug-in for XDMOD, to present value proposition metrics for HPC centers [18]. This demonstrates the cost-effectiveness of supercomputers for most educational institutions.

2.3 Usability of High-Performance Computing in STEM and other disciplines

High-performance computing (HPC) is increasingly essential in higher education, HPCs offer new opportunities but pose challenges in teaching. Educators and researchers need to adopt innovative teaching approaches to leverage HPC effectively. Computing Education faces new demands, requiring motivated students and adaptive learning systems like SAIL- a System for Adaptive Interest-based Learning [9]. The Higher Education System needs comprehensive training to address the limitations of integrating technology into teaching and the constraints of using supercomputers. HPC is crucial for STEM and other disciplines despite challenges like limited experience and high costs [6]. Supercomputing education fosters innovation and creates a new generation of professionals. It's essential for various fields that use computational tools. Many educational institutions are integrating computer science concepts into basic training, raising questions about the importance of supercomputing education for students' careers [8]. ICT integration is crucial for modernizing education and transitioning to a knowledge society. Supercomputers will be vital for solving complex data problems in STEM and non-STEM disciplines [6].

2.4 Demystifying High-performance computing for non-Linux users: Introducing Remote computing desktops.

Implementing web-based portals for end users to manage and use HPCs has been remarkable. It has improved users' time navigating around HPCs to complete a task. This is not only limited to the time HPC centers spend training new users, especially non-Linux users, to adapt to the Command Line Interface (CLI) to complete their tasks. Not only have end users been relieved, but system admins and research computing facilitators have also been relieved; these portals give them a one-stop shop to manage these clusters and easily support end users. In this era of computing, more emphasis on graphical and intuitive interfaces is essential to clear barriers away from researchers from all respective fields to use computing resources with ease and not deal with rigid legacy system designs. Mastering the CLI is an essential skill, but it becomes intimidating for researchers who want to use the High-Performance computing resources of a university [23].

According to [25], Over 75% of users say that desktop services are either moderately or extremely important for their ability to use HPC resources. During training sessions at Indiana University, they observed that researchers who were constant Linux users who did not know about HPCs were enthusiastic about learning and trying out the HPC system and wanted to know more about its capabilities and features. However, they mostly complained about the Job Schedulers, e.g., SLURM and longer wait times for their jobs to be in the queue. Non-Linux researchers/users have more reservations about using HPC systems; most had no choice but to learn the command line and how to eventually use the job schedulers due to the nature of their research and the size of their data sets. In Summary, most HPC Users were demotivated to use the HPC system for their research. This concluded that the command line and the batch job scheduling are significant barriers for potential HPC users [25].

2.4.1 Web-based Portals for Managing High-Performance Computing Clusters. Web-based portals for HPCs were introduced in the 1990s with Java applets, making it easy to access HPC resources remotely using a variety of use cases. This identified the key features and functional and non-functional requirements of these portals [5]. Web-based management portals aid in reducing barriers that limit the adoption of high-performance computing (HPC), which addresses an important challenge affecting the STEM community. Through the Cyber-Infrastructure grant, the NSF has supported Higher Education institutions and National Supercomputing Labs to develop web-based HPC management portals to help researchers with little computer science skills adapt and foster the use of HPCs in their research [26].

Some benefits of having Web-based portals include but are not limited to:

(1) Increasing the use of HPC resources among disciplines that are not well represented in the community but desire the need for HPC resources.
(2) It is also beneficial to HPC stakeholders because it is easier to adopt advanced features to monitor and visualize system needs and upgrades.

(3) This creates a working model for training and workforce development to expand the cyberinfrastructure space for computer science. In 2016, the first version of Open On-demand launched the Open On-demand project, funded through the NSF Cyberinfrastructure grant ([27]) to make accessing HPC resources easier. The Open on Demand Project was an Open-source software project hosted by the Ohio Super Computing Center that enables HPC centers to install and deploy advanced web and graphical interfaces for their users [12].

2.5 Case Study: Advanced Research Computing Centre at the University of Cincinnati

The University of Cincinnati's Advanced Research Computing (ARC) Centre offers a readily accessible hybrid CPU/GPU computing cluster, supporting computational and data science researchers while developing a highly competitive workforce amongst the university community.

[3]. ARC partners with researchers to utilize the core HPC services and resources suite. With ARC's resources, researchers advance theoretical knowledge and expand the realm of discovery, generating leading-edge research and applications suitable for innovation and commercialization in line with the University of Cincinnati's "Next Lives Here" strategic direction. The center has a sustainable high-performance computing (HPC) infrastructure with technical support that leverages HPC services to accelerate the time to discovery and enables sophisticated and increasingly realistic modeling, simulation, and data analysis, which helps to bridge users to the local, regional, and national HPC ecosystem [3].

3 METHODS

To understand user needs at the High-Performance Computing Center, we combined methods used by [16]. We conducted an online survey, qualitative think-aloud interviews, and a heuristic evaluation. These methods provided comprehensive feedback on user needs and design effectiveness. The survey and interviews collected user insights, while the heuristic evaluation assessed adherence to usability principles. All methods were IRB-approved.

3.1 Research Participants

The selection of participants in this research was guided by the study's nature and desired outcomes, as outlined in [22]. Our target population for the survey were students and faculty members of the University of Cincinnati, especially current and future non-Linux researchers who intend to use or use the resources of the High-Performance Computing cluster.

The targeted survey audience was the Science, Technology, Engineering, and Mathematics (STEM) students and faculty members. We chose not to limit our audience to a particular group and included all potential users who were interested in using the cluster regardless of their technical skills or background. This diversified our participants, which gave us more information to feed into our future design.

Participants were recruited in two stages or categories:

- **Current University Researchers:** HPC researchers at the University of Cincinnati, both current and potential users.
- **Students (Undergraduate, Masters, PhD):** STEM students from various departments, including Engineering, Information Technology, Medicine, Art, and Planning, as well as non-STEM students.
- **Faculty:** Faculty teaching courses requiring HPC resources and principal investigators leading research labs.
- **Staff:** University staff using or potentially using HPC resources for departmental work.

3.2 Data Collection

For quantitative data, a 27-question survey was shared with participants by casting a wide net to all the users of the cluster, either past or present. These questions were carefully designed to capture feedback from highly skilled and novice users with little skills for genuine feedback on their needs and issue severity.

For Qualitative data, we analyzed and interpreted data, identified themes, and understood the study's phenomenon. We collected cognitive and constructive feedback through a think-aloud activity. [21].

3.2.1 *Think-Aloud Activity*.

The method of collecting think-aloud data involves participants spontaneously verbalizing words that come to mind as they complete an activity [28]. This makes it essential to conduct usability testing, as it allows users to verbalize their thinking as they use a new system. This allows evaluators to infiltrate the minds of users and acknowledge individual differences.

To further understand users' needs in-depth, We selected five users out of the respondents of the Survey who voluntarily opted in, three of whom were super users, and two non-superusers. We opened the conversation for users to have cognitive thinking, which allowed them to expand their thoughts while using the portal; we further asked questions in two categories, for Superusers and Non-Superusers:

- **Superusers** were asked what connection methods to the cluster they preferred and whether they needed an intuitive Graphical User Interface portal. Also, what were their needs for the graphical user interface portal? Furthermore, what could be done to improve the usability of the portal?
- **Non-Superusers**, mostly Non-Linux users, were asked about their concerns with the HPC portal and what they expected from the HPC cluster to avoid using the CLI.

In order not to influence users' thoughts and have their ideal participation while offering their honest thoughts spontaneously, we did not prompt users during the conversation; however, to avoid too much silence or stiltedness, we had signs with important information, e.g., keep talking, Go on, that is interesting, etc. For pre-orientation and housekeeping, we dedicated the first fifteen minutes of the session to explain how the think-aloud activity would be facilitated, and an overview of the 'think-aloud' method was provided in a pre-interview facilitator script [28].

Our participants were users who were allowed to reflect more on the best options that worked for them when using the cluster and why they preferred them. After that, we focused on the features of each connection method to understand their perceptions and whether they had HPC experience from their previous organizations or personal use. We also observed and noted their behaviours in real time.

3.2.2 Heuristic Evaluation. In 1990, a published article by web usability pioneers Jakob Nielsen and Rolf Molich defined Heuristic Evaluation as a set of principles used in "Improving a Human-Computer Dialogue" [15]. Heuristics evaluation is a process to systematically determine or certify a design or product's usability; researchers review the product's interface and compare it with other usability principles, and the results are accompanied by recommendations to improve the current system [7].

We based our Heuristic evaluation on the ten most fundamental principles that were given by the Nielsen Norman Group, which are listed in 1.

Figure 1: Nielsen and Molich's 10 User Interface Design Heuristics [7].

A group of examiners evaluated the University of Cincinnati's High-performance Computing cluster interface's user interface according to a set of heuristics rules based on 1.

We observed users' actions and interpreted their connection to usability challenges. We recorded the evaluation results and the evaluator's comments about the user interface; however, there was no need to interpret the evaluator's actions [**?**]. We provided a severity rate scale of 5 for each violation to measure the severity of design problems; at the end of the heuristic evaluation, we rated these design problems accurately to reflect the responses and evaluation of our participants. The schema for the severity ratings is listed below :

- **0** - don't agree that this is a usability problem
- **1** - cosmetic problem
- **2** - minor usability problem
- **3** - Major usability problem; important to fix
- **4** - usability catastrophe; imperative to fix

3.2.3 Data Analysis . Due to the nature of our study, we employed a mixed methods approach to collecting data, which involved both qualitative (survey) and quantitative (Think-aloud activity) methods. According to [24] mixed methods offer confirmation and complementarity feedback that "enables the researcher to simultaneously answer confirmatory and exploratory questions, and therefore verify and generate theory simultaneously. We found similarities between topics discussed by the participants, which were generated as themes by performing a Qualitative Thematic analysis.

We employed a mixed methods approach to collecting data and focused more on participant engagements, which influenced our research: both positive and negative feedback, participants describing experiences with using the HPC Cluster during our conducted aloud activity, and ideas to design an intuitive, customized graphical user Interface.

After manually reviewing collected data, which included surveys, transcribed audio, and videos, the data were categorized into codes and potential themes. For thematic data analysis by [13], we grouped the themes into broad sub-themes in relation to the research questions and analyzed the data with regard to the identified themes related to the research questions.

4 RESULTS

In this section, we present our research findings and in-depth analysis of participant recommendations.

4.1 Description of Users and Results from the Survey

The survey had 21 responses; two were voided, and the remaining 19 responses were from 11 super users and 8 non-super users. Most users were graduate students, followed by faculty and staff. Superusers had more experience with HPC clusters than non-superusers. Most users used the cluster once or twice a week for research projects. Key reasons for using the cluster included research projects, personal projects, grant requirements, thesis/dissertation, and class projects. Accessibility, data security, and cost-effectiveness were major factors in choosing the on-campus cluster. 47% of respondents didn't need UNIX/LINUX training, while 38% did. 42% used FTP for data transfer, 25% used an API, and 16.5% used Globus. Details about the participant demographics are shown in Table 5.

4.2 Identified Codes and Themes from the Think-Aloud activity

For our Thematic Coding analysis, as suggested by [16], we analyzed data from the Think aloud activity and created a thematic coding scheme to identify common patterns in the experiences of

Table 1.1: Participant Demographics and Descriptive Reported by Expertise

Descriptor	All Participants	Super Users	Novice Participants
N	19	11	8
No. of Years at University			
1 – 3	9	2	5
3 – 5	4	3	3
5 – 10	6	6	0
Role at University	*n* (%)	*n* (%)	*n* (%)
Faculty Researcher or PI	3 (1.5)	3 (10.0)	0(0)
Masters Student	4 (2.2)	1 (2.5)	3(7.5)
PhD Student	11 (5.78)	7 (6.3)	5(3.7)
Staff	1 (0.52)	1 (10.0)	0(0)
Previous Use of HPC	*n* (%)	*n* (%)	*n* (%)
Never Used	0(0)	0 (0)	0(0)
Seasonal User (several uses total)	1 (0.54)	1 (10.0)	0(0)
Once per Month	3 (1.59)	2 (6.3)	1(3.7)
Two or Three Times per Week	10 (5.26)	5 (5.0)	5(5.0)
Everyday	5 (2.6)	3 (6.0)	2(4.0)
Type of Project			
Personal Project	4	3	1
Class Project	1	0	1
Thesis or Dissertation	4	3	1
Internally Funded Research	7	3	4
Externally Funded Research	3	2	1

Figure 2: Participant Demographics and Descriptive Reported by Expertise

our participants using the HPC Cluster; after identifying key themes and concepts mentioned by our participants via the interview, we searched for patterns to in these codes after reaching saturation point. We identified key themes/codes from our participants, which we explained in the subsections below. There were five interviewees for the Think-aloud activity, and to be more specific in identifying them, we referenced participants as P1, P2, P3, etc. Three superusers and two non-supers users participated in the Think-aloud activity. However, to incorporate qualitative responses from our survey participants, we referenced them as S1, S2, S3, etc.

4.2.1 Theme 1: Connection Methods to the HPC Cluster . For theme (1), we explored user preferences regarding connection methods (e.g., CLI, SSH, Open On-Demand) to the HPC Cluster. Generally, users relied on instructions provided by the Advanced Research Computing Center (ARCC) [3] through their website or the Research Computing Facilitator. The Facilitator served as a crucial resource for users facing challenges or requiring direct information about the cluster (e.g., specific software, GPU/CPU allocations, job queues).

This reliance is exemplified by the following comment from a participant (P1): *"It was OK because I think I was following some of the instructions. You see, IT had some instructions on that. About like first users. So it was. I had thought that that was pretty smooth."*

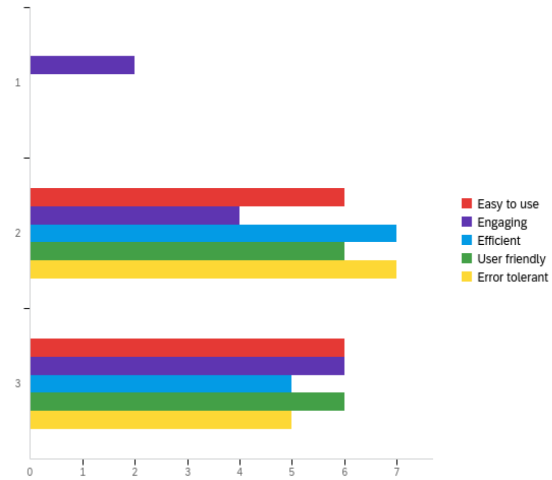

Figure 3: A Visualized report of respondents ratings of the current HPC Web Portal

Users found the Advanced Research Computing Center's website provided sufficient instructions for connecting to the cluster. On-campus connections were generally smooth due to the University's network resources. However, off-campus access via VPN was disliked by some (P1, P2, P3) due to perceived hindrance to collaboration, despite its security benefits.

All users favored the Open OnDemand interface. Non-superuser users (P3, P4) relied on it exclusively due to a lack of UNIX/Linux skills required for CLI or SSH access. Open OnDemand also streamlined file transfers and provided access to necessary software (e.g., Matlab for P1).

Both superusers and non-superusers appreciated Open OnDemand's intuitive interface, particularly for first-time or less experienced users. The quarterly *Introduction to HPC* training was also highly valued, serving as an onboarding tool for superusers and providing basic command and cluster usage knowledge for non-superusers.

4.2.2 Theme 2: Jobs Management (Job Scheduling, Queue Management and Execution of Jobs) . Job submission and queue management are essential steps when using the HPC Cluster. All jobs on the cluster would achieve the desired results based on the accuracy and correctness of your Job submission file. The Primary tool for Job Scheduling, queue management and execution of Jobs on the HPC Cluster is called "The Slurm Workload Manager." SLURM takes some information about the requirements of resources and sends these calculations to the compute nodes that run the jobs to satisfy the criteria or requirements. It ensures resources are allocated to all users fairly and based on priority.

Two ways to submit jobs on the HPC Cluster currently are using SLURM via the Command Line Interface "CLI" or ARCC Desktop on Open On-Demand.

Our participants (P1-P4) appreciated using the ARC Desktop through the Open on-Demand (Interactive Desktop) portal, which made it easy and accessible to submit their jobs by quickly uploading your Python file. This also allows users to select their preferred resources, (e.g., Number of hours, Number of nodes and partitions); it also has an email notifications option for when your session starts if selected. Open On-demand has a File Transfer button, Job Management, shell access and the ARCC Desktop.

SLURM is an excellent tool for Job scheduling, queuing and submission; However, our participants, both superusers and non-superusers, appreciated working with an interactive graphical user interface desktop as expressed in the following quotes from study participants:

P4 *"I think using the SLURM platform is a little bit challenging for me as a beginner. The Open On-Demand is more helpful and straight forward."* and *"I think the On-Demand platform helped me a lot when I used the advanced research computing for my dissertation."*

P3 *"Too difficult to or like I can still run things and get an output, so it's so that so I put it in the middle.Helpful, but it was still usable, so it wasn't too difficult to use. I think it was.Sometimes it does seem that something more interactive could be. Well, like, I feel like I'd be in the middle. Like it wasn't too difficult, but I feel like there are some aspects about like, troubleshooting that could be easier. I with the way that I've been studying up. So I kind of just like send a job and then."*

4.2.3 Theme 3: Navigating the HPC Web portal . The core reasons users choose the web portal to access the cluster other than other preferred methods were ease of access, efficiency and optimized performance. Our results from the Think-aloud activity showed that both Super and Non-superusers preferred to use the interactive graphical user interface portal since it came with File Transfer, Job Management, shell access and the ARCC Desktop features. This made it very accessible for users to interact with the portal for their activities.

Though it was evident users preferred the Interactive Desktop, the current Graphical user interface had some cosmetic issues our participants identified and recommended to make the portal more accessible and usable as demonstrated through the following representative quotes on this theme:

P4 *"You know, Welcome to the ARC?, you know follow these steps one upload your data too you know and then this is how to use open on demand. I think just having that on this front page there's other things I think would make it better but."*

P3 *"I'm using about the same amount of cores every time, so I'm typically it's like I'm. I'm running the same setup every time, so it seems like going in and then having to type in the stuff I usually have.* and *"Visual Studio code or Mathematica through a portal on open on demand I there's a I'd say there's a decent likelihood I would use it."*

4.3 Results from Heuristic Evaluation

A heuristic evaluation was conducted to assess the severity of usability issues in the Interactive web-based Graphical User Desktop for managing the HPC Cluster. The evaluation identified specific violations directly impacting functionality and general violations related to visual and functional design that could lead to usability problems. The violations and their severity ratings are presented in Figure 5.

In total, 18 problems with a severity rating of 0 or higher were found, with 8 having a severity of 2 or higher. The main issues involved complex connection methods and user interface design problems, such as large blank spaces and the lack of an intuitive job status button. Other identified issues were minor cosmetic problems.

Figure 4: A Text Analysis of Keywords from the Quantitative Data

Table 1.2: Results of Heuristic Evaluation.

Violations by Task	Heuristic	Severity
Task 1 Going through long procedures to get an account set up	User Control and Freedom	2 -minor usability problem
Task 2 Users reported difficulty knowing the status of a job and whether an error occurred after submission	Visibility of System Status	4 -usability catastrophe; imperative to fix
Task 3 No violations	none	N/A
Task 4 we found violations with page alignments, consistency in terms of alignments and fonts	Atheistic and Minimalist Design	1- Cosmetic problem
Task 5 we found violations were users didn't know either if their submitted jobs were running or an estimated completion time	Flexibility and efficiency of use	4 - usability catastrophe; imperative to fix
Task 6 Users could not edit their already submitted jobs	User Control and Freedom	2 - minor usability problem
Task 7 Landing page underutilized, could contain tips or documentation section	Help and Documentation	1- cosmetic problem
Task 8 Users not able to use other software installed on the Cluster on the web portal	Consistency and Standards	4- usability catastrophe; imperative to fix
Task 9 Users wanted a more intuitive and interactive web portal	Aesthetic and minimalist design	3 - major usability problem; important to fix

Figure 5: Results of Heuristic Evaluation.

5 RECOMMENDATION AND DISCUSSION

Our results show that users of the HPC Cluster managed by the Advanced Research Computing Center were satisfied with the services provided by ARCC, as shown in Figure 3. This further proved that

HPCs played a significant role in fostering the research capabilities of researchers at the University of Cincinnati. Although there were multiple alternatives for users to connect to the clusters, Open on-Demand was the preferred connection method for job submission, file transfer, and managing job queues.

To address our research question of how HPCs can be made more accessible across disciplines in institutes of higher education, specifically at the University of Cincinnati as our case study and also providing standardized HPC tools, we recommend taking the following actions:

- **Develop a customized graphical user interface (GUI) web HPC management portal**. Considering the research and user environment, this would be tailored to users' needs at the University of Cincinnati. Open On-Demand is an excellent HPC portal, but it is a generic system that other HPC centers use to manage their clusters. However, designing an accessible and customized GUI portal for the University of Cincinnati would improve users' needs and make their use of the High-Performance Computing cluster more accessible.
- **Allow the HPC to be internet-facing**. This would allow users to connect to the cluster off-campus, benefiting researchers who need to access the cluster from home or other locations. Also, this goes further to enable researchers, especially industry companies that require the cluster's services, to integrate third-party systems, e.g. APIs, to improve their research needs. The clusters can be more secure by implementing 2-factor authentication methods, e.g., the current DUO system used university-wide to serve security needs.
- **Workforce Development** To increase workforce development, ARCC should increase the frequency of training sessions and introduce Advanced Level and speciality training for users. Users want to acquire academic credits when these courses are taken, which we believe will encourage novice and frequent users to attend the training sessions. Training can be made mandatory regardless of your experience, especially new users, as an onboarding method for new users of the cluster; this would help to ensure that all users, irrespective of their experience level, have the skills they need to use the HPC cluster effectively.
- **Student Cluster Competition** ARCC should introduce the student cluster competition on campus, primarily partnering with the School of Information Technology (SoIT) and Engineering Faculty during the annual expo; this would be a means of encouraging users to draw more interest to using the High-Performance Computing Cluster.

Feedback from our survey respondents gave a general picture of the current usage of the HPC Cluster. Graduate students from departmental or research labs use the Cluster for their dissertations and funded research. This proves the importance of designing a customized HPC web portal for the University of Cincinnati to make the cluster more accessible for researchers across all disciplines. This would make onboarding and usage of the cluster very easy and direct. Training plays a pivotal role in onboarding users with no experience and users with some experience with Linux/UNIX. However, though some users have experience with Unix, using a cluster requires some special skills, which is beneficial for all users. Frequent training and introducing special programs like the "Student Cluster Competition" will be an excellent initiative for ARCC to solidify its name across the University of Cincinnati ecosystem.

6 LIMITATIONS AND FUTURE WORK

Our primary limitation was participant recruitment. Many super-users were unavailable for various reasons, and non-super-users were often hesitant about being recorded on video while performing basic tasks. Greater incentives might be necessary to encourage user participation in future data collection. In future work, we would expand our research criteria and participant pool to several universities in various regions to encompass the broader HPC community. This broader focus would support our goal of making HPCs accessible to researchers across all disciplines.

7 CONCLUSION

In this research, we conducted a usability study on how to make HPCs accessible to users across all disciplines, and our case study involved the users of the Advanced Research Computing Center at the University of Cincinnati. Our results highlight several factors that back the need for the Advanced research computing cluster to build a customized, intuitive HPC management web portal. Overall, our results indicated that both super and non-superusers preferred using the Web-based HPC Management portal for job submission and queue management, and implementing our recommendation and future work would improve the ease of access to the cluster and increase the confidence of our non-superusers since the portal would make it easy for them to manage the cluster and implement several technologies through a web-based portal. A critical contribution from this study to workforce development is to expand the training sessions, which have been very instrumental to many users; this training is a resource to a lot of first-time users, especially non-superusers and a few of our respondents, who were on-boarded through the training and have been using the cluster which has made them superusers. This is a suitable means to encourage the use of HPCs across all disciplines, which addresses the challenge in our research question.

REFERENCES

[1] Amy Apon, Linh Ngo, Michael Payne, and Paul Wilson. 2014. Assessing the effect of high performance computing capabilities on academic research output. *Empirical Economics* 48 (02 2014). https://doi.org/10.1007/s00181-014-0833-7

[2] Amy W. Apon, Linh B. Ngo, Michael E. Payne, and Paul W. Wilson. 2015. Assessing the effect of high performance computing capabilities on academic research output. *Empir. Econ.* 48, 1 (Feb. 2015), 283–312. https://doi.org/10.1007/s00181-014-0833-7

[3] Advanced Research Computing at UC. 2019. Advanced Research Computing at UC. https://arc.uc.edu/citing-the-arc/

[4] Michael A. Bauer. 2007. High Performance Computing: The Software Challenges. In *Proceedings of the 2007 International Workshop on Parallel Symbolic Computation (PASCO '07)*. Association for Computing Machinery, New York, NY, USA, 11–12. https://doi.org/10.1145/1278177.1278180

[5] Patrice Calegari, Marc Levrier, and Paweł Balczyński. 2019. Web Portals for High-performance Computing: A Survey. *ACM Trans. Web* 13, 1 (Feb. 2019), 1–36. https://doi.org/10.1145/3197385

[6] Álvaro Fernández, Camino Fernández, José-Ángel Miguel-Dávila, and Miguel Á. Conde. 2021. Integrating supercomputing clusters into education: a case study in biotechnology. *J. Supercomput.* 77, 3 (March 2021), 2302–2325. https://doi.org/10.1007/s11227-020-03360-5

[7] Interaction Design Foundation. 2021. What is Heuristic Evaluation?

[8] Michail N. Giannakos. 2014. Exploring Students Intentions to Study Computer Science and Identifying the Differences among ICT and Programming Based Courses. *Turkish Online Journal of Educational Technology - TOJET* 13, 4 (Oct. 2014), 36–46. https://eric.ed.gov/?id=EJ1043181

[9] Abid Haleem, Mohd Javaid, Mohd Asim Qadri, and Rajiv Suman. 2022. Understanding the role of digital technologies in education: A review. *Sustainable Operations and Computers* 3 (Jan. 2022), 275–285. https://doi.org/10.1016/j.susoc.2022.05.004

[10] Stephen Harrell, Hai Nam, Verónica Larrea, Kurt Keville, and Dan Kamalic. 2015. Student Cluster Competition: A Multi-disciplinary Undergraduate HPC Educational Tool. https://doi.org/10.1145/2831425.2831428

[11] HPCC 2021. 2021. The 23rd IEEE int'l Conference on HPC and Communications. Retrieved May 26, 2023] from http://www.ieee-hpcc.org/2021

[12] David E. Hudak, Douglas Johnson, Jeremy Nicklas, Eric Franz, Brian McMichael, and Basil Gohar. 2016. Open OnDemand: Transforming Computational Science Through Omnidisciplinary Software Cyberinfrastructure. In *Proceedings of the XSEDE16 Conference on Diversity, Big Data, and Science at Scale*. Association for Computing Machinery, New York, NY, USA, Article 43, 7 pages. https://doi.org/10.1145/2949550.2949644

[13] Barikisu Issaka. 2021. *Alienating: how the portrayal of Muslim women in US media affects Muslim women's social identities*. Ph.D. Dissertation. https://krex.k-state.edu/handle/2097/41509

[14] Peter Kogge, S. Borkar, Dan Campbell, William Carlson, William Dally, Monty Denneau, Paul Franzon, William Harrod, Jon Hiller, Stephen Keckler, Dean Klein, and Robert Lucas. 2008. ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems. *Defense Advanced Research Projects Agency Information Processing Techniques Office (DARPA IPTO), Techinal Representative* 15 (01 2008).

[15] Rolf Molich and Jakob Nielsen. 1990. Improving a human-computer dialogue. *Commun. ACM* 33, 3 (1990), 338–348.

[16] Svati Murthy, Jess Kropczynski, and Shane Halse. 2022. Understanding Decision-Making Needs of Open Government Data Users. https://doi.org/10.24251/HICSS.2022.331

[17] Rajendra K. Raj, Carol J. Romanowski, John Impagliazzo, Sherif G. Aly, Brett A. Becker, Juan Chen, Sheikh Ghafoor, Nasser Giacaman, Steven I. Gordon, Cruz Izu, Shahram Rahimi, Michael P. Robson, and Neena Thota. 2020. High Performance Computing Education: Current Challenges and Future Directions. In *Proceedings of the Working Group Reports on Innovation and Technology in Computer Science Education (ITiCSE-WGR '20)*. Association for Computing Machinery, New York, NY, USA, 51–74. https://doi.org/10.1145/3437800.3439203

[18] Olga Scrivner, Gagandeep Singh, Sara Bouchard, Scott Hutcheson, Ben Fulton, Matthew Link, and Katy Borner. 2018. XD Metrics on Demand Value Analytics: Visualizing the Impact of Internal Information Technology Investments on External Funding, Publications, and Collaboration Networks. *Frontiers in Research Metrics and Analytics* 2 (01 2018). https://doi.org/10.3389/frma.2017.00010

[19] Craig A. Stewart, David Y. Hancock, Wernert Julie, Matthew R. Link, Nancy Wilkins-Diehr, Therese Miller, Kelly Gaither, and Winona Snapp-Childs. 2018. Return on Investment for three Cyberinfrastructure facilities: a local campus supercomputer; the NSF-funded Jetstream cloud system; and XSEDE (the eXtreme Science and Engineering Discovery Environment). https://scholarworks.iu.edu/dspace/handle/2022/22590 [Online; accessed 26. May 2023].

[20] Craig A. Stewart, Stephen Simms, Beth Plale, Matthew R. Link, David Y. Hancock, and Geoffrey C. Fox. 2010. What is Cyberinfrastructure? *ACM* (Oct. 2010). https://scholarworks.iu.edu/dspace/handle/2022/12967

[21] Anna Sutton. 2016. Heasuring the Effects of Self-Awareness: Construction of the Self-Awareness Outcomes Questionnaire - PMC. Europe's journal of psychology. https://doi.org/10.5964/ejop.v12i4.1178

[22] Jane Sutton, AnnaSutton and Zubin Austin. 2015. Qualitative Research: Data Collection, Analysis, and Management. In *The Canadian journal of hospital pharmacy (vol. 68,3)*. The Canadian journal of hospital pharmacy, Article , numpages = 68,3,.

[23] High Performance Systems. 2023. HPC everywhere: A quick-access web portal for all things HPC-related at IU.

[24] Tasbakkori Teddlie, C. 2023. *A. (2003). Major Issue and Controversies in the Use of Mixed Methods in Social and Behavioral Sciences. In C. Teddlie, & A. Tasbakkori (Eds.), Handbook of Mixed Methods in Social and Behavioral Sciences (pp. 3-50). Thousand Oaks, CA SAGE. - References - Scientific Research Publishing*. Ph.D. Dissertation. https://www.scirp.org/(S(lz5mqp453edsnp55rrgjct55))/reference/referencespapers.aspx?referenceid=2220642

[25] Abhinav Thota, Le Mai Weakley, Ben Fulton, H. E. Cicada Brokaw Dennis, Laura Huber, Scott Michael, Winona Snapp-Childs, Stephen Lien Harrell, Alexander Younts, Daniel T. Dietz, Christopher Phillips, and Xiao Zhu. 2019. Research Computing Desktops: Demystifying Research Computing for Non-Linux Users. In *Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning) (PEARC '19)*. Association for Computing Machinery, New York, NY, USA, Article 54, 8 pages. https://doi.org/10.1145/3332186.3332206

[26] U.S. National Science Foundation. 2023. Frameworks: Software NSCI-Open On-Demand 2.0: Advancing Accessibility and Scalability for Computational Science through Leveraged Software Cyberinfrastructure. Retrieved May 13, 2023] from https://www.nsf.gov/awardsearch/showAward?AWD_ID=1835725 NSF Award Search - Award # 1835725.

[27] U.S. National Science Foundation. 2023. SI2-SSE: Open OnDemand: Transforming Computational Science through Omnidisciplinary Software Cyberinfrastructure. Retrieved May 13, 2023] from https://www.nsf.gov/awardsearch/showAward?AWD_ID=1534949 NSF Award Search - Award # 1534949.

[28] Maaike J. Van den Haak, Menno D. T. de Jong, and Peter Jan Schellens. 2003. Retrospective vs. concurrent think-aloud protocols: Testing the usability of an online library catalogue. *Behaviour & IT* 22, 5 (Sept. 2003), 339–351. https://doi.org/10.1080/0044929031000

[29] XSEDE. 2022. A Brief History of High-Performance Computing (HPC) - XSEDE Home - XSEDE Wiki. Retrieved 04/08/2022) from https://confluence.xsede.org/pages/viewpage.action?pageId=1677620

[30] Il-Chul Yoon, Alan Sussman, and Adam Porter. 2005. And Away We Go: Understanding the Complexity of Launching Complex HPC Applications. In *Proceedings of the Second International Workshop on Software Engineering for High Performance Computing System Applications (SE-HPCS '05)*. Association for Computing Machinery, New York, NY, USA, 45–49. https://doi.org/10.1145/1145319.1145333

[31] Jill Zarestky and Wolfgang Bangerth. 2014. Teaching High Performance Computing: Lessons from a Flipped Classroom, Project-Based Course on Finite Element Methods. 34–41. https://doi.org/10.1109/EduHPC.2014.10