

Evaluating the Effectiveness of an Online Learning Platform in Transitioning Users from a High Performance Computing to a Commercial Cloud Computing Environment

Dhruva Chakravorty
High Performance Research
Computing
Texas A&M University
College Station, TX, USA
chakravorty@tamu.edu

Minh Tri Pham
High Performance Research
Computing
Texas A&M University
College Station, TX, USA
phamminhtris@tamu.edu

ABSTRACT

Developments in large scale computing environments have led to design of workflows that rely on containers and analytics platform that are well supported by the commercial cloud. The National Science Foundation also envisions a future in science and engineering that includes commercial cloud service providers (CSPs) such as Amazon Web Services, Azure and Google Cloud. These twin forces have made researchers consider the commercial cloud as an alternative option to current high performance computing (HPC) environments. Training and knowledge on how to migrate workflows, cost control, data management, and system administration remain some of the commonly listed concerns with adoption of cloud computing. In an effort to ameliorate this situation, CSPs have developed online and in-person training platforms to help address this problem. Scalability, ability to impart knowledge, evaluating knowledge gain, and accreditation are the core concepts that have driven this approach. Here, we present a review of our experience using Google's Qwiklabs online platform for remote and in-person training from the perspective of a HPC user. For this study, we completed over 50 online courses, earned five badges and attended a one-day session. We identify the strengths of the approach, identify avenues to refine them, and consider means to further community engagement. We further evaluate the readiness of these resources for a cloud-curious researcher who is familiar with HPC. Finally, we present recommendations on how the

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Copyright ©JOCSE, a supported publication of the Shodor Education Foundation Inc.

© 2018 Journal of Computational Science Education
DOI: <https://doi.org/10.22369/issn.2153-4136/x/x/x>

large scale computing community can leverage these opportunities to work with CSPs to assist researchers nationally and at their home institutions.

CCS CONCEPTS

•CS→Computer Science; •Cybertraining→training on using cyberinfrastructure; •HPC→high performance computing

Keywords

HPC training, cloud computing, assessment strategies, best practices, diversity

1. INTRODUCTION

The growing computing needs of researchers in data science and engineering have led to increasing use of cloud computing resources, coupled with innovative web-based platforms for data analysis. This is observable in National Science Foundation (NSF) investments in projects such as JetStream [1,2], Aristotle [3], and CloudLab [4]. Indeed, in recent years, the research and engineering (R&E) community has adopted commercially available cloud resources and services (CACRS, also known as commercial service providers) in research and education. This is particularly true for the computational biology community, led by the National Institutes of Health (NIH) Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) initiative [5]. The NIH STRIDES initiative explores new models of data stewardship as it provides access to cloud services such as compute, storage and training on Amazon Web Services (AWS) and Google Cloud Platform (GCP). Indeed, STRIDES is on its way to hosting the tens of petabytes of NCBI data set on the cloud. In recent times, the NSF has also taken steps to offer researchers cloud-based resources. The NSF Big Data program [6] partnered with IBM, GCP, Microsoft Azure and AWS to make a number of cloud-based awards as well. Similar partnerships led to awards in the decade-long NSF CC* program in 2019 [7]. In perhaps a sign of the growing role of CACRS in R&E, the NSF launched the NSF Cloud Access program [8] that created a CSP-based cloud-exchange for NSF researchers requiring CACRs for their NSF-funded projects. [NSF Cloud bank] The scientific community has relied on high performance computing (HPC) to meet its large-scale scientific computing needs. Owing to the rapid proliferation of cloud computing in science technology engineering and mathematics (STEM) disciplines, migration to

CACRS requires training on both, how to configure and use these resources.

Online training platforms offer opportunities to scale that are beyond the capabilities of instructor-led, in-person classrooms. To meet the growing demand for trained and certified cloud- and data-science engineers, a number of online training platforms have emerged to fill the void between existing curricula at institutions of higher education and the urgent needs of these sectors. As such, informal online education platforms have taken the lead in preparing aspiring engineers for certifications and careers in these fields. In addition to Coursera [9] and Linux Academy [10], Qwiklabs is an online training platform that offers hands-on, lab-learning environments on using cloud-based services. At the time of this study, Qwiklabs, provided temporary credentials to Google Cloud Platform (GCP) and Amazon Web Services (AWS). A study of recent awards made by the NSF Big Data, CC* and E-CAS programs (in collaboration with Internet2) [11] suggest that AWS and GCP remain the platforms of choice for the large scale computing community. Furthermore, researchers (and faculty) may readily access Qwiklabs via free credits from generous programs offered by GCP. Furthermore, instructors specializing in in-person training events could use Qwiklabs to develop, host and broadcast their hands-on training labs and lecture notes. Taken together, these factors make Qwiklabs a strong candidate for an online platform that can provide existing resources to a cloud-hungry group of researchers. It should be noted that while Qwiklabs started as an independent effort, it was acquired by Google in 2016. [12,13] While GCP-related courses continue to be added, Qwiklabs continues to offer AWS training courses as well. Support for instructor-led training activities continues on this platform. In this paper, we explore the readiness this learning platform in transitioning HPC users to the commercial cloud.

2. EXPERIMENTAL DESIGN

At the onset, we acknowledge that proficiency in computing widely differs from student to student. Furthermore, with a wide range of computing constructs available, there is a need to define a standard for introductory, intermediate, and advanced, activities proficiencies in training. Teaching computing practices to HPC users presents additional challenges. Researchers have a diverse set of skills and varying degrees of experience using computing in their research. The Qwiklabs training platform helps address some of these concerns by including scaffolded instruction methods that support learners with varied skill sets. Research find that active learning is more effective than procedural training. [14-25] This approach is adopted across the platform. Since Qwiklabs is owned by Google, here we evaluate the efficacy of its GCP-centric labs in helping HPC system administrators and users adopt to the GCP environment.

2.1 The Qwiklabs Approach

Qwiklabs, as the name suggests, is compilation of quick labs that give participants hands-on experience on the “real” cloud. Qwiklabs leverage community expertise to develop these tutorials. These labs are web-based. As such, they are platform agnostic and can be run on any computer! These tutorials/labs themselves have an easy-to-comprehend format. All labs are templated and follow a similar format. While some labs are free, others require credits. A collection of labs on a single topic are combined to create a quest. Depending on the learning levels, quests are identified as Introductory, Fundamental, Advanced and Expert. Community experts typically create these labs, and some labs have an associated tutorial entry on github. By employing community

experts to create the labs, the curriculum on Qwiklabs can extend in complexity, while simultaneously supporting analysis across various STEM disciplines. While templating labs allows for a large number of contributors, the labs themselves may be of variable quality. Participants can grade a lab on completing it as a measure of feedback.

Informal effort that use well-reviewed pedagogical approaches to education have been found to encourage participation and adoption of computational thinking. [14-25] We have found that teaching students new computing concepts, such as navigating a Linux environment, using a command line, and writing code, is more productive when done through an interactive format. [26-29] With an interactive format, students are more motivated to follow along with the instructor and other students by participating in the activities. In agreement with existing literature, our data indicate that students with varying degrees of programing are best suited with scaffolded learning approaches like Jupyter notebooks for application specific training. [DKC citations] A problem-solving approach, though slower, encourages greater interactions and deeper learning of the subject matter. Including a scaffolded learning approach helps learners grasp complex concepts better, and helps reduce the barriers to computing enablement. This is particularly relevant for informal efforts supported by HPC centers that support users with a diverse range of research needs and computing prowess. The relatively inexpensive nature of these labs and the quests open possibilities for educators to couple activities with classroom and HPC training. Indeed, coupled with adequate assessment techniques, Qwiklabs can serve as a scalable platform on which students can develop prowess on GCP.

Here, we explore the usability of Qwiklabs, and the pedagogical approaches used to introduce users to cloud computing services and environments on GCP for R&E. In particular, we report on the use of employing scaffolded instruction, tiered instruction techniques and innovative active learning exercises in the context of teaching Cloud Computing via Quests and their associated labs. In this study, we do not expand on the challenges available in Qwiklabs. The programs goals were to assess the usability of these existing labs in (a) increasing participant engagement in cloud computing, (b) teaching relevant knowledge for current HPC users in order to help them migrate their workflows, and (c) providing participants with a learning environment that employed hands-on exercises.

2.2 Using Qwiklabs

The lead author on this publication (Chakravorty) has worked with HPC technologies since 2004 and did not use Qwiklabs prior to this exercise. Tri Pham has adopted HPC and cloud computing practices in recent years. The authors have completed over 50 Qwiklabs courses in 2019. These are listed in Table 1. The primary author was first introduced to the platform at the Google Next 2019 conference in San Francisco, where conference attendees were encouraged to explore Qwiklabs, with assistance provided by “Googlers.” Since then, the authors have completed various GCP labs online. To understand the tiered training approach, we completed learning labs in quests at the Introductory, Fundamental, Advanced, and Expert levels. In this study we did not participate in Challenge labs. In addition to taking the online labs described in Table 1, the author attended an introductory level in-person Google Cloud Platform Fundamentals class offered by a third-party company on behalf of Google. This session provided an overview of Platform products and services, and demonstrated how to incorporate GCP solutions into business strategies. In addition to providing an introductory overview of the GCP suite of products, the course covered topics related to storage, virtual machines,

containers and applications in the cloud. The in-person section also covered topics in developing, deploying, and monitoring processes in the cloud, as well as the use of machine learning and Big Data technologies in the cloud. In addition, the author also attended a session taught by a Googler that demonstrated how APIs can be integrated with G Suite for Big Data analysis.

Table 1. Completed courses and quests for GCP along with all labs associated with the quests and learning level are listed. Quests that are in progress are denoted with an asterisk (*). Only completed labs are reported for quests that are yet to be completed. A list of completed labs that are not associated with Quests are provided under the category of “Independent Labs Completed”.

Quest	Courses
GCP Essentials <i>Level: Introductory</i>	A Tour of Qwiklabs and the GCP
	Creating a Virtual Machine
	Compute Engine: Qwik Start - Windows
	Getting Started with Cloud Shell & gcloud
	Kubernetes Engine: Qwik Start
	Set up Network and HTTP Load Balancers
Google Cloud Platform Fundamentals: Core Infrastructure <i>Level: Fundamental</i>	Getting Started With Cloud Marketplace
	Getting Started with Compute Engine
	Getting Started with Cloud Storage and Cloud SQL
	Getting Started with Kubernetes Engine
	Getting Started with App Engine
	Getting Started with Deployment Manager and Stackdriver
Baseline Infrastructure <i>Level: Introductory</i>	Cloud Storage: Qwik Start - Console
	Cloud Storage: Qwik Start – CLI/SDK
	Cloud IAM: Qwik Start
	Stackdriver: Qwik Start
	Cloud Functions: Qwik Start - Console
	Cloud Functions: Qwik Start – Command Line
	Google Cloud Pub/Sub: Qwik Start - Console
	Google Cloud Pub/Sub: Qwik Start – Command Line
	Google Cloud Pub/Sub: Qwik Start - Python
Security & Identity Fundamentals <i>Level: Fundamental</i>	Cloud IAM: Qwik Start
	IAM Custom Roles
	Service Accounts and Roles: Fundamentals
	Install a Forseti Server on GCP
	VPC Network Peering
	User Authentication: Identity Aware Proxy
	Getting Started with Cloud KMS
	Setting up a Private Kubernetes Cluster
GKE Migrating to Containers	
Google Kubernetes Engine Best Practices <i>Level: Advanced</i>	Monitoring with Stackdriver on Kubernetes Engine
	Tracing with Stackdriver on Kubernetes Engine
	Logging with Stackdriver on Kubernetes Engine
	Connect to Cloud SQL from an Application in Kubernetes Engine
	Quiz: Kubernetes Engine Best Practices Quiz
Baseline: Data, ML and AI <i>Level: Introductory</i>	Cloud ML Engine: Qwik Start
	Dataprep: Qwik Start
	Dataflow: Qwik Start- Templates
	Dataflow: Qwik Start - Python
	Dataproc: Qwik Start - Console
	Dataproc: Qwik Start – Command Line
	Cloud Natural Language API: Qwik Start
Google Cloud Speech API: Qwik Start	
Video Intelligence: Qwik Start	
Cloud Architecture <i>Level: Fundamental</i>	Orchestrating the Cloud with Kubernetes
	Deployment Manager- Full Production
Kubernetes in the Google Cloud* <i>Level: Advanced</i>	Introduction to Docker
	Kubernetes Engine: Qwik Start
	Orchestrating the Cloud with Kubernetes
Windows on GCP* <i>Level: Expert</i>	Managing Deployments Using Kubernetes Engine
	Compute Engine: Qwik Start - Windows
	Deploy Microsoft SQL Server to Compute Engine
	Running Windows Containers on Compute Engine
G Suite: Integrations* <i>Level: Advanced</i>	Stackdriver: Qwik Start
	Introduction to APIs in Google
Independent Labs Completed <i>Level: Various</i>	Creating a Gmail add on
	Hardening Default GKE Cluster Configurations
	Google Kubernetes Engine Security: Binary Authorization
	Provision Services with GCP Marketplace
	Using Role-based access controls in Kubernetes Engine

3. RESULTS

An important consideration while evaluating this learning platform is that the author had self-selected himself to participate in these exercises, and has worked in the field of informal computing

education. Each approach and the associated exercises were appropriate for the author and were found to be equally engaging.

3.1 Pedagogical Approach

Complex computing projects can overwhelm the new learner. Qwiklabs facilitates in a tiered format where information is provided, comprehended, analyzed and employed before moving to the next step. As described above, quests may be at the Introductory, Fundamental, Advanced and Expert levels. The exercises in these approaches build on each other. In a quest, labs typically begin with an introduction to basic concepts, ensuring that learners are familiar with common technologies and understand the language. This is especially important for users who have had little to no experience of using CACRS. These baseline activities also provide a useful review for students who had some experience with CACRS. Each topic in the lab contained a small activity to allow for immediate application of the respective topic. To synthesize knowledge gained in the lesson, learners were tasked with activities and performed small assignments. Each segment scaling up in difficulty and building on previously-learned concepts, as per a scaffolded instruction approach. The activities are designed to engage learners at varied skill levels. For example, in labs that required software installation or deployment of pods, the labs took steps to reduce the complications arising from having to install software and associated libraries. Labs provide an installation script that automated most of the process. Furthermore, in Data Science labs, learners learn to how to call data sets and leverage existing modules to solve real world problems. At increased levels of complexity, activities that provided lesser scaffolding were harder to grasp than those that employed more scaffolding. The approach included structured components, and lecture notes (or tutorials) that can be accessed by the learner after completing the exercise.

3.2 The Learning Environment

Like many other online learning platforms, Qwiklabs quests and labs are timed and self-paced. QL labs can take anywhere from 30 minutes to 2 hours to complete, and a quest may contain between 5 to 7 labs. Rather than requiring a commitment of multiple hours to complete a module, the relatively shorter duration of these labs is a strength as learners are able to work them into their daily schedule with relative ease. For example, in order to maintain a continuous learning process, the authors strived to successfully complete one lab every day.

The labs rapidly introduce learners to the richness of the GCP environment. The menu is expansive and includes all attributes present in the GCP production environment. This includes the current analytics packages, orchestration environment and software stack. Mindful of traditional HPC users, the interface offers both command line and GUI-driven management options that may be used. While this rich environment remains a strength for experienced cloud users, we noted that this could be overwhelming for new users, particularly those who were not guided through the process. Users can select an option on the GUI and see the corresponding command on the shell. The advantage of this approach is that learners who are not familiar with the command line can get to work using the GUI instantly, while those more comfortable with using the command line and inbuilt technologies are accommodated as well. In addition to simultaneously learning how to use the GUI and command line to perform select functions, we found this two-pronged approach to be particularly useful when the instructions in the Qwiklabs tutorial did not map to the options on the GUI.

Cross-listed labs, i.e. labs associated with more than one quest, are used to create a cross-pollinating effect that encourages the learner to start an additional quest. For example, the lab, "Stackdriver: Qwik Start" is cross-listed on the "Stack Driver" and "Windows on GCP" quests. While there are significant gains associated with this approach, it can also cause learners to run into a cascading effect if a cross-listed Qwiklabs course is not well-developed or has issues. Such a situation could negate the learning environment for multiple quests. Though the author did run into less refined labs on the platform, he didn't encounter these issues with cross-listed labs. An additional issue is that Qwiklabs labs may go offline or be placed under maintenance. For example, the cross-listed course entitled "Continuous Delivery with Jenkins in Kubernetes Engine" has stopped progress on the "Cloud Architecture", and the "Kubernetes in the Google Cloud" quests.

3.3 Challenges

Technology plays an important role in the success of online learning environments. Glitches in an online training platform are likely to deter new users. The time limits, coupled with the for-credits (dollars) nature of these activities, makes the learner more sensitive to issues on these platforms. While the Qwiklabs platform is highly usable, here we describe some of the issues that we discovered in the process. These issues did not prevent us from completing a lab or exercise. Some of these issues were: (a) on starting a lab, the session would fail to launch; (b) the corresponding shell session would fail to launch for a session; (c) virtual machines took a large amount of time to spin up; (d) limits on resources supporting virtual machines (VMs) were not established in the exercise; (e) the automatic progress checks would not approve a step, while latter steps were successfully completed and authenticated by the system; and (f) labs previously completed would appear as incomplete in the learner's learning profile if restarted. A working solution for problems (a) and (b) was to refresh the web-browser, while for (c) we recommend choosing a smaller VM.

4. FUTURE DIRECTIONS

In this section we discuss ways in which Qwiklabs may be integrated into the HPC user-education, and staff professional-development environment.

4.1 Supporting a HPC Quest

Qwiklabs supports enterprise IT and web-based services. While these services overlap with the needs of the HPC community in areas such as cybersecurity and resource management, the specifics are different. In terms of research pathways, labs for Big Data analytics and artificial intelligence are perhaps best aligned to the needs of the HPC community. We once again find synergies in hosted data sets, use of shared server/notebooks (example Jupyter Notebooks, Codelab), and containers but the specifics such as the type of containers (docker on GCP, but Singularity in HPC workloads) may change. On our introduction to Qwiklabs, we tried to find an appropriate quest that focused on HPC technologies such as deploying clusters, job schedulers (SLURM and IBM Spectrum LSF), setting up SSH tunnels, how to enable cloud-burst and setting up virtual private networks. While Qwiklabs doesn't host a dedicated quest for HPC workflows, GCP has a collection of tutorials and resources that are available on Github. We direct the interested user to these resources listed in Table 2. Qwiklabs offers a number of training labs for installing and running AI-relevant software such as Tensorflow. This level of support extends to analytics such as Google BigQuery as well. To the best of our

knowledge, this resource-rich training ecosystem does not, however, extend to the use of some of the largest use-cases of HPC software stacks such as multi-physics codes (examples LAMMPS, GROMACS). A Qwiklabs quest that focuses on installing this software in a Singular container, followed by applications and best data- and computer-management practices would be useful for researchers in the HPC community. These modules could be extended to include ways to share data and results among researchers.

Table 2. Resources for Cloud Bursting [30-35].

Resource	Courses
Cloud Burst Using Slurm <i>Workflow tutorial</i>	Deploy an Auto-Scaling HPC Cluster on GCP with Slurm
	Introduction to GCP Python Client for Compute Engine
	Setting up VPN tunnel between strongSwan VPN and GCP Cloud VPN
Documentation	GCP Compute Engine: Instance Template
	Slurm Elastic Computing (Cloud Bursting)
	GCP Cloud DNS: Overview

4.2 Developing a Direct Pathway to GCP Certifications

Offering certifications to system administrators fills a key need in the cloud arena. Such certifications allow sites to help with employment and system production decisions. While Qwiklabs provides a quick way to jump onto GCP's platform, the correlation between completing quests and current GCP certifications such as the Google Cloud Associate Engineer is not clear. As discussed previously, a person may complete a lab by merely copying and pasting text from the instruction set to the lab, thus making it hard to assess the value of an individual completing a quest. A defined (or prescriptive) pathway toward achieving this certification would be meaningful. A casual look at sample questions for the Google Cloud Associate Engineer suggests that the certification exam is geared for folks who have months of experience on GCP supplementing the knowledge gained from Qwiklabs exercises. While this approach to training users is likely to work for researchers who are merely trying to find ways to migrate their workflows to GCP, it may create a "chicken or egg" situation for employers keen to adopt GCP. In the absence of lower-level certifications demonstrating an applicant's skill level, employers will be forced to make hiring decisions without knowing *a priori* whether their future employees will be able to be certified.

4.2 Surveys and Assessments

Qwiklabs relies on the honor system to report a learner's progress. The hand-holding (concierge service) nature of the learning labs is such that solutions to problems are visible to the users. Student success is generally evaluated based on their ability to complete the session's hands-on activities. In specific cases a student's progress is marked by an in-built checkpoint. While focusing on ease of use, the platform allows learners to copy paste the solution and complete a lab within minutes. Only a few labs have associated quizzes that test the knowledge gained by the exercise. This could be improved on as an extension of the work. There are opportunities to perform formative and summative assessments to gauge a learner's learning gains and the platform's overall delivery. Formative evaluation can be built into activities through short, informal student quizzes that evaluate a learners' understanding of the presented concepts. In

turn, a summative assessment would evaluate a student's learning, and collect feedback on the utility of the labs to a user's learning. It is important to note that in contrast to the Learning labs discussed in this review, Challenge labs only provide objectives rather than step-by-step instructions. These labs test a participant's ability to achieve them. Challenge labs could be associated with the quests themselves, thus functioning as capstone projects that required the students to apply the knowledge gained from the session. Taken together this data could be used to build a quantitative model for evaluating course success and to develop a profile of the kind of students (or groups of students) that are most likely to benefit from these labs.

5. CONCLUSIONS AND LESSONS LEARNED

Student training in computing is a critical area where demand currently outweighs supply. By design, the scaffolded nature of the labs in quests provide few opportunities to solve problems by developing hypotheses and validating them. In contrast, to learning labs, Challenge labs present a general objective and require participants to complete a series of tasks. In such labs, participants may choose one of many approaches to solve the problem at hand. Our experience from this training program shows that intermediate-level coding can be effectively combined with a number of interesting activities. The largest challenges lies in that these activities are not connected directly to HPC users. An abundance of free-tutorials on Github, hands-on activities and scaffolding educational technique helped reduce *a priori* knowledge that a user is required to have in order to get started on traditional HPC resources. The strategies described in this work are likely to help specific aspects of undergraduate curricula. These approaches present exciting opportunities to engage HPC users in Cloud computing, a critical step, to get them to use CACRS effectively.

6. SUPPORTING INFORMATION

All training materials used in this study are available to the community at Qwiklabs (<https://www.qwiklabs.com>). The author's progress on the Qwiklabs platform can be followed on his LinkedIn profile. Labs used in the in-person training exercise by a GCP external provider are the property of the instructor and may be accessed for reviewing purposes with explicit permission of the instructor. Surveys, and review exercises that will be developed as part of this longitudinal study may be requested from the author. Please send us feedback about your adoption experience via an email to help@hprc.tamu.edu.

7. ACKNOWLEDGEMENT

The authors thank staff at Texas A&M HPRC and Google Cloud Platform for assisting with the research related to this study. We thank the GCP training crew at Google Next 2019 for helping get started on Qwiklabs and offering introductory credits. Additional credits for Qwiklabs were kindly provided on request by GCP. We gratefully acknowledge support from the National Science Foundation (NSF). We thank the NSF for award #1649062, "*NSF Workshop: Broadening Participation in Chemical and Biological Computing at the Early Undergraduate Level*", award #1730695, "*CyberTraining: CIP: CiSE-ProS: Cyberinfrastructure Security Education for Professionals and Students*", and award # 1925764, "*NSF CC* SWEETER: South West Expertise in Expanding Training Education and Research*"

8. REFERENCES

- [1] C.A. Stewart, T.M. Cockerill, I. Foster, D. Hancock, N. Merchant, E. Skidmore, D. Stanzione, J. Taylor, S. Tuecke, G. Turner, M. Vaughn, and N.I. Gaffney, "Jetstream: a self-provisioned, scalable science and engineering cloud

- environment,” In Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure. St. Louis, Missouri. ACM: 2792774. p. 1-8. <http://dx.doi.org/10.1145/2792745.2792774>
- [2] J. Towns, T. Cockerill, M. Dahan, I. Foster, K. Gaither, A. Grimshaw, V. Hazlewood, S. Lathrop, D. Lifka, G.D. Peterson, R. Roskies, J.R. Scott, N. Wilkins-Diehr, "XSEDE: Accelerating Scientific Discovery", *Computing in Science & Engineering*, vol.16, no. 5, pp. 62-74, Sept.-Oct. 2014, doi:10.1109/MCSE.2014.80
- [3] Aristotle - <https://federatedcloud.org/about/index.php>
- [4] CloudLab - <https://cloudlab.us/>
- [5] NIH Strides - <https://datascience.nih.gov/strides>
- [6] Critical Techniques, Technologies and Methodologies for Advancing Foundations and Applications of Big Data Sciences and Engineering (BIGDATA) - https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504767
- [7] NSF Campus Cyberinfrastructure (CC*) - https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504748
- [8] Enabling Access to Cloud Computing Resources for CISE Research and Education (Cloud Access) - https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505591
- [9] Coursera- <https://www.coursera.org/>
- [10] Linux Academy - <https://linuxacademy.com/>
- [11] NSF and Internet2 to explore cloud computing to accelerate science frontiers - https://nsf.gov/news/news_summ.jsp?cntn_id=297193
- [12] Google acquires Qwiklabs to teach Developers cloud skills, November 21, 2016. This article may be accessed online at URL - <https://techcrunch.com/2016/11/21/google-acquires-qwiklabs-to-teach-developers-cloud-skills/>
- [13] Google acquires Qwiklabs to teach Developers cloud skills, November 21, 2016. This article may be accessed online at URL - <https://blog.qwiklabs.com/skill-up-the-world/>
- [14] Texas Regional Collaboratives, "Building the Texas Computer Science Pipeline Strategic Recommendations for Success | theTRC.org," 2014
- [15] K. Saichaie, D. C. Brooks, P. Long, R. Smith, R. Holeyton, C. Meyers, A. Finkelstein, S. Dugdale, "7 Things You Should Know About Research on Active Learning Classrooms," in *ELI 7 Things You Should Know*, Educause Learning Initiative (ELI), 2017
- [16] P. Baepler, J. D. Walker, D. C. Brooks, K. Saichaie, C. I. Petersen, B. A. Cohen, "A Guide to Teaching in the Active Learning Classroom: History, Research, and Practice," Stylus Publishing, ISBN-13: 978-1620363003, 2016
- [17] Student-Centered Active Learning Environment with Upside-down Pedagogies: <http://scaleup.ncsu.edu/>
- [18] A. Yadav, N. Zhou, C. Mayfield, S. Hambrusch, and J. T. Korb, "Introducing computational thinking in education courses," in Proceedings of the 42nd ACM technical symposium on Computer science education, pp. 465–470, ACM, 2011.
- [19] J. J. Lu and G. H. Fletcher, "Thinking about computational thinking," in Proceedings of the 40th ACM Technical Symposium on Computer Science Education, SIGCSE '09, (New York, NY, USA), pp. 260–264, ACM, 2009.
- [20] K. Brennan and M. Resnick, "New frameworks for studying and assessing the development of computational thinking," in Annual American Educational Research Association meeting, (Vancouver, BC, Canada), 2012
- [21] J. M. Wing, "Computational thinking and thinking about computing," *Philosophical transactions of the royal society of London A: mathematical, physical and engineering sciences*, vol. 366, no. 1881, pp. 3717–3725, 2008.
- [22] M. Prince, "Does active learning work? a review of the research," *Journal of engineering education*, vol. 93, no. 3, pp. 223–231, 2004.
- [23] J. Parsons and L. Taylor, "Improving student engagement," *Current issues in education*, vol. 14, no. 1, 2011.
- [24] Larkin, M. (2002). Using scaffolded instruction to optimize learning. <http://www.vtaide.com/png/ERIC/Scaffolding.htm>
- [25] C. Holdgraf, A. Culich, A. Rokem, F. Deniz, M. Alegro, D. Ushizima. "Portable Learning Environments for Hands-On Computational Instruction" in Proceedings of PEARC17, New Orleans, LA, USA, July 09-13, 2017, 9 pages. DOI: 10.1145/3093338.3093370
- [26] D. K. Chakravorty, M. Pennings, H. Liu, Z. Wei, D. M. Rodriguez, Levi T. Jordan, D. F. McMullen, N. Ghaffari, and S. D. Le. "Effectively Extending Computational Training Using Informal Means at Larger Institutions," *Journal of Computational Science Education* 2018, 40-47 DOI 10.22369/issn.2153-4136/10/1/7.
- [27] D. K. Chakravorty, M. Pennings, H. Liu, Z. Wei, D. M. Rodriguez, L. T. Jordan, D.F. McMullen, N. Ghaffari, S. D. Le, D. Rodriguez, C. Buchanan, and N. Gober. "Evaluating Active Learning Approaches for Teaching Intermediate Programming at an Early Undergraduate Level," *Journal of Computational Science Education* 2018, 61-66 DOI 10.22369/issn.2153-4136/10/1/10.
- [28] D. K. Chakravorty, D. F. McMullen, N. Gober, J. H. Seo, M. Bruner, and A. Payne. "Using Virtual Reality to Enforce Principles of Cybersecurity," *Journal of Computational Science Education* 2018, 81-87 DOI 10.22369/issn.2153-4136/10/1/13.
- [29] D. K. Chakravorty, M. Pennings, H. Liu, X. Thomas, D. M. Rodriguez, and L.M. Perez. "Incorporating Complexity in Computing Camps for High School Students – A Report on the Summer Computing Academy Program at Texas A&M University," *Journal of Computational Science Education* 2019. (Accepted).
- [30] Slurm on Google Cloud Platform - <https://github.com/SchedMD/slurm-gcp>
- [31] Using the Cloud Client Libraries for Python - <https://cloud.google.com/compute/docs/tutorials/python-guide>
- [32] How to set up a VPN between strong Swan and Cloud VPN - <https://cloud.google.com/community/tutorials/using-cloud-vpn-with-strongswan>
- [33] GCP Compute Engine: Instance Template - <https://cloud.google.com/compute/docs/instance-templates/>

[34] Slurm Elastic Computing (Cloud Bursting) -
https://slurm.schedmd.com/elastic_computing.html

[35] GCP Cloud DNS: Overview The URL may be found at -
<https://cloud.google.com/dns/docs/overview>

REPRODUCIBILITY APPENDIX

Computational results are not part of this paper. (Please refer to <https://sc18.supercomputing.org/submit/sc-reproducibility-initiative/>)