Strengthening the US Department of Energy's Recruitment Pipeline: The DOE/NNSA Predictive Science Academic Alliance Program (PSAAP) Experience

John K. Holmen
holmenjk@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

Verónica G. Vergara Larrea
vergaravg@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

Erik W. Draeger
draeger1@llnl.gov
Lawrence Livermore National Laboratory
Livermore, California, USA

Eric T. Phipps
etphipp@sandia.gov
Sandia National Laboratories
Livermore, California, USA

Philip J. Smith
philip.smith@utah.edu
University of Utah
Salt Lake City, Utah, USA

Martin Berzins
mb@sci.utah.edu
University of Utah
Salt Lake City, Utah, USA

Sean T. Smith
smith_sean@lanl.gov
Los Alamos National Laboratory
Los Alamos, New Mexico, USA

Jeremy N. Thornock
thornock2@llnl.gov
Lawrence Livermore National Laboratory
Livermore, California, USA

Suzanne Parete-Koon
paretekoonst@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

ABSTRACT

The US Department of Energy (DOE) oversees a system of 17 national laboratories responsible for developing unique scientific capabilities beyond the scope of academic and industrial institutions. These labs strive to keep America at the forefront of discovery and are home to some of the Nation's best minds and the world's best scientific and research facilities. Collaborations between national laboratories and academic institutions are critical to develop and recruit talent for the DOE workforce. Academia's cooperative education model poses challenges for DOE recruitment pipelines centered around traditional internships. This paper discusses a promising DOE recruitment pipeline, the National Nuclear Security Administration's (NNSA) Predictive Science Academic Alliance Program (PSAAP) initiative. As a part of this, experiences capturing the successes and challenges faced by the University of Utah's Carbon Capture Multidisciplinary Simulation Center (CCMSC) through their participation in the PSAAP-II initiative are shared. These experiences demonstrate the success of Utah's PSAAP center as a recruitment pipeline with approximately 43% of CCMSC students going to a national laboratory after graduation. Potential opportunities to strengthen the DOE's recruitment pipeline are also discussed.

CCS CONCEPTS

• **Social and professional topics** → **Computing education programs**; **Employment issues**; **Funding**.

PEARC '23, July 23–27, 2023, Portland, OR, USA

ACM ISBN 978-1-4503-9985-2/23/07.
https://doi.org/10.1145/3569951.3593606

KEYWORDS

Cooperative Education, Experiential Education, Recruitment Pipeline, Workforce Development

ACM Reference Format:
John K. Holmen, Verónica G. Vergara Larrea, Erik W. Draeger, Eric T. Phipps, Philip J. Smith, Martin Berzins, Sean T. Smith, Jeremy N. Thornock, and Suzanne Parete-Koon. 2023. Strengthening the US Department of Energy's Recruitment Pipeline: The DOE/NNSA Predictive Science Academic Alliance Program (PSAAP) Experience. In *Practice and Experience in Advanced Research Computing (PEARC '23), July 23–27, 2023, Portland, OR, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3569951.3593606

1 INTRODUCTION

The US Department of Energy (DOE) oversees a system of 17 national laboratories responsible for developing unique scientific capabilities beyond the scope of academic and industrial institutions. These laboratories strive to keep America at the forefront of discovery and are home to some of the Nation's best minds and the world's best scientific and research facilities. Figure 1 shows where these laboratories are located throughout the United States. Locations span from East Coast to West Coast with laboratories ranging from multi-purpose security to single-purpose physics facilities.

In recent years, the DOE has made several efforts to broaden workforce development initiatives. Examples include the Exascale Computing Project's (ECP) Broadening Participation Initiative [7], which includes the Sustainable Research Pathways for High Performance Computing (SRP-HPC) initiative [8], and the Reaching a New Energy Sciences Workforce (RENEW) initiative [6]. Additionally, the DOE has recently announced $56 million in funding to provide research opportunities to historically underrepresented groups and institutions in STEM [14].

These efforts are encouraging as they demonstrate the DOE's interest in and commitment to establishing strong recruitment

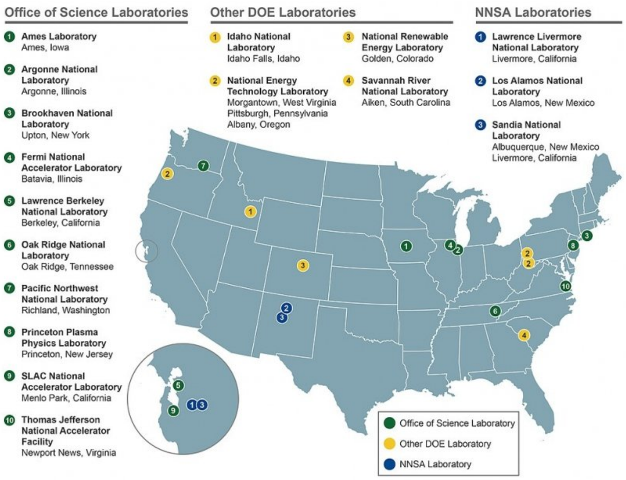

Figure 1: Map of the 17 national laboratories. [Public domain], via DOE. (https://www.energy.gov/articles/map-explore-national-labs).

pipelines that are both supportive and inclusive of underrepresented groups and institutions. SRP-HPC defines underrepresented groups as Black or African American, Hispanic/Latin, American Indian, Alaska Native, Native Hawaiian, Pacific Islanders, women, persons with disabilities, first generation scholars, and students from community and smaller colleges and universities. In addition to students from community and smaller colleges and universities, students from co-op programs are also underrepresented.

Co-op programs differ from internship experiences in that they require a one-to-many year experience rather than a one-off experience. This can be problematic when funding is required as it often involves a longer running project than a traditional internship experience. An existing funding model with potential to align well with co-op programs due to the associated multi-year project is the NNSA's Predictive Science Academic Alliance Program (PSAAP) initiative, which aims to train the next generation of laboratory staff.

This paper describes experiences capturing the successes and challenges faced by the University of Utah's Carbon Capture Multidisciplinary Simulation Center (CCMSC) through their participation in the PSAAP-II initiative. These experiences demonstrate the success of Utah's PSAAP center as a recruitment pipeline with approximately 43% of CCMSC students going to a national laboratory after graduation. For comparison, Westat reports that 32% of DOE Computational Science Graduate Fellowship recipients began their first position at a DOE laboratory after completing their fellowship [27]. The Westat data is a publicly available benchmark worth referring to as it provides a reasonable measure of calibration that is otherwise missing. Potential opportunities to strengthen the DOE's recruitment pipeline are also discussed.

Note, the PSAAP experience captured here is limited to that of a single center and may not be representative of the typical PSAAP experience. To help avoid bias and imbalance, care was taken in selecting perspectives. Perspectives include external reviewers, mentors, and employers, internal faculty, and recipients. To fully evaluate the successes and challenges of the PSAAP initiative, a comparison of experiences across multiple PSAAP centers would be appropriate. Such an evaluation is beyond the scope of this paper. Further, this paper does not present detailed comparisons of recruitment and retention rates across recruitment pipelines.

The remainder of this paper is structured as follows. Section 2 describes DOE-related efforts to broaden workforce development initiatives. Section 3 provides an overview of academia's cooperative education model. Section 4 provides an overview of co-op hiring challenges. Section 5 describes the DOE/NNSA Predictive Science Academic Alliance Program (PSAAP). Section 6 describes the University of Utah's PSAAP-II experience. Section 7 describes potential opportunities to strengthen DOE recruitment pipelines. Section 8 concludes this paper.

2 WORKFORCE DEVELOPMENT INITIATIVES

In recent years, the DOE has made several efforts to broaden workforce development initiatives. Examples include the Exascale Computing Project's (ECP) Broadening Participation Initiative [7], which includes the Sustainable Research Pathways for High Performance Computing (SRP-HPC) initiative [8], and the Reaching a New Energy Sciences Workforce (RENEW) initiative [6]. The ECP Broadening Participation Initiative is establishing a sustainable plan to recruit and retain a diverse HPC workforce by fostering a supportive and inclusive culture within the computing sciences at DOE national laboratories. The initiative has three complementary thrust areas: (1) HPC Workforce Development and Retention (HPC-WDR), which is an action group working to improve the DOE's HPC workforce culture in creative ways, (2) Intro to HPC, which is an action group working to develop training materials for educating HPC newcomers, and (3) the Sustainable Research Pathways for High Performance Computing (SRP-HPC) initiative, which is an internship and mentoring program. SRP-HPC is based on a program started in 2015 at the Lawrence Berkeley National Laboratory and developed by the Sustainable Horizons Institute. The original SRP concept was scaled up across the ECP community and the DOE by the ECP Broadening Participation Initiative in 2021. SRP-HPC aims to connect students and faculty working among underrepresented groups with DOE national laboratory scientists to encourage lasting collaborations, jump start careers, and build inclusive workplace environments. The RENEW initiative aims to leverage the Office of Science's unique national laboratories, user facilities, and other research infrastructures to provide training opportunities for undergraduate and graduate students, postdoctoral researchers, and faculty at academic institutions not currently well represented in the U.S. science and technology ecosystem.

3 THE COOPERATIVE EDUCATION MODEL

The cooperative education model (co-op) was developed at the University of Cincinnati by Dr. Herman Schneider in 1906 [4]. Co-op is defined as an educational methodology in which periods of classroom instruction alternate with periods of paid discipline-related work experience [2]. Co-op students alternate between classroom and work experiences throughout a portion or all of their academic career. Early adopters of the cooperative education model include the University of Pittsburgh, the University of Detroit,

Table 1: Kettering University's Academic Calendar

Term	A Section	B Section
Summer (July - September)	School	Co-op
Fall (October - December)	Co-op	School
Winter (January - March)	School	Co-op
Spring (April - June)	Co-op	School

Georgia Institute of Technology, Rochester Institute of Technology, Drexel University, and Kettering University.

A detailed example of the co-op experience can be found in a recent article [21]. This article captures Kettering University's co-op experience and discusses the academic program, work experience, and related thesis project. At Kettering University, undergraduate students are able to earn up to 2.5 years of work experience before graduation with a minimum of 1.25 years of work experience required to graduate. This is accomplished using a term-based academic calendar that is divided into two sections, A and B. Table 1 describes when students participate in school or co-op depending upon their designated section. These alternating sections allow for a steady stream of students to be both on campus and at work throughout the year. This is advantageous for co-op employers, where Section A students can replace Section B counterparts, allowing the co-op employer to have the equivalent of a year-round full-time employee.

The cooperative education model offers value not only for the student but also for the employers and institutions. For students, the primary value in co-op lies in the ability to experience one's field of interest firsthand prior to committing to a career. For employers, the primary value in co-op lies in the ability to "grow your own" employee through the extensive work experience a student will accumulate throughout their academic career. For institutions, the primary value in co-op lies in the ability to cultivate lasting relationships and collaborations between co-op employers. More discussion on benefits and challenges related to the cooperative education model can be found in recent articles [9, 15, 22, 28].

4 CO-OP HIRING CHALLENGES

Co-ops differ from internships in that they are typically a one-to-many year experience rather than a one-off experience. While beneficial for improving a student's workplace-readiness, the co-op experience can be problematic when funding is required as it often involves a longer running project than a traditional internship experience. This challenge has been experienced firsthand at the Oak Ridge National Laboratory (ORNL), where student employment opportunities are primarily focused on one-off experiences. As a result, the ORNL authors have had limited avenues to onboard students from co-op programs until recently with the creation of a co-op partnership with Tennessee Tech University (TTU).

In order to establish the co-op program with TTU, ORNL had to explore different avenues to identify the best alternative to onboard students. One challenge was that TTU has four co-op plans for students: (1) Plan A (Traditional), where a student works full-time for an employer for up to 12 months, (2) Plan B (Alternating), where a student works alternate semesters at the employer's site (work, return to school, work, etc.), (3) Plan C (Parallel), where a student will attend college and work locally approximately 20 hours per week for the employer, and (4) Plan D (Summer), where a student works during the summer semester only. The diversity of plans available to students and the unique and varied timelines that students operate on pose a challenge for employers such as ORNL that are looking to establish co-op opportunities.

To further complicate matters, other institutions, such as Kettering University, offer an entirely different timeline which would require establishing a separate employment type and funding structure to accommodate students.

5 THE PREDICTIVE SCIENCE ACADEMIC ALLIANCE PROGRAM (PSAAP)

An existing funding model with potential to align well with the cooperative education model due to the associated multi-year project is the NNSA's PSAAP initiative, which aims to train the next generation of laboratory staff. PSAAP is the primary mechanism by which the NNSA's Advanced Simulation and Computing (ASC) program engages the U.S. academic community in advancing science-based modeling and simulation. In 2020, nine PSAAP-III centers were established, focusing on three major integrated areas: (1) Discipline-focused research needed to further predictive science and enabled by effective Exascale computing technologies; (2) Developing and demonstrating technologies and methodologies to support effective Exascale computing in the context of science/engineering applications; and (3) Predictive Science based on verification and validation and uncertainty quantification (V&V/UQ) for large-scale simulations [20].

PSAAP funds three types of Centers: (1) Multi-disciplinary Simulation Centers (MSCs), (2) Single-Discipline Centers (SDCs), and Focused Investigatory Centers (FICs). MSCs and SDCs focus on scalable application simulations, targeting either large-scale, integrated multidisciplinary problems or a broad single science/engineering discipline, respectively. MSCs and SDCs develop and demonstrate computer science technologies and methodologies that will advance Exascale computing, and demonstrate integrated, verified, validated predictive simulation with uncertainty quantification. FICs are tightly focused on a specific research topic in either a science/engineering discipline, or an Exascale enabling technology, of interest to NNSA's mission [20].

NNSA-funded graduate students at each center are required to complete a 10 consecutive week visit to one of the three NNSA National Laboratories during their graduate career. During their visits, graduate students shall conduct research related to their responsibilities at their home institution. In addition, they shall take advantage of opportunities provided by the Laboratories that expose them to Laboratory research programs. These visits may occur during summers, but often do not. [20].

Though only a single visit is required, multiple visits are possible due to projects lasting many years. This lends itself well to the cooperative education model as it allows students to take on increasingly more responsibility in the multi-year project during subsequent lab visits. This is important for co-op programs with a project and/or thesis component. Though possible, such continuity is difficult to achieve through traditional internships.

For students, the key difference between a PSAAP experience and an internship experience is that the latter typically ends after the lab visit. In the case of PSAAP, the experience begins at the home institution even before the lab visit and continues throughout the student's involvement with the PSAAP center, which potentially spans their graduate career. For example, students typically return to their home institution and continue their collaborations, make use of national laboratory resources (e.g., high performance computing (HPC) systems), participate in HPC training (e.g., leadership computing facility training for users), and participate in bi-annual center reviews on campus. This lends itself well to the cooperative education model as it allows students to pursue a continuous line of research throughout their academic career regardless of whether they are at school or on co-op. This is important for co-op programs with a project and/or thesis component.

6 THE PSAAP-II EXPERIENCE

Sections 6.1 through 6.8 discuss takeaways related to the University of Utah's Carbon-Capture Multidisciplinary Simulation Center (CCMSC) and their participation in the DOE/NNSA PSAAP-II initiative. Care was taken in selecting perspectives to help avoid bias and imbalance towards PI and co-PI views. Perspectives include external reviewers, mentors, and employers, internal faculty, and recipients. The CCMSC used the Uintah Computational Framework [1] to predict performance of a 1000 MWe ultra-supercritical clean coal boiler. More details on the CCMSC can be found on their website [3] and in the final report [26].

6.1 Tri-Lab Support Team Lead (External)

This section discusses external takeaways from the perspective of Erik W. Draeger from Lawrence Livermore National Laboratory. Erik participated in the PSAAP-II initiative as the Tri-Lab Support Team (TST) Lead for Utah's CCMSC. As the TST Lead, he served as the technical interface between the CCMSC and the tri-labs and worked to insure the success of the CCMSC.

From Erik's perspective, the CCMSC was a textbook example of the value of initiatives like PSAAP-II from the perspective of workforce development. The Center had ambitious, well-defined multiphysics simulation goals that effectively leveraged the expertise and abilities of faculty, staff and students across institutions and disciplines. While there were many underlying research challenges that had to be addressed, the structure of the predictive simulation plan meant that participants could easily see where and how their contributions had impact on the larger-scale results throughout the project. This level of coordinated research effort requires both multi-year funding and excellent management. One of the most successful strategies to enable cohesion across the Center was the use of hierarchical error quantification and top-down V&V/UQ to regularly and quantitatively identify the dominant sources of error in the multiphysics simulations to motivate key research drivers. Similarly, from the Computer Science side of the project, the roadmap of upcoming computer architectures and the need for task-based parallelism and hardware abstraction was clearly articulated and regularly communicated at all levels of the project. This allowed all participants, but particularly students, to have a consistent view into the need for and immediate impact of their individual contributions. This approach to team-based problem-solving closely matches how NNSA programs attack multiphysics challenge problems, demonstrating how PSAAP-II Centers can naturally complement and augment the work at the Labs.

The CCMSC PSAAP-II project was also an unqualified success as a demonstration of the value of high performance computing (HPC) as an enabling resource for industry. The Center's demonstration calculations of different coal-fired boiler designs showcased the efficiencies that can be realized in both construction and operational costs. Simulation-driven design optimization is a major growth area for industry and one that is needed to ensure economic competitiveness, but without a workforce trained in how to effectively use HPC it will be difficult or impossible to realize. Centers like the CCMSC are therefore essential in providing working models for how and when to use simulation most effectively, on top of the foundational research in physics and CS needed to make such efforts possible.

The biggest challenge of the PSAAP model is being able to sustain and replicate the successful projects like the CCMSC at the conclusion of the project. Best practices and lessons learned are challenging to impart across institutions and research domains under ideal circumstances, but doubly so in cases like the CCMSC where talented people came together to form a team whose effectiveness far exceeded the sum of its parts.

6.2 Tri-Lab Support Team Mentor (External)

This section discusses external takeaways from the perspective of Eric T. Phipps from the Sandia National Laboratory. Eric participated in the PSAAP-II initiative as a Tri-Lab Support Team (TST) Mentor for Utah's CCMSC. As a TST Mentor, he coordinated internships for J.K. Holmen and other students at the Sandia National Laboratory.

Each Multi-disciplinary Simulation Center funded by the PSAAP program includes a Tri-Lab Support Team (TST) consisting of two NNSA lab staff members from each of the three NNSA labs. The purpose of the TST is to (1) provide technical advice to the MSC on research directions and scope to ensure MSC success and relevance to NNSA missions, (2) foster research collaboration opportunities between the MSC and lab personnel, and (3) facilitate internships for MSC students at each lab to help fulfill the recruitment goals of the PSAAP program.

How successful the TST is at achieving these goals depends greatly on the overlap between TST members' technical knowledge/experience and the research directions being pursued by the MSC. Each MSC is undertaking leading-edge research in the areas of high-performance computing, physical simulation, uncertainty quantification, verification, and validation, and thus TST members must be well-versed in these topics to provide effective advice and have the research connections within their respective labs to facilitate collaborations. This is particularly important in arranging internships that are fruitful for the student, MSC, and labs as each lab is a very large institution consisting of thousands of researchers. It is impossible for any TST member to be aware of all of the research being conducted at each lab and who is doing it, and so it is only practical to arrange good matches between lab personnel and students in research areas the TST members are familiar with.

Over the course of the PSAAP-II program, it became clear this model is most successful when TST members are themselves involved in research directly related to MSC technical objectives as they are much more likely to have deep technical knowledge of the field as well as be aware of relevant networks of researchers at their home lab. This was found to be true with the CCMSC as many of the TST members themselves conducted research in areas highly relevant to the CCMSC objectives, leading to numerous joint academic publications, cross fertilization of research ideas between the CCMSC and the labs, and eventual recruitment of many CCMSC students for positions at NNSA and DOE labs.

6.3 Academic Principal Investigator (Internal)

This section discusses internal takeaways from the perspective of Philip J. Smith from the University of Utah. Philip participated in the PSAAP-II initiative as the Principal Investigator (PI) and Director for Utah's CCMSC. As the PI and Director, he led the predictive science, V&V/UQ, education, and outreach efforts.

The CCMSC was created for the purpose of developing and demonstrating the use of formal uncertainty quantification (UQ) methodologies in conjunction with scalable and portable high performance computing (HPC) strategies for solving large practical problems. The parallels between CCMSC objectives and the NNSA lab objectives were strong. NNSA oversight and collaborations changed the way the CCMSC approached their problem, organized their teams and conducted their research. The application selected by the Center was the demonstration of positive societal impact of HPC with UQ for the deployment of low-cost, low-carbon energy solutions for power generation. To accomplish this mission, the CCMSC developed a multi-physics, large-eddy simulation (LES) code (Arches/Uintah) to run at scale on world-class computational resources made available to them by NNSA. To guide their application, they partnered with two industrial collaborators, General Electric (GE Power) and Ontario Power Generation (OPG). These industrial partners and their applications provided purpose and focus to the methodologies developed in the Center. With GE Power, the CCMSC objective was to demonstrate the advantages of HPC with hierarchical UQ for design decisions. Specifically, the objective was to predict the heat flux profile to a validated level of uncertainty for a full-scale, pulverized coal, thermal power generation boiler. The CCMSC capstone project was accomplished in partnership with OPG, where they deployed all the methodologies of the Center to demonstrate dynamic, online artificial intelligence (AI) for operating a biomass-fired power generation boiler. Seventy-eight engineers and scientist worked together in three teams to complete the Center mission: the computer science team, the computational physics team, and the UQ team.

The constant focus on an application-driven objective brought clarity to decision making in both resource allocation and research task decisions. They learned lessons on the importance of achieving predictivity through hierarchically-driven, science-based model development in tight conjunction with formal validation and uncertainty quantification. They learned the value of a well defined quantity of interest to keep tangential inquiry from sabotaging a mission-driven objective. They learned the value of quantifying what is good enough (engineering vs. science). They learned the value of high performance computing that allows for scientific exploration of real operational space. They learned how crucial it is to have integration between multiple disciplines (ie. computer scientists, physical scientists and data scientists). They learned how important it is to have people who are comfortable working at the interface of these many disparate disciplines. They learned that most academic educational environments build silos instead of the multi-disciplinarity needed for large societal problems.

The PSAAP program changed who Philip is today. It changed how he conducts his research and directs his research team. It changed what and how he teaches in the classroom. It changed his interactions in the non-technical world too. He has had many conversations with former center personnel, who each expresses a similar appreciation for what they learned through interactions in the PSAAP program.

6.4 Exascale Computing Lead (Internal)

This section discusses internal takeaways from the perspective of Martin Berzins from the University of Utah. Martin participated in the PSAAP-II initiative as the Exascale Computing Lead for Utah's CCMSC. As the Exascale Computing Lead, he led the software effort, which started with a proven computational platform, the Uintah Computational Framework, and sequentially moved towards multi-petaflop and eventually exascale computing. Additionally, he advised J.K. Holmen through his doctoral dissertation.

PSAAP-II provided both unique opportunities and some challenges for Utah's Computer Science research. For this project, there was a clear division of labor among Uintah application development as tasks were written by the Physics team and the core task-management infrastructure and solver components were written by the Computer Science team to run at scale. Uintah's task-based approach made this clean separation possible. The need to perform production runs while at the same time moving the core infrastructure forward was also possible and indeed was achieved by the core developers and students. In addition to runtime infrastructure, a ray tracing-based radiation model was also developed in parallel by Alan Humphrey [17]. This work also led to a number of Ph.D. dissertations [10, 16, 24, 25].

A key success of this division of labor was the agile pattern of rapid deployment of recent research ideas, which was closer to a production environment in a DOE laboratory than is usual in a University setting. A key challenge of this division of labor is that it required a careful delineation of who did what at dissertation proposal defense time. Another challenge was that this division of labor led to the development of distinctive software styles that differed greatly between the Computer Science and Physics teams. In hindsight, there should have been more software engineering coordination by the Computer Science lead and team across the project. Examples of where this had an impact were the creation of too many lightweight tasks for which host-device transfer overhead dominated [12] and the creation of an application-friendly C++ abstract naming approach that made it difficult to follow variables while debugging [11]. However, such relatively minor differences did not have a serious impact on the overall success of the PSAAP-II center.

Another key success of Utah's PSAAP-II center relates to the final aim to move to Exascale as an early user of the Aurora/A21 system through the Aurora Early Science Program. Though delayed by the changes and evolving nature of the architecture, a Uintah benchmark developed by J.K. Holmen and D. Sahasrabudhe and run across the DOE Summit and NSF Frontera systems[13] is now almost ready to run portably on multiple GPU architectures at scale as of this writing. This is in no small part due to Utah's participation in PSAAP-II, Uintah's task-based approach, DOE developments such as the Kokkos performance portability layer, and many other software developments for Aurora and forthcoming exascale systems.

6.5 Graduate Research Assistant (Recipient)

This section discusses recipient takeaways from the perspective of John K. Holmen from the University of Utah (now at Oak Ridge National Laboratory). John participated in the PSAAP-II initiative as a Graduate Research Assistant with Utah's CCMSC. As a Graduate Research Assistant, he pursued research aligned with the CCMSC's exascale computing and software goals that contributed to his doctoral dissertation.

When applying for Ph.D. programs, Utah's PSAAP-II center was advertised as one of the options for dissertation research at Utah's Scientific Computing and Imaging Institute. Seeing great value in PSAAP after his experiential learning experience at Kettering University, John chose to attend the University of Utah solely for the opportunity to participate in Utah's PSAAP-II center. He is forever grateful for having made this decision as PSAAP enabled invaluable opportunities for him to learn, grow, and network.

One of the key successes of his participation was the ability to participate in several years of comprehensive training on how to effectively use HPC. This was made possible through extensive use of HPC systems and hands-on experience with workflows similar to those of a national laboratory. This preparation for a career in HPC was foundational in easing his transition to the Oak Ridge National Laboratory. Another success was the ability to build a strong network of connections to the national laboratories. This was made possible through bi-annual reviews where students had the opportunity to interact with laboratory staff and present their research to the multidisciplinary simulation center. Another success was having the opportunity to experience a national laboratory firsthand through the mandatory internship. This allowed students to experience what it is like to work at a national laboratory without having to wait for graduation. Despite participating in a single internship, the PSAAP-II experience as a whole felt comparable to his Kettering University co-op experience due to the bi-annual reviews and multi-year project.

One of the key challenges of his participation was the disconnect between his internship experience and dissertation research. The project chosen for the 10-week visit was loosely related to his dissertation research and did not contribute to the end dissertation. Care must be taken when selecting a project for the 10-week visit. Nevertheless, the internship experience was invaluable for the opportunity to experience what it was like to work at a national laboratory firsthand.

6.6 Post-PSAAP Employer (External)

This section discusses external takeaways from the perspective of Verónica G. Melesse Vergara from the Oak Ridge National Laboratory. Verónica did not participate in the PSAAP-II initiative. She hired J.K. Holmen as an HPC Engineer in the System Acceptance and User Environment group after successfully defending his doctoral dissertation.

The System Acceptance & User Environment group is comprised of HPC engineers with broad interests and diverse backgrounds. Experienced HPC engineers are often well-established at their current institutions and less likely to make an organizational switch. On the other hand, hiring recent graduates often is done with the expectation that they will need a longer period of time to gain the required HPC knowledge to fully contribute to the projects in the group. Based on Verónica's recent experience with John, it was clear that, even though he was joining the team right out of school, he had the needed experience in HPC and scientific computing to start working on projects in short order. The experience and connections he gained via PSAAP-II are clearly an advantage as he now has been able to establish multiple cross-institution collaborations within a few months of starting in the team.

6.7 Post-PSAAP Recruit (Internal)

This section discusses internal takeaways from the perspective of Sean T. Smith from the University of Utah (now at Los Alamos National Laboratory). Sean participated in the PSAAP-II initiative as a Research Professor for Utah's CCMSC. As a Research Professor, he contributed to the predictive science and V&V/UQ efforts.

From the V&V/UQ perspective, he has had many people at Los Alamos National Laboratory curious to hear the story of CCMSC. What they've found compelling is that the center's problem was surprisingly analogous to the lab's mission yet significantly smaller. A key success of this problem size was that it allowed the center, in five years, to work through every aspect of the problem and consider it in a wholistic way. In contrast, the lab is progressing through a many-decades-long program in which the significantly larger number of individuals are only exposed to small pieces. A major challenge he has faced at the lab has been determining how each piece should be executed in a way that best allows the pieces to fit together in the end while those performing the execution have an extremely limited window of understanding.

6.8 Post-PSAAP Recruit (Recipient)

This section discusses recipient takeaways from the perspective of Jeremy N. Thornock from the University of Utah (now at Lawrence Livermore National Laboratory). Jeremy participated in the PSAAP-II initiative as a Research Professor for Utah's CCMSC. As a Research Professor, he contributed to large-eddy simulation capability development, working with the computer science team to promote portability and scalability while facilitating inclusion of necessary physics into code and algorithms for the CCMSC target problem.

One takeaway to the success of the PSAAP-II CCMSC, in Jeremy's observation, was a collection of people who enjoyed excelling in the overlapping technical regions of the center's goal. These interface/overlap people included professional staff and faculty as well as students who were willing to leave their respective technical comfort zone. An important ingredient to this success was coordinating communication and efforts across a broad set of distinct skillsets. The coordination relied on finding the right balance of overlap, allowing each technical group enough freedom to influence the other without becoming heavy-handed in one approach, but together with collective investment in the center's overarching goal.

The PSAAP-II initiative provided a DOE/NNSA microcosm of a compelling technical challenge along with a diverse work environment that broadened Jeremy's technical exposure. This experience facilitated a positive career transition into a DOE/NNSA laboratory. The attention to the V&V/UQ concepts developed within CCMSC has had a particular impact on his current position. This exposure has led him into another interface role, working with a diverse set of people to tackle the Lab's various UQ missions. His experience with the hierarchical organization of a UQ objective has provided a compelling method for communicating a clear UQ workflow accompanied with examples of the CCMSC's V&V/UQ successes.

7 OPPORTUNITIES TO STRENGTHEN THE DOE RECRUITMENT PIPELINE

Sections 7.1 through 7.6 discuss opportunities to strengthen the DOE recruitment pipeline. These opportunities are informed by the authors' collective experience with various DOE recruitment pipelines.

7.1 Adopt PSAAP-Like Models at Other Labs

PSAAP is managed by the NNSA Office of Advanced Simulation and Computing (ASC) in collaboration with Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratories. These are 3 of the 17 national laboratories in the DOE's research system. Considering the extension of PSAAP or PSAAP-like initiatives to other national laboratories is valuable for strengthening DOE recruitment pipelines elsewhere. An example could be creating an academic alliance initiative based on PSAAP for the Office of Science laboratories.

7.2 Increase Visibility to Students

PSAAP provides students with a one-of-a-kind opportunity to experience lab-aligned research firsthand through a multi-year project. This is made possible through close collaboration between academia and the tri-labs as well as the mandatory lab visit. Considering explicit advertisement of the student experience is valuable for increasing awareness of the program and student interest in the national labs. An example could be advertising PSAAP-funded student positions at career fairs either locally at PSAAP-sponsored schools or nationally through events such as the SuperComputing or Tau Beta Pi career fair.

7.3 Support Co-Op Students

NNSA-funded graduate students are required to complete a 10-week visit to one of the NNSA national laboratories during their graduate career. This visit is more so an internship experience rather than a co-op experience. Considering the incorporation of a co-op track would be valuable for improving alignment with academia's cooperative education model. An example could be increasing the visit requirement from one 10-week visit to between three and five 10-week visits.

7.4 Support Early Career Students

PSAAP funds graduate students. Considering extending the PSAAP experience to undergraduate and high school students could help strengthen the recruitment pipeline by exposing students to lab life earlier on in their academic careers to help better inform their next steps. An example for supporting such students could be through a mentorship experience such as the Oak Ridge Leadership Computing Facility offers through their Next Generation Pathways to Computing Program [23], where high school students can experience lab life through collaboration with an undergraduate mentor.

7.5 Support a Diverse Set of Students

As was laid out in a DOE Advanced Scientific Advisory Committee (ASCAC): Workforce Subcommittee Letter [5], "All large DOE national laboratories face workforce recruitment and retention challenges in the fields within Computing Sciences that are relevant to their mission. Future projections indicate an increasing workforce gap and a continued underrepresentation of minorities and females in the workforce unless there is an intervention." Studies like the AIP Team-UP report [18] and Leung 2018 [19] show that there are specific factors that can help support the success of culturally diverse students in STEM fields. These factors include ensuring that students develop a sense of belonging within their cohort, and an identity within their academic discipline, and that students have academic, personal, and financial support. Additionally, it is important that institutional leadership prioritize creating environments, policies, and structures that support students from under-represented populations. To further foster a sense of belonging, it is important for students to have opportunities to interact with researchers both inside and outside of their national lab work. To develop a sense of identity within computing, mentors should encourage students to participate in computing conferences, internships, and professional societies. Research groups should take time to consider ground rules for mindfully treating their members and students with respect. Leaders should consider taking culturally aware mentor training. Measures like these make the workplace more supportive of all staff and students.

7.6 Lab Staff Sabbaticals

PSAAP centers pursue a multi-year project in collaboration with national laboratory staff. Considering adding a sabbatical component for national laboratory staff to visit the PSAAP center could help provide a more in-depth perspective of the techniques applied and a better understanding of how they fit together. Further, it would provide significantly more insight into potential employees while also providing an opportunity to recruit locally at the PSAAP-sponsored school. An example for such an experience could be incorporating a 10-week visit during the final years of the center to help with center ramp down. Such timing would increase the effectiveness of the visit through exposure to lessons learned over the life of the center.

8 CONCLUSIONS

Strengthening DOE recruitment pipelines is important for helping to develop and recruit talent for the national laboratories. A key challenge in strengthening such pipelines is identifying underrepresented groups and institutions among current workforce development initiatives. One such underrepresented group is students from co-ops, who have more demanding employment requirements than traditional internship experiences.

An existing funding model with potential to align well with co-op programs due to the associated multi-year project is the NNSA's Predictive Science Academic Alliance Program (PSAAP) initiative, which aims to train the next generation of lab staff. This paper described experiences capturing the successes and challenges faced by the University of Utah's Carbon Capture Multidisciplinary Simulation Center (CCMSC) through their participation in the PSAAP-II initiative. These experiences demonstrated the success of Utah's PSAAP center as a recruitment pipeline with approximately 43% of CCMSC students going to a national laboratory after graduation. Potential opportunities to strengthen the DOE's recruitment pipeline were also discussed.

Next steps at Oak Ridge National Laboratory include identifying how to onboard both Tennessee Tech University and Kettering University co-op students. Once a process is in place, students will be recruited for co-op positions. Long-term, the goal is to establish a pipeline of co-op students and faculty to collaborate with. For Kettering University co-op students, this goal also extends to collaborating with students on their bachelors theses.

ACKNOWLEDGMENTS

This material is based upon work originally supported by the Department of Energy, National Nuclear Security Administration, under Award Number(s) DE-NA0002375. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. We would like to thank the NNSA Office of Advanced Simulation and Computing (ASC) for managing the PSAAP program and providing an invaluable on-ramp to the national laboratories. We would also like to thank Ashley Barker for the constructive feedback on this submission.

REFERENCES

[1] M. Berzins, J. Luitjens, Q. Meng, T. Harman, C. A. Wight, and J. R. Peterson. 2010. Uintah: A Scalable Framework for Hazard Analysis. In *Proceedings of the 2010 TeraGrid Conference* (Pittsburgh, Pennsylvania) *(TG '10)*. Association for Computing Machinery, New York, NY, USA, Article 3, 8 pages. https://doi.org/10.1145/1838574.1838577

[2] C. Cates and K. Cedercreutz. 2008. *Leveraging cooperative education to guide curricular innovation: The development of a corporate feedback system for continuous improvement.* Center for Cooperative Education Research and Innovation, University of Cincinnati.

[3] CCMSC. 2022. *The Carbon-Capture Multidisciplinary Simulation Center.* Retrieved December 29, 2022 from https://ccmsc.utah.edu/

[4] K. Cedercreutz and C. Cates. 2010. Cooperative education at the University of Cincinnati: A strategic asset in evolution. *Peer Review* 12, 4 (2010), 20–24.

[5] B. Chapman, H. Calandra, S. Crivelli, J. Dongarra, J. Hittinger, S. A. Lathrop, V. Sarkar, E. Stahlberg, J. S. Vetter, and D. Williams. 2014. DOE Advanced Scientific Advisory Committee (ASCAC): Workforce Subcommittee Letter. (7 2014). https://doi.org/10.2172/1222711

[6] DOE. 2023. *Predictive Science Academic Alliance Program.* Retrieved January 19, 2023 from https://science.osti.gov/initiatives/RENEW

[7] ECP. 2023. *ECP Broadening Participation Initiative.* Retrieved January 11, 2023 from https://www.exascaleproject.org/hpc-workforce/

[8] ECP. 2023. *Sustainable Research Pathways.* Retrieved January 11, 2023 from https://shinstitute.org/sustainable-research-pathways-2022/

[9] M. Haddara and H. Skanes. 2007. A reflection on cooperative education: From experience to experiential learning. *International Journal of Work-Integrated Learning* 8, 1 (2007), 67.

[10] J.K. Holmen. 2022. *Portable, Scalable Approaches For Improving Asynchronous Many-Task Runtime Node Use.* Ph. D. Dissertation. School of Computing, University of Utah.

[11] J. K. Holmen, B. Peterson, and M. Berzins. 2019. An Approach for Indirectly Adopting a Performance Portability Layer in Large Legacy Codes. In *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)* (Denver, CO, USA). 36–49. https://doi.org/10.1109/P3HPC49587.2019.00009

[12] J. K. Holmen, D. Sahasrabudhe, and M. Berzins. 2021. A Heterogeneous MPI+PPL Task Scheduling Approach for Asynchronous Many-Task Runtime Systems. In *Proceedings of the Practice and Experience in Advanced Research Computing 2021 on Sustainability, Success and Impact (PEARC21).* ACM.

[13] J. K. Holmen, D. Sahasrabudhe, and M. Berzins. 2022. Porting Uintah to Heterogeneous Systems. In *Proceedings of the Platform for Advanced Scientific Computing Conference* (Basel, Switzerland) *(PASC '22).* Association for Computing Machinery, New York, NY, USA, Article 11, 10 pages. https://doi.org/10.1145/3539781.3539794

[14] HPCWire. 2023. *DOE Announces $56M for Traineeships Supporting Historically Underrepresented Groups and Institutions.* Retrieved January 11, 2023 from https://www.hpcwire.com/off-the-wire/doe-announces-56m-for-traineeships-supporting-historically-underrepresented-groups-and-institutions/

[15] J. K. Huggins. 2010. Engaging computer science students through cooperative education. *ACM SIGCSE Bulletin* 41, 4 (2010), 90–94.

[16] A. Humphrey. 2019. *Scalable Asynchronous Many-Task Runtime Solutions to Globally Coupled Problems.* Ph. D. Dissertation. School of Computing, University of Utah.

[17] A. Humphrey, T. Harman, M. Berzins, and P. Smith. 2015. A Scalable Algorithm for Radiative Heat Transfer Using Reverse Monte Carlo Ray Tracing. In *High Performance Computing*, J. M. Kunkel and T. Ludwig (Eds.). Springer International Publishing, Cham, 212–230.

[18] M James, E Bertschinger, B Beckford, T Dobbins, S Fries-Britt, SJ Gates, J Isler, M Ong, A Richardson, Q Williams, et al. 2020. The time is now: systemic changes to increase African Americans with bachelor's degrees in physics and astronomy. *American Institute of Physics* (2020). https://www.aip.org/sites/default/files/aipcorp/files/teamup-full-report.pdf

[19] M. A. Leung. 2018. Developing Sustainable Methods for Broadening Participation by Transforming Mainstream Science and Technology Communities Through the Normalization of Inclusion. *American Behavioral Scientist* 62, 5 (2018), 683–691. https://doi.org/10.1177/0002764218768844

[20] LLNL. 2022. *Predictive Science Academic Alliance Program.* Retrieved December 29, 2022 from https://psaap.llnl.gov/

[21] B. J. McCartin, J. J. Salacuse, and D. Green Jr. 1998. A Well-Kept Secret: The Kettering University Experience. *SIAM News* 31, 6 (1998).

[22] C. R. Noyes, J. Gordon, and J. Ludlum. 2011. The academic effects of cooperative education experiences: Does co-op make a difference in engineering coursework?. In *2011 ASEE Annual Conference & Exposition.* 22–1428.

[23] OLCF. 2022. *Pathways to Supercomputing.* Retrieved December 29, 2022 from https://www.olcf.ornl.gov/community/pathways-to-supercomputing/

[24] B. Peterson. 2019. *Portable and Performant GPU/Heterogeneous Asynchronous Many-task Runtime System.* Ph. D. Dissertation. School of Computing, University of Utah.

[25] D. Sahasrabudhe. 2021. *Enhancing Asynchronous Many-Task Runtime Systems for Next-Generation Architectures and Exascale Supercomputers.* Ph. D. Dissertation. School of Computing, University of Utah.

[26] P. J. Smith. 2020. *PSAAP-II Historical Report: The UQ-Predictive Multidisciplinary Simulation Center for High Efficiency Electric Power Generation with Carbon Capture.* Technical Report. University of Utah. https://home.chpc.utah.edu/~u0033047/CCMSC_Final_Report.pdf

[27] J. Wells, H. Bozeman, B. Keller, A. Miyaoka, V. Sergueev, A. von Glatz, and K. Wills. 2020. *U.S. Department of Energy Computational Science Graduate Fellowship: 1991-2021, A follow-up study of recipients and programmatic outcomes.* Technical Report. Westat. https://www.krellinst.org/doecsgf/docs/2022_DOE_CSGF_Longitudinal_Study-Web.pdf

[28] R. Wyonch. 2020. Work-ready graduates: The role of co-op programs in labour market success. *CD Howe Institute Commentary* 562 (2020).