Rosen Center for Advanced Computing



Annual Report 2025

PURDUE UNIVERSITY | Rosen Center for Advanced Computing

Director's Welcome

Dear colleagues at Purdue and beyond,

I'm excited to welcome you to this year's report highlighting the impact of our Rosen Center for Advanced Computing to computational science happening both within Purdue and across the nation. Through 2025 and forward, my message to our partners and potential partners centers around ***"we have resources here at Purdue",*** as we support Purdue Computes, One Health, Purdue in Indianapolis, and the Daniels School of Business. Computing is a critical ingredient in STEM education in the 21st century, and these resources are not only used by researchers across Purdue, but 32% of last year's PhD graduates used HPC during their career!

20+ years into our community cluster program, we provide powerful but cost-effective high-performance computing resources that now have pivoted into AI, with our new system "Gautschi", named after CS luminary Walter Gautschi. Compute resources are coupled with a diverse portfolio of data offerings that even as I type are seeing the 3rd generation refresh of our Data Depot system that will be larger and faster than ever. But the most important value is our people: a stable of domain science experts to help apply computing to a wide variety of subjects ranging from AI, biology, computational chemistry, and augmented reality; research scientists; and research software engineers (RSEs).

Our campus offerings now support restricted data, such as data subject to NIH's now heightened cybersecurity requirements, and we continue to blaze the trail among higher education in supporting controlled unclassified information (CUI) and ensuring research security in our ever-evolving landscape.

Researcher needs always grow with the size and complexity of the science, however, and we're prepared to serve as your on-ramp to national-scale platforms via the NSF's ACCESS and NAIRR infrastructures, the Department of Energy's leadership class systems via our partnership with Argonne National Lab's Lighthouse program, and commercial clouds.

Finally, these pages will also outline the impact of RCAC's role as a partner – our Anvil system supports the nation's AI capabilities now as a NAIRR resource provider, and our research scientists and RSEs have had several successes this year with new grant awards and fresh partnerships supporting Purdue PIs. I hope that some of these stories may plant the seed for a new collaboration idea!

Thanks again for making us your cyberinfrastructure partner!

Preston Smith

Mission & Vision

Mission

Purdue IT provides the technology infrastructure, services, solutions, and information security that support teaching and learning, enhance research, and enable faculty and staff to achieve their objectives while providing a positive student experience.

Vision

To empower giant leaps across Purdue by providing safe, efficient and reliable services in our pursuit to become the benchmark for IT in higher education.

Table of Contents

Gautschi, Supporting the Next Giant Leap in AI Research

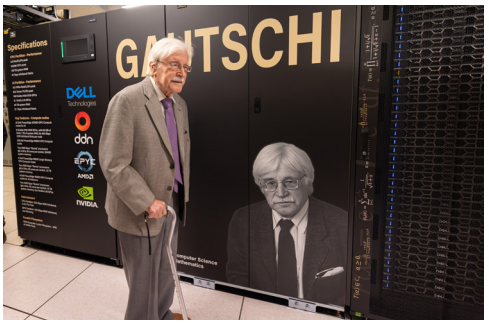

In late 2024, RCAC unveiled its newest and most powerful supercomputer to date, Gautschi. The Gautschi cluster was designed to provide Purdue researchers with a world-class computing resource capable of driving the university toward its next giant leap. Eponymously named in honor of Walter Gautschi, Professor Emeritus of Computer Science and Professor Emeritus of Mathematics at Purdue University, the Gautschi supercomputer consists of two partitions—a traditional HPC resource focused on providing next-generation CPUs and a dedicated AI partition containing Nvidia H100 SXM GPUs.

Gautschi-AI Stats:

25 *PIs from* **11** *departments have purchased allocations*

11 *IPAI affiliate PIs applied for matching allocations*

148k *GPU hours delivered in past quarter*

Average of $2.17M to procure at commercial cloud prices

The Gautschi cluster was built through a partnership with Dell, AMD, and Nvidia, thanks to support from Purdue Computes and the Institute for Physical AI (IPAI). It debuted at number 157 on the Top500 list of the world's most powerful supercomputers and number 43 on the Green500list of the most energy-efficient supercomputers. The system is currently a Top-7 academic system in the US, and the 5th most powerful campus-serving cluster.

System Specs:

In total, the Gautschi cluster consists of:

- 338 Dell compute nodes with two 96-core AMD Epyc "Genoa" processors (192 cores per node) and 384 GB of memory,
- Six large memory nodes with 1.5 TB of memory,
- Six AI inference nodes with two Nvidia L40S GPUs, and
- 20 Dell PowerEdge XE9680 compute Nodes w/ 2x 56-Core Intel Xeon Platinum 8480+ CPUs and 8 Nvidia H100 SXM GPUs with 80 GB of HBM3 memory. Each card has its own non-blocking 400 Gbps NDR link.

The AI partition, Gautschi-AI, is comprised of 20 Dell PowerEdge XE9680 compute nodes, each with Dual 56-Core Intel Xeon Platinum 8480+ CPUs, 8 Nvidia H100 SXM GPUs with 80 GB of RAM, and 8x non-blocking 400 Gbps NDR links. This massive acquisition of H100 GPUs provides Purdue researchers with a whopping 10.7 PetaFLOPS of peak performance. These state-of-the-art GPUs utilize NVIDIA's Hopper architecture and a Transformer Engine in order to provide training and speeds that are four times faster than previous generation models. And the unique, 8-way connected configuration of the nodes means that a researcher can leverage tens of GPUs, if that is what their job requires. In anticipation of rising demand from AI researchers, Gautschi-AI is designed to be expandable up to 32 nodes/256 GPUs. The system also features a high-speed all-flash parallel filesystem with 7PB of capacity, which will ensure that its CPUs and GPUs remain fed with data.

Scientific fields are constantly evolving. As knowledge is expanded and new technology emerges, research methods are in a continual state of flux. AI-based workflows are one of the newest methodological advancements and have quickly been adopted by every field of scientific computing. From Large-Language Models (LLMs) and Convolutional Neural Networks (CNNs), to protein structure prediction, robotics, and computer vision, intensive AI workloads are becoming integral to research across the scientific spectrum. An AI cluster like Gautschi will ensure that Purdue researchers have the necessary tools to make new discoveries.

HPC Impact on Courses

During the 2024-25 Academic Year (Fall 2024 – Summer 2025 semesters), 121 courses used HPC for instructional purposes, impacting a total of over 4700 students. (121 courses, 5500 students used HPC in AY23-24)

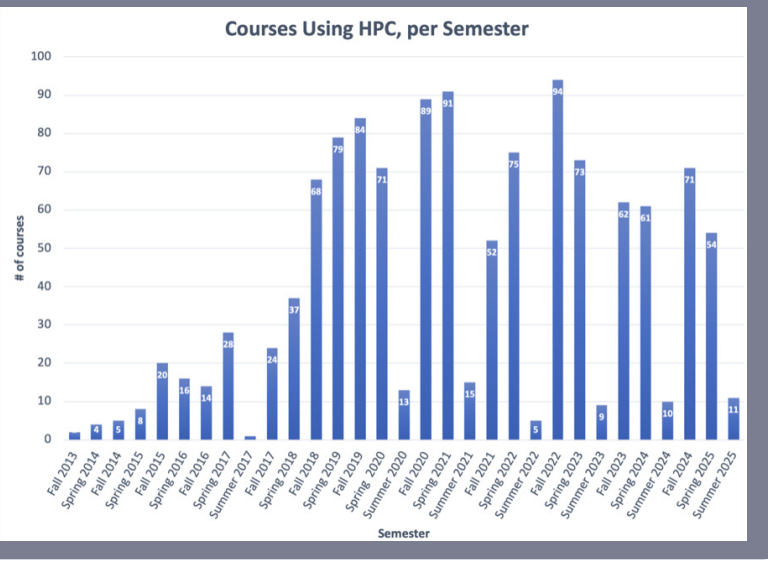

Anvil Year Four

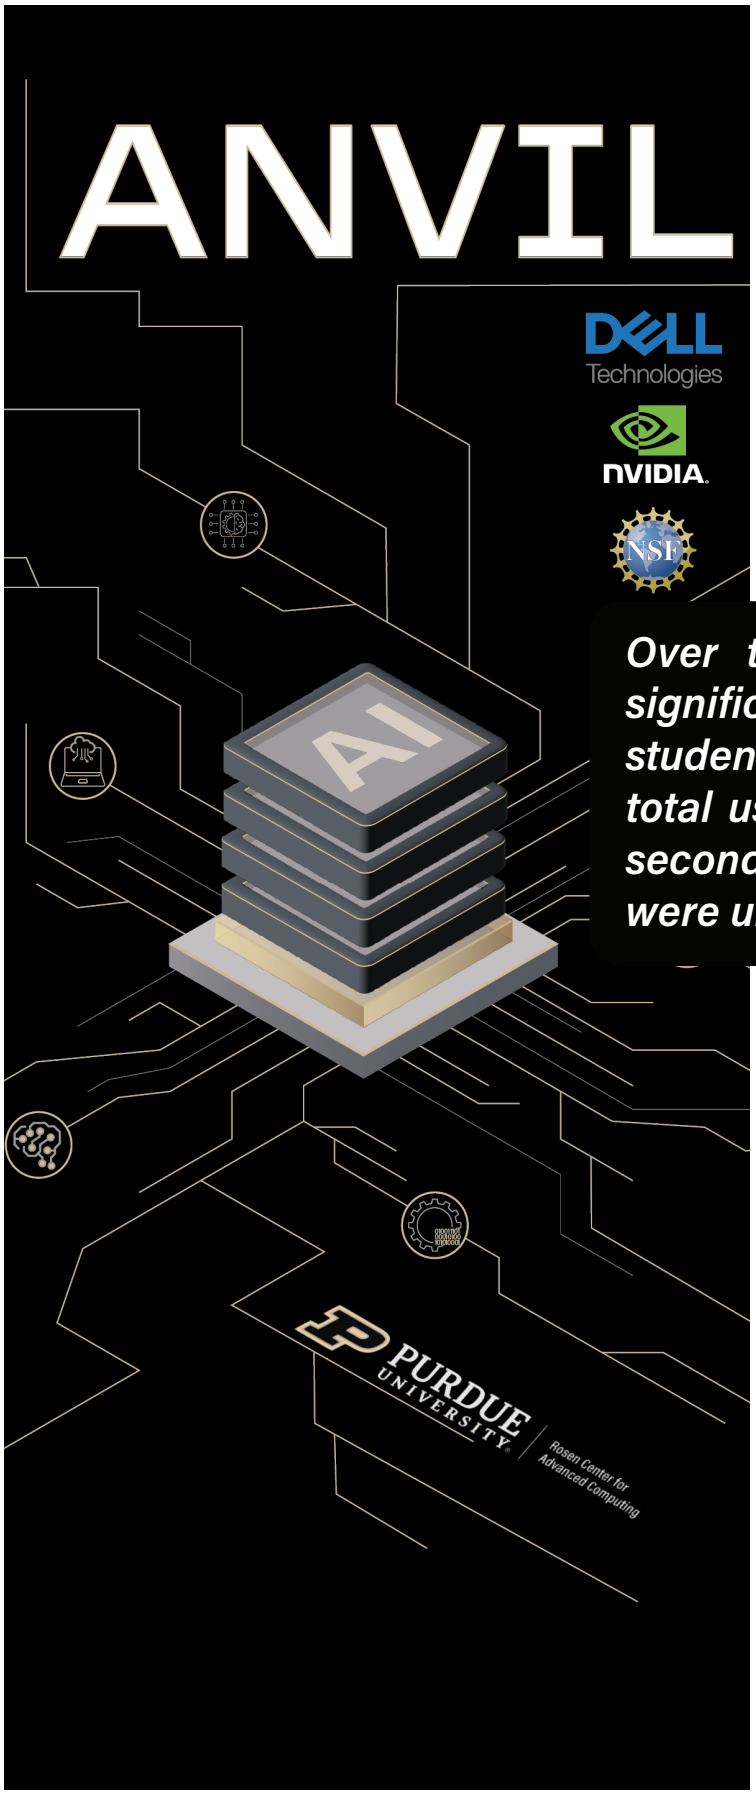

Anvil, one of Purdue's most powerful supercomputers, continues its pursuit of excellence in HPC as it enters its fourth year of operations. Funded by a $10 million acquisition grant from the National Science Foundation (NSF), Anvil began early user operations in November 2021 and entered production operations in February 2022. After three years online, Anvil has more than proven its value. The supercomputer has been used to help over 12,000 researchers push the boundaries of scientific exploration in a variety of fields, including artificial intelligence, astrophysics, climatology, and nanotechnology. This past year was also marked by an explosion of growth for Anvil, both in machine size and usage statistics. Thanks to supplemental funding from the NSF's National Artificial Intelligence Research Resource (NAIRR) Pilot, the Anvil AI partition was added to the supercomputer and brought online. A total of 84 Nvidia H100 SXM GPUs were procured and added to the system. With this upgrade, Anvil is now poised to deliver a world-class AI supercomputing resource to researchers nationwide.

Over the past three years, Anvil has had a significant impact on scientific research and student development. With more than 12,000 total users thus far (double the number from its second year of operations), of which over 6,000 were undergraduate students.

Over the past three years, Anvil has had a significant impact on scientific research and student development. With more than 12,000 total users thus far (double the number from its second year of operations), of which over 6,000 were undergraduate students (another twofold increase), Anvil is not only helping meet the growing need for high-performance computing (HPC) within the realms of research, but also actively assisting with the development of cyberinfrastructure professionals of tomorrow. Overall, Anvil has allowed users access to 1.8 billion CPU hours and 1.8 million GPU hours, supporting research across 65 diverse scientific domains. In 2024 alone, 165 research publications (a ~2.3x increase from 2023) cited Anvil usage. Aside from the supercomputer itself, the Anvil team has been hard at work promoting the benefits of HPC and ensuring the nation has a workforce trained in the use, operation, and support of advanced cyberinfrastructure. In its three years of operations, the Anvil team has participated in 72 outreach events and conducted 34 training sessions, with a multitude already planned for the coming year. These training sessions are designed to deliver working knowledge of HPC systems and teach users how to get the most out of their research time on Anvil. The team also provided hands-on training to students through initiatives such as the Anvil Summer REU program and RCAC's CI-XP student program, which allowed the students to gain much-needed knowledge and experience in the field of HPC.

ANVIL

FORGING THE FUTURE OF COMPUTING

1/1/2022 - 8/1/2025

2.4 Billion
Allocated CPU hours

2.4 Million
Allocated GPU hours

1211
Allocations

14,319
Anvil users

72
Fields of Science

6,302
Undergraduate student users

153
Outreach events

36
Students directly supported

279,803
Interactive Jobs

2,611
People impacted

79
Training sessions

Open OnDemand Users **5,091**

PURDUE UNIVERSITY | Rosen Center for Advanced Computing

Anvil is funded by NSF grant #2005632

Anvil Tech Specs

Anvil is a supercomputer deployed by Purdue's Rosen Center for Advanced Computing (RCAC) in partnership with Dell and AMD. The system was created to significantly increase the computing capacity available to users of the NSF's Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS), a program that serves tens of thousands of researchers across the United States. Before the new expansion, Anvil's system consisted of 1,000 Dell compute nodes, each with two 64-core third-generation AMD EPYC processors, 32 large memory nodes with 1 TB of RAM per node, and 16 GPU nodes, each with four NVIDIA A100 Tensor Core GPUs, all of which are interconnected with 100 Gbps Nvidia Quantum HDR Infiniband. The new NSF NAIRR funding has added 21 Dell PowerEdge XE9640 compute nodes, each with 4 Nvidia 80GB H100 SXM GPUs, as well as an additional 1 PB of flash-based object storage integrated into Anvil's composable subsystem. The new GPU nodes also feature an additional NDR Infiniband fabric to support larger AI workloads.

> **This H100 GPU expansion not only gives Anvil a significant boost to the amount of resources available to the NAIRR Pilot users,but also provides a major increase in Anvil's GPU computing power.**

"Anvil joined the NAIRR Pilot as a resource provider in May of 2024" says Rosen Center Chief Scientist Carol Song, principal investigator and project director for Anvil. "We made available Anvil's discretionary capacity, which was allocated entirely to researchers, right away. This H100 GPU expansion not only gives Anvil a significant boost to the amount of resources available to the NAIRR Pilot users, but also provides a major increase in Anvil's GPU computing power. The H100 GPU outperforms the current A100 GPU in Anvil by as much as nine times in computing speed. Many workloads, especially AI model training and inference, will run much faster, reducing the time-to-results for researchers."

In 2024, GPU capabilities were upgraded for the Anvil Composable Subsystem of the Anvil supercomputer. The Anvil Composable Subsystem now hosts eight composable nodes, each with 64 cores and 512 GB of RAM, and multiple GPU nodes with a total of 4 NVIDIA A100 80GB GPUs and 4 NVIDIA H100 96GB GPUs. The Anvil Composable Subsystem is a Kubernetes-based private cloud managed with Rancher that provides a platform for creating composable infrastructure on demand. This cloud-style flexibility allows researchers to self-deploy and manage persistent services to complement HPC workflows and run container-based data analysis tools and applications. The composable subsystem is intended for non-traditional workloads, such as science gateways and databases, and the addition of the composable GPU node supports tasks such as AI inference services and model hosting.

Anvil Innovations

The Anvil supercomputer has been host to a number of innovations throughout the past year. From on-premises generative AI, to a Jupyter Notebook platform, to increased datasets and a streamlined, user-friendly dashboard, the Anvil team has strived to provide researchers with the best cutting-edge tools to help advance their work. These innovations include:

AnvilGPT: AnvilGPT is a large language model (LLM) service that makes open-source LLM models like LLaMA accessible worldwide to ACCESS researchers. Unlike other LLM services, AnvilGPT is hosted entirely with on-premises (on-prem) resources at Purdue. This means researchers have more democratized access to LLMs, as well as more control. AnvilGPT is hosted on the Anvil Composable Subsystem and leverages the powerful H100 GPUs for rapid processing. The service was designed to provide a secure, central, and flexible AI platform tailored for Anvil users. Anyone with an Anvil allocation has access to AnvilGPT for free.

Anvil Notebook Service: The Anvil Notebook Service is a cloud-based, scalable platform for web-based Jupyter Notebooks. It offers access to CPU and GPU resources through a variety of Jupyter notebooks supporting Python, R, Julia and popular machine learning frameworks like Tensorflow and PyTorch. The notebook service is also tightly integrated with the Anvil HPC system, allowing users to interact with data stored on Anvil and submit jobs to Anvil's batch system.

Scaling Anvil Composable: With the addition of AnvilGPT and the Anvil Notebook Service, as well already hosting 12 Science Gateways with various scaling requirements, Anvil has seen an ever-increasing demand for Kubernetes infrastructure. To combat this heightened demand (which often exceeded the Kubernetes resource capacity), the Anvil team has developed an automated batch to Kubernetes conversion process. This process utilizes idle batch nodes on the Anvil HPC system to increase Kubernetes resources, which not only allows Kubernetes to perform at scale, but also maximizes the use of Anvil's 1000+ node capacity. The ability to scale the composable system has already been used to great effect:

- **NanoHUB STARS Workshop**
  - NanoHUB staff integrated their hub with Anvil Composable to scale out tool sessions
  - Supported 75 participants launching tool sessions with 4C and 16GB RAM
- **CyberFACES (NSF CyberTraining)**
  - Custom JupyterHUB supporting 100s of participants
- **Purdue DataMine (Anvil Notebook Service, 2025)**
  - 1200+ students currently using Anvil batch to launch notebooks

Open OnDemand Dashboard: As part of their Anvil REU experience, undergraduate students Richie Tan and Anjali Rajesh developed an Anvil web dashboard to highlight Anvil usage metrics and make complex information more accessible to Anvil users. By creating this dashboard, Rajesh and Tan provided Anvil users the ability to effortlessly tap into relevant metrics that can help them understand how they are using their computational resources and how they can improve their performance without any coding or command-line confusion. The dashboard has been so successful that it has been shared with the OpenOnDemand project (which it was built on) for broader use. Some of its key features include:

- **Homepage widgets** showing service units, disk usage, queued jobs, etc.
- **My Jobs** page for a comprehensive view of recent jobs on Anvil.
- **Performance Metrics** page for job performance summary over specific periods of time.
- **In-memory caching** for API requests.

Datasets: The Anvil team has incorporated popular domain-specific datasets onto the system to optimize user workflows. A module system enables searching by dataset category; e.g. hydrological models, geospatial models, etc. The team also included automatic web-based documentation generation for future discoverability and search function. Perhaps the biggest innovation within the datasets is the conversational search made available through the dataset query. This enables a context-sensitive chat function that summarizes information from various dataset documents and works across multiple domains.

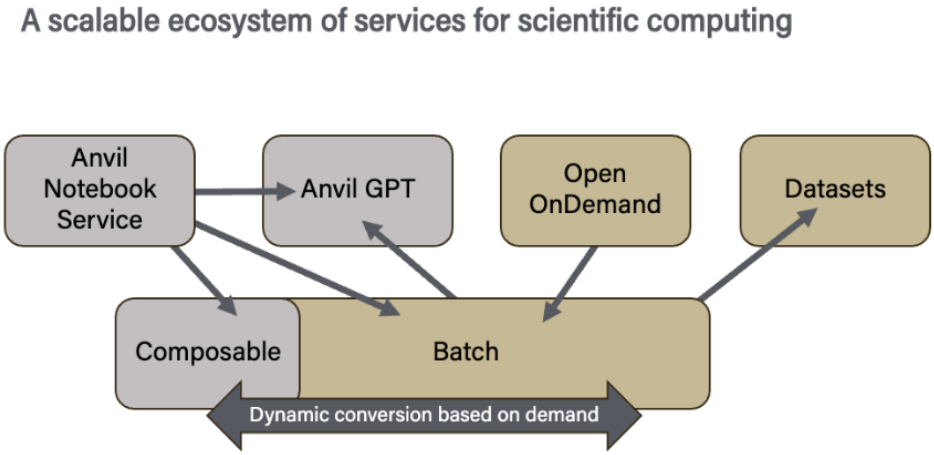

Research Software Engineering

The Purdue Center for Research Software Engineering (aka the RSE center) is a university-approved center within the Rosen Center for Advanced Computing. Its official establishment recognizes the increasingly vital role that software plays in all fields of scientific research, and formalizes RCAC's software engineering efforts at RCAC to better support research at Purdue. The RSE center's mission is to help accelerate research and increase its impact through the creation of innovative, robust, and sustainable research software. Two teams form the RSE Center: the Scientific Solutions Group (SSG) and the Envision Center (EC). SSG focuses on developing innovative cyberinfrastructure solutions, while the EC provides novel solutions to effectively communicate complex research concepts. Often, both teams will work together on a project, providing partners with world-class software design and visualization tools, and enabling them to take their next giant leap in research.

What is a Research Software Engineer?

A Research Software Engineer (RSE) combines professional software engineering expertise with an intimate understanding of research. These individuals design, write, and maintain software that is used to support, disseminate, or undertake research.

- **Proposal collaboration**
- **Software design**
- **Visualization**
- **Code Development**
- **Website and Web Application Development**
- **Training**
- **Consulting**
- **Software configuration and deployment**
- **Research software**
- **High-performance computing**

Notable Projects

ACID-R: ACID-R (Automated Commercial Industry Data – Repository) is designed to help the Air Force identify promising new technologies with more efficiency by harnessing the power of AI, but with no risk of hallucination (AI-fabricated false information). Currently, the Air Force receives many new technology proposals through what could be considered spam emails, with the relevant information buried within PDFs that are upwards of 20 pages in length. ACID-R is an all-encompassing platform for connecting commercial vendors with the DOD department. Vendors will be able to upload their capability statements, the Air Force will be able to view, search, and filter through them all (thousands of proposals), and the AI-powered software will extract the relevant information from each proposal to highlight what the Air Force needs to know without having to manually review each statement. ACID-R will also enable vendors to create more effective proposals by automatically informing them of any missing information that the Air Force needs.

Biobattery: Biobattery is designed to enhance autism diagnosis through real-time eye tracking technology. The tool works by inputting the eye-tracking data acquired during an autism evaluation, analyzing the data, combining it with clinical observations from a doctor or other professional, and generating a report to guide a clinician in diagnosing autism. Phase 1 of this project was a success, with phase 2 due to begin soon. Phase 2 will see the creation of a fully realized, prototype, autism diagnosis via eye-tracking software product that integrates with specific hardware and is suitable for use in a clinical setting by trained, unsupervised technicians.

RSE Stats

Since its inception as a university approved center last year, the RSE Center has completed more than 20 projects, with 23 others currently active. These projects have required a wide-range of specialized skillsets, needing help from our AI and data scientists, visualization experts, and research software engineers. The RSE Center is committed to ensuring that the software used in research is robust, reliable, and sustainable, and that research projects can achieve their goals as efficiently and effectively as possible.

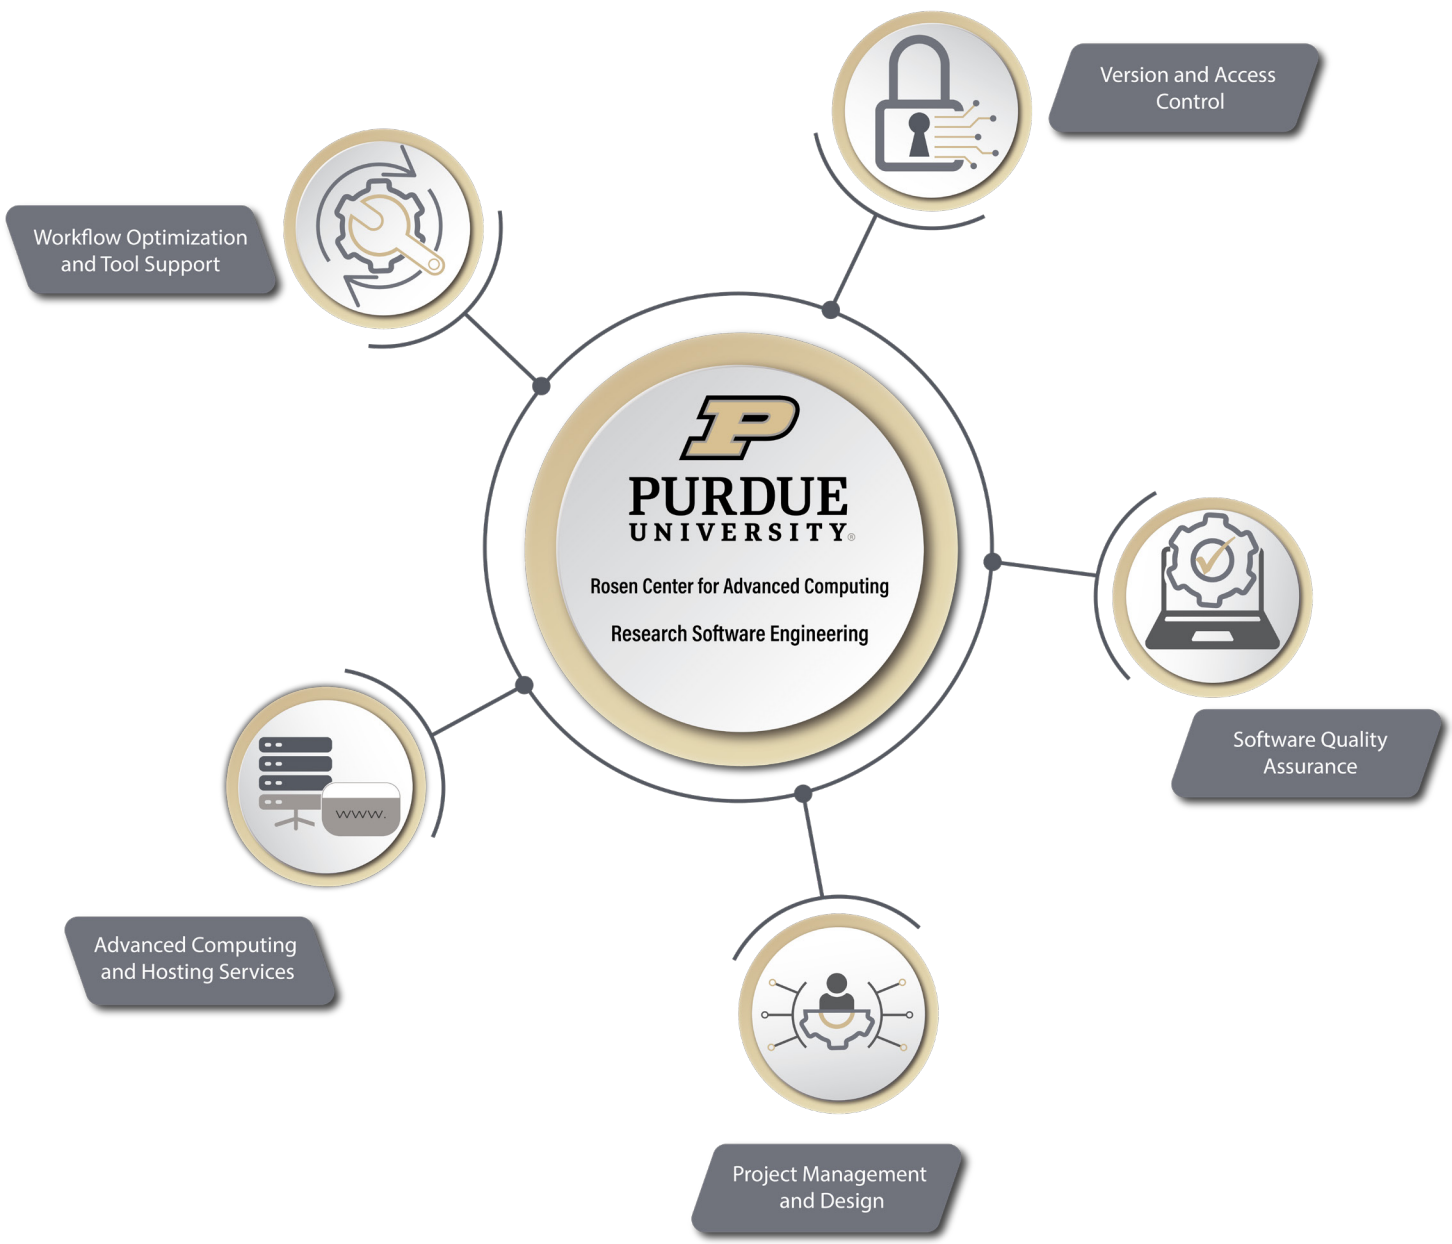

Industry Partners

This past year saw an explosion of growth for RCAC's Industry Partnership program. Part of this program is designed to enable sponsorship and support for RCAC events, such as the CI Symposium. Some of the biggest sponsors throughout the year have been:

- **DDN**: Engaged as a partner in a year-long sponsorship at the "First National Supercomputer Resource, 1984" level. Their gift supported events throughout the year, starting with our third annual CI Symposium last October. They have been a sponsor at this level for 2 years consecutively.
- **Pier Group**: Sponsored the October CI Symposium. This was their second year sponsoring this event, and their gift was instrumental in providing food and drinks for our networking reception.
- **Lenovo**: Sponsored the Spring CI Symposium held at Purdue in Indianapolis. Their gift enabled the post-symposium reception at High Alpha. Lenovo also provided a speaker for the event.

The Industry Partnership program allows industry users to utilize the Anvil supercomputer for their business needs, but at a fraction of the cost of private HPC companies. Examples of some of the current Industry Partnership users, as well as projects under discussion, include:

- **MyRadar:** high resolution weather prediction
- **BlueWave AI Labs:** AI/ML for nuclear plant operational, regulatory efficiency improvements
- Smart building technology company (Kubernetes GPU workloads)
- Non-profit engaged in clean energy research (I/O intensive workloads)
- LLM-based tool for conversational assistance during emergencies
- AI-driven platform for airport power infrastructure management for electric aircraft
- Electromagnetic propulsion systems
- Generative AI for personalized content
- Medical research company working on blood test-based cancer detection
- Life sciences diagnostics company
- Technology company aimed at early detection of TBI and cognitive impairment

Envision Center

The Envision Center—one of the two branches of Purdue's RSE Center—is dedicated to assisting, supporting, and collaborating with faculty, students, and industry in scientific visualization, virtual and augmented reality, and media creation. The center's staff and student employees work with faculty partners and external clients to create virtual reality (VR) and data visualization tools for research and educational use. Throughout its 21-year history, Envision has worked on a wide variety of projects, from creating virtual labs for aeronautics and electrical engineering to developing a visualization platform that can illustrate 3D supernova remnants for classroom usage and more. The center also collaborates on grant proposals and develops promotional media such as publication-quality stills and animated videos.

CollabXR

The Envision Center made waves this past year with the development and real-world testing of its CollabXR platform. CollabXR is a shared learning platform that puts advanced visualizations at the fingertips of professors and students, immersing them in a learning environment like never before. Throughout the Spring 2025 semester, six different university instructors took advantage of an Innovation Hub-supported pilot program to utilize CollabXR and deliver world-class educational experiences to students. These six projects marked the first large-scale implementation of CollabXR in the classroom.

"Utilizing CollabXR in these different classroom settings was an amazing opportunity for us," says George Takahashi, Principal Visualization Scientist at the Envision Center. "Not only was the platform a huge success, enabling the professors to relay complex information in a very fun and impactful manner, but we took away lessons that will allow us to improve functionality for future courses."

The CollabXR (XR=extended reality) platform is a shared environment that allows anyone in a headset to view and interact with the same virtual content pieces together, viewed in fully virtual (the entire field of vision is virtual, obscuring all of the real environment, and other participants appear as virtual avatars) or passthrough augmented reality (cameras feed the real room and people to the headset view, overlaying virtual content to appear to exist in the same space). Developed in conjunction with Professor Danny Milisavljevic, the CollabXR platform is the first of its kind, merging virtual reality (VR), augmented reality (AR), and mixed reality (MR) capabilities with training, research, and education.


PURDUE

"The CollabXR experience demonstrates the great potential of applying advanced technology for teaching and learning. To the best of my knowledge, this is the first time to show the complicated 3D flow structure—which is very tricky to explain to students in standard lectures—vividly in the eyes of students. The students are amazed to see what they observe and develop an increased interest in learning more about the subject. I am looking forward to working with the Envision Center to further develop this capacity to be adopted into our curriculum."

Jun Chen – Professor, School of Mechanical Engineering

Innovation Highlights

Rossman Community Cluster

Rossman cluster deployed with Protected Data Filesystem as a resource for sensitive or restricted data.

New Data Storage options

RCAC offers updated data storage options which stem from a CC* project that was aimed at closing the gap between campus and cloud storage through the deployment of a central, shared, multi-petabyte Ceph distributed storage system.

Rowdy Testbed

Rowdy serves as a technology testbed designed for the evaluation and benchmarking of cutting-edge computing, networking, and storage architectures, including accelerators and CPUs. It aims to advance computational research by allowing invited users to test and explore new technologies in a controlled environment.

Gautschi Community Cluster

Gautschi is a system optimized for communities running traditional, tightly-coupled science and engineering applications. Gautschi is equipped with both CPU and GPU compute nodes, each designed for specific computational tasks.

Purdue GenAI Studio

An LLM service that makes open-source LLM models like LLaMA accessible to anyone at Purdue. Unlike other LLM services, Purdue GenAI Studio is hosted entirely on-premises using resources within Purdue's supercomputers, providing researchers with more democratized access to LLMs, as well as more control.

AnvilGPT

The forebearer of Purdue GenAI Studio, AnvilGPT was developed to provide an LLM service that researchers worldwide could easily access and use, and which negates the concern of leaking intellectual property or proprietary data.

ALCF Lighthouse Initiative

RCAC joins the ALFC Lighthouse Initiative, a newly launched initiative from the U.S. Department of Energy's Argonne National Laboratory. Through this partnership, RCAC can now offer Purdue researchers a direct path to access GPU-accelerated nodes on both the Polaris and Aurora supercomputers, enabling innovative science at the largest scale.

Anvil AI

Thanks to the new NSF NAIRR funding, Anvil AI is comprised of 21 Dell PowerEdge XE9640 compute nodes, each with 4 Nvidia 80GB H100 SXM GPUs, as well as an additional 1 PB of flash-based object storage integrated into Anvil's composable subsystem. The new GPU nodes also feature an additional NDR Infiniband fabric to support larger AI workloads.

Hypershell

RCAC staff members developed HyperShell, the ultimate workflow automation tool. This software helps divide and conquer large volumes of discrete tasks to assist with scheduling and managing what's known as "many task computing."

RSE Services expansion

Researchers can now leverage the RSE center's AI specialists and data processing experts, along with taking advantage of its robust grant-funding consultation service.

Argonne Leadership Computing Facility Lighthouse Initiative

The Rosen Center for Advanced Computing (RCAC) is joining forces with the U.S. Department of Energy's Argonne National Laboratory, through the newly launched Argonne Leadership Computing Facility (ALCF) Lighthouse Initiative. Thanks to this partnership, Purdue can contribute its research strengths in areas such as semiconductors, quantum information science, advanced manufacturing, and artificial intelligence to the nation at the highest possible scales, helping to accelerate the United States' ability to translate breakthrough discoveries into real-world innovations.

What is the ALCF Lighthouse Initiative?

The Argonne Leadership Computing Facility (ALCF) Lighthouse Initiative is a program designed to expand ALCF's user community and propel computational science forward by partnering with academic institutions. Created in 2024, the program's goal is to broaden the reach of HPC and actively engage with the next generation of computing professionals. Purdue University is the fourth institution to join the initiative. By partnering with ALCF and becoming a Lighthouse Institution, Purdue now offers researchers access to the most powerful open science resources in the country.

Argonne NATIONAL LABORATORY | Argonne Leadership Computing Facility

"Between our powerful campus resources, the National Science Foundation's ACCESS program where we operate Anvil, and now the ALCF Lighthouse partnership, there's a clear path for researchers to scale up to the most powerful machines in the nation. Our team is ready to assist researchers who are ready to scale up, all the way from their laptop to Aurora." Preston Smith, Executive Director, RCAC

How will Purdue researchers benefit?

As an ALCF partner and Lighthouse Institution, Purdue researchers have additional computational tools to conduct innovative science at scale, enabling Purdue to advance discovery and achieve significant breakthroughs previously unimaginable. Purdue will be given a discretionary allocation for ALCF machines, and the RCAC staff will receive intensive training on the use of these systems. RCAC will then be able to provide support for Purdue faculty and researchers looking to scale their science to leadership-class systems. This includes:

- **Preparing, porting, and testing applications for suitability of ALCF resources**
- **Scaling applications for use on ALCF systems**
- **Aiding in migrating projects from Lighthouse to an ALCF larger allocation, such as INCITE/ALCC/Director's Discretionary for graduating projects**
- **Offering access to specialized training tailored for the Purdue researchers offered by ALCF**
- **Providing information about summer internship opportunities for students**

RCAC will identify and onboard users whose research projects could best utilize the advanced HPC resources provided by ALCF. These users will then be granted access to GPU-accelerated nodes on both the Polaris and Aurora supercomputers. Aurora is one of the world's most powerful supercomputers. As an exascale system, it is able to perform over a quintillion calculations per second.

Excellence at Scale

With this newfound partnership, Purdue has reached the pinnacle of delivering excellence at scale for computational researchers. RCAC can now provide access to the full spectrum of computing performance possible with today's technology. No matter your job size or computational needs, RCAC has the appropriate resource for you.

"We're excited to add the ALCF Lighthouse to the portfolio of resources available to Purdue researchers to solve their scientific problems using both physics-based simulation and AI," says Preston Smith, Executive Director of RCAC and Director of Computing Infrastructure for Purdue's Institute for Physical Artificial Intelligence. "Between our powerful campus resources, the National Science Foundation's ACCESS program where we operate Anvil, and now the ALCF Lighthouse partnership, there's a clear path for researchers to scale up to the most powerful machines in the nation. Our team is ready to assist researchers who are ready to scale up, all the way from their laptop to Aurora."

By providing the university with an avenue for access and support on these leadership-class systems, RCAC is helping Purdue in its persistent pursuit of innovation, driving the university to the forefront of science and technology.

ALCF Resources:

- **Aurora:** Argonne's exascale supercomputer which contains a whopping 10,624 nodes. The machine is equipped with 63,744 GPUs and 8 fabric endpoints per node (84,992 total!).The system is built on the HPE Cray EX supercomputer platform and includes Slingshot 11 interconnect as the network backbone. Please visit the Aurora website to learn more about the system.

- **Polaris:** Polaris is a leading-edge supercomputer initially designed as a testbed for researchers and developers to optimize their codes for use on Aurora. Alongside its role as a testbed, Polaris now fully supports scientific research in its own right. Polaris contains 560 nodes and is built on the HPE Apollo 6500 Gen 10+ platform. Each node has a single 2.8 GHz AMD EPYC Milan 7543P 32-core CPU, complemented by 512 GB of DDR4 RAM, as well as four NVIDIA A100 GPUs. To learn more about the system, please visit the Polaris website.

Science Highlights

Computational materials scientist relies on Negishi cluster for quantum simulations, training machine learning models

Negishi is playing a key role in the research of Purdue assistant professor of materials engineering Arun Mannodi Kanakkithodi and his lab. Mannodi, who is a computational materials scientist currently working on compositional-, structural-, and defect-engineering of semiconductors for enhanced optoelectronic performance, needs to perform quantum mechanics-based density functional theory (DFT) simulations that require much more computational power than a personal computer can deliver. Negishi provides the computing power needed to run these models.

"The fact that we have these very nice clusters and that we have dedicated nodes for my group as well as standby nodes we can use for short jobs has been very important to our success," – Arun Mannodi

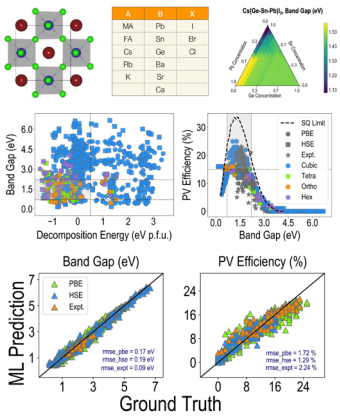

Purdue professor uses Bell Cluster to study college course shutouts

It's a problem familiar to every college student – courses are oversubscribed, and you can't get into a class you need to stay on track – or even worse, to graduate. Now Kevin Mumford, a professor of economics, is using Bell to simulate thousands of course registration outcomes. His goal is to analyze how Purdue's batch registration system impacts students, particularly when they are shut out of essential courses.

By using the Bell cluster, Mumford can conduct large-scale simulations that would be impossible with standard computing methods. With these large-scale simulations, Mumford can quantify the probability of a course shutout for each course requested by each student and estimate how this affects their educational path, particularly in STEM fields.

Bell powers notable studies of seismology, volcano dynamics

Xiaotao Yang, an assistant professor of earth, atmospheric and planetary sciences and a seismologist who uses seismological methods to solve geological and tectonic problems and understand natural hazards, is using Bell for seismic imaging and to better understand the lithosphere, the outermost layer of the Earth, and how earthquakes and volcanoes occur.

Yang and his research team try to understand the long-term evolution of the lithosphere going back hundreds of millions of years and study how the structure of the Earth has been modified by tectonic processes. On a shorter time scale, they also study active volcanoes to understand the magma plumbing system, including its geometry dynamics and how it changes over time before and after eruption.

"We're working on projects that are very computationally intensive, so we need a lot of computational power. This work wouldn't be possible without a supercomputer like Bell," - Xiaotao Yang

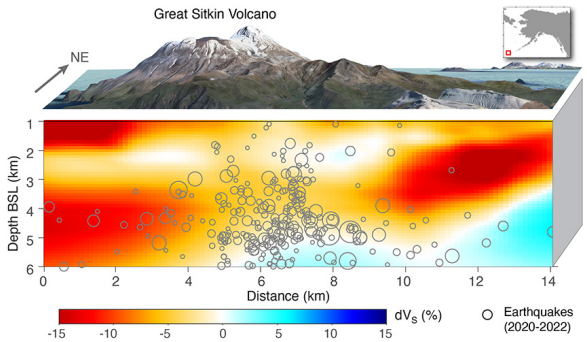

Purdue-connected software company uses Weber cluster in AI, quantum computing tool awarded SBIR DLA grant

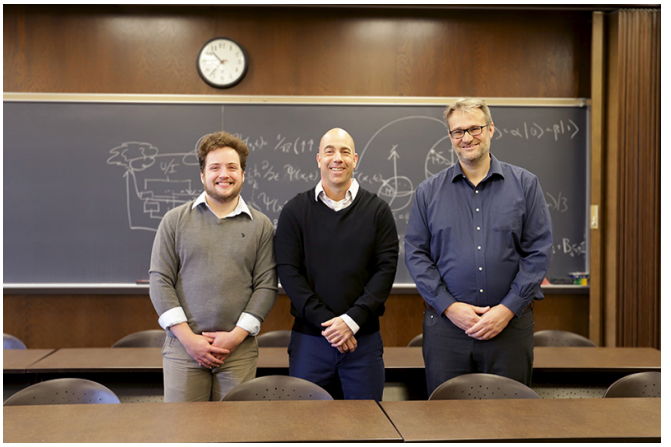

A Purdue-connected software company used Weber to develop a tool that resulted in a Small Business Innovation Research grant from the Defense Logistics Agency (DLA), which contracts, purchases, distributes, stores and disposes of items for the Department of Defense.

The company, Quantum Research Sciences, developed its Product Obsolescence Prediction tool to help the US Air Force with inventory management. The Air Force has many pieces of equipment, all of which have many parts to keep track of with multiple vendors and distribution centers. Figuring out how to manage repair and replacement of all these parts is a very complicated problem that can't be solved with classical computing resources.

The QRS team is able to use RCAC's specialty Weber cluster to develop and train machine learning algorithms. Output from the Weber cluster is capable of improving inventory management technology. Another DoD project in Jung's research group also utilizes the Weber cluster, then fed the output into a quantum computer to get answers back on the best procurement and inventory management strategy.

RCAC partners with faculty member on autism diagnostic tool

A new collaboration between the Envision Center and Brandon Keehn, associate professor of speech, language and hearing sciences and psychological sciences, and the director of Purdue's Autism Research Center, aims to make it easier to detect autism earlier and more accurately.

Together they've developed a tool that can input the eye-tracking data acquired during an autism evaluation, analyze the data, combine it with clinical observations from a doctor or other professional and generate a report to guide a clinician in diagnosing autism. Eye-tracking is a tool for measuring how individuals deploy their attention – what captures their gaze, what they find most engaging. Keehn's lab has developed videos that elicit patterns of gaze that are sensitive to autism and can aid in diagnosis. Usually this eye-tracking data would be analyzed in batches on a powerful computer, necessitating a delay between collection of the data and a diagnosis of autism. But the tool developed by Keehn and RCAC analyzes what is essentially a large, complex dataset and produces a report in seconds or minutes, allowing a clinical provider to make a diagnostic decision in real time.

RCAC's Negishi cluster helps nuclear engineering researcher parallelize and scale electroporation model

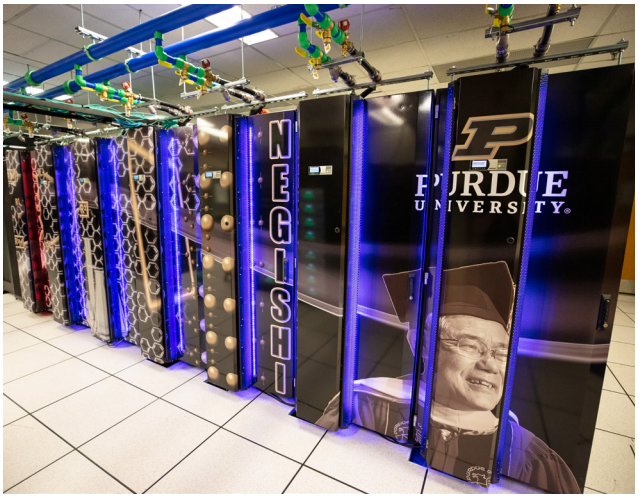

Samuel Wyss, who recently graduated with his master's degree in nuclear engineering under the supervision of Allen Garner, studies a phenomenon known as electroporation, where cell membranes open up in response to a high intensity electromagnetic field. This has important applications to healthcare, allowing for more efficient drug delivery and better cancer treatments. Wyss has collaborated with RCAC staff to refine and scale his electroporation models. While electroporation models exist in the literature, Wyss and his team are the first ones to parallelize them to model cellular level effects.

Running the models on a personal computer, it would take days to model a single set of parameters. Wyss had the idea to parallelize the models on Negishi to accelerate the process and run the model for many different parameters, quickly getting more information that can be used by experimentalists doing studies on real cells.

National Research

Anvil used to advance research on two-phase flows

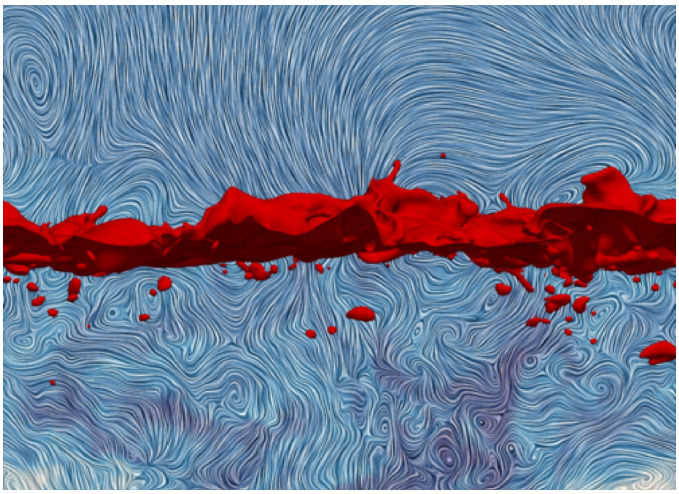

Researchers from The George Washington University used Purdue's Anvil supercomputer to simulate fluid flows in order to elucidate the physics of turbulent bubble entrainment. Understanding this process will lead to practical applications in a variety of fields, including oceanography, naval engineering, and environmental science.

Andre Calado is a Graduate Research Assistant at the George Washington University, working to complete his PhD in computational fluid dynamics. He, alongside his advisor Elias Balaras—a professor in the Department of Mechanical and Aerospace Engineering—wanted to advance the study of two-phase flows (air and water), specifically how turbulence underneath the water interacts with the water's surface and the role this plays in air entrainment. The pair used Anvil to run direct numerical simulations in order to produce high-fidelity simulations of the physics of bubble entrainment. Their work has pushed the boundaries of what has been accomplished so far within two-phase flow research.

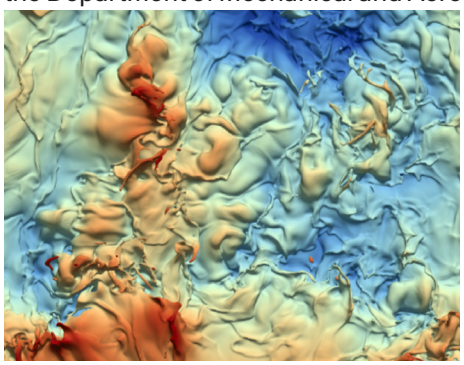

Anvil helps researchers study particle dynamics for hypersonic vehicles

Image courtesy of: https://www.hermeus.com

Dr. Qiong Liu, an Assistant Professor in the Department of Mechanical and Aerospace Engineering at New Mexico State University, along with Irmak Karpuzcu, Akhil Marayikkottu, and Deborah Levin, all from the Department of Aerospace Engineering at the University of Illinois, Urbana-Champaign, have used the Anvil supercomputer to elucidate exactly what happens when particles hit the surface of a hypersonic vehicle in flight.

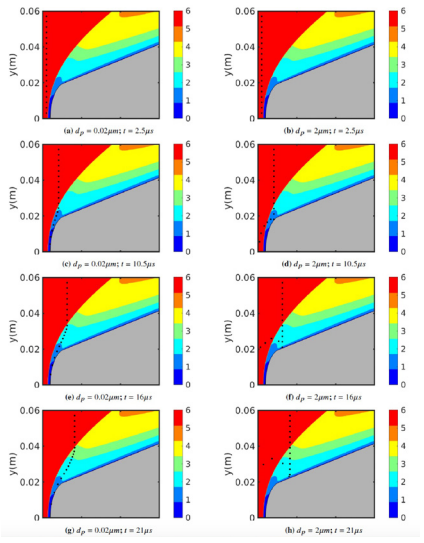

Using the direct simulation Monte Carlo (DSMC) method, the researchers studied the fundamental flow physics and particle trajectory in the flow field around a blunted cone, the most common forebody shape used in hypersonic flight. The simulations looked at particles ranging from .01 micrometers up to 2 micrometers in size. With these simulations, the group was able to determine both the effects that a single particle of varying sizes had on the bow shock and the statistical characteristics of those particles. They found that lighter particles (less than .02 micrometers) could not penetrate the bow shock wave, and so could never directly impact the vehicle. However, heavier particles (greater than .2 micrometers) passed through the bow shock, directly impacted the vehicle, ricocheted upstream, and then traveled downstream in the flow. This unique interaction and motion of heavier particles led to bow shock distortions.

With this study, the research team delivered some much-needed clarity regarding the physics of hypersonic flight, but it was only possible with help from supercomputing resources like Anvil. Liu was thrilled with Anvil's performance throughout the project.

"We are really, really happy with computing on Anvi," says Liu. "The code was well parallelized, so we had no problems running that. Also, the queue was very short, so we were able to submit jobs and get results very quickly. I've actually encouraged many of my new colleagues to apply to use Anvil because I had such a good experience."

Half of the entire Anvil supercomputer used to challenge traditional turbulence theory for space and climate modeling

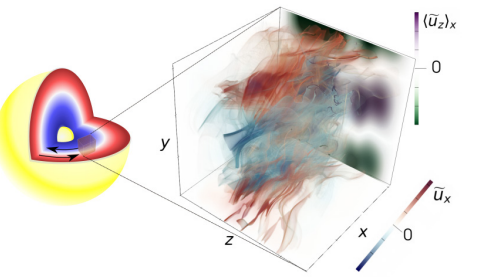

Researchers from the University of Wisconsin (UW)–Madison used Purdue's Anvil supercomputer to study turbulence and turbulent transport in astrophysical plasmas. This research seeks to elucidate the fundamental physics of turbulence, which will have applications across the fields of fluid and plasma dynamics. The group not only pushed the boundaries of scientific research with their work, but also tested the performance limits of Anvil, utilizing upwards of half the machine (512 nodes at once) to run a single simulation.

Bindesh Tripathi, who spearheaded the project, is working toward finishing his doctoral dissertation in the Department of Physics at UW–Madison. Under the joint supervision of advisors Dr. Paul Terry and Dr. Ellen Zweibel, both of whom are professors at the university, Tripathi conducts research involving astrophysics and plasma physics, mathematical/theoretical physics, and numerical methods. Tripathi used Anvil to shed light on the underlying physics of stable-mode excitations within fluid and plasma dynamics, a little understood phenomenon that occurs at large (galactic) scales. To accomplish this task, Tripathi first had to make several bespoke changes to a 3-dimensional (3D) magnetohydrodynamics simulation software known as Dedalus. Then, in order to run the code successfully, the group needed access to an extraordinary amount of computing power, which the Anvil supercomputer was able to provide. To support the researchers' work, the Anvil team set up a special allocation that allowed the group to utilize 512 nodes at once. The group routinely used 30,000 to 40,000 cores simultaneously. To be clear, this was a parallel code, so one single simulation required the use of all of the cores at the same time. This level of computation for a real-world research problem had not yet been tested on Anvil, but the computer was able to handle it with no issues. Tripathi's code ran seamlessly, even at such a large scale, and he was thrilled with the performance of the system.

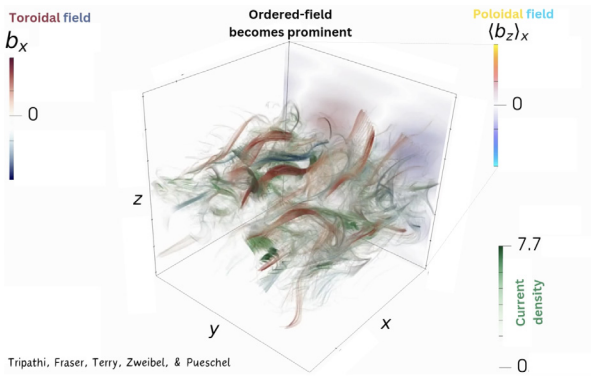

I ran the Dedalus code, and I found it running beautifully well," says Tripathi. "Anvil has a large number of cores, and the queue time was relatively short, even for the very large resources that I was requesting, and the jobs would run quite fast. So it was a quick turnaround, and I got the output pretty quickly. I have had to wait a week or even longer on other machines, so Anvil has been quite useful and easy to run the code. Anvil has also generously provided us with storage of a large dataset, which now amounts to 125,000 gigabytes from my turbulence simulations."

Research Scientists from RCAC were awarded $1.89M of new sponsored research in FY 2025.

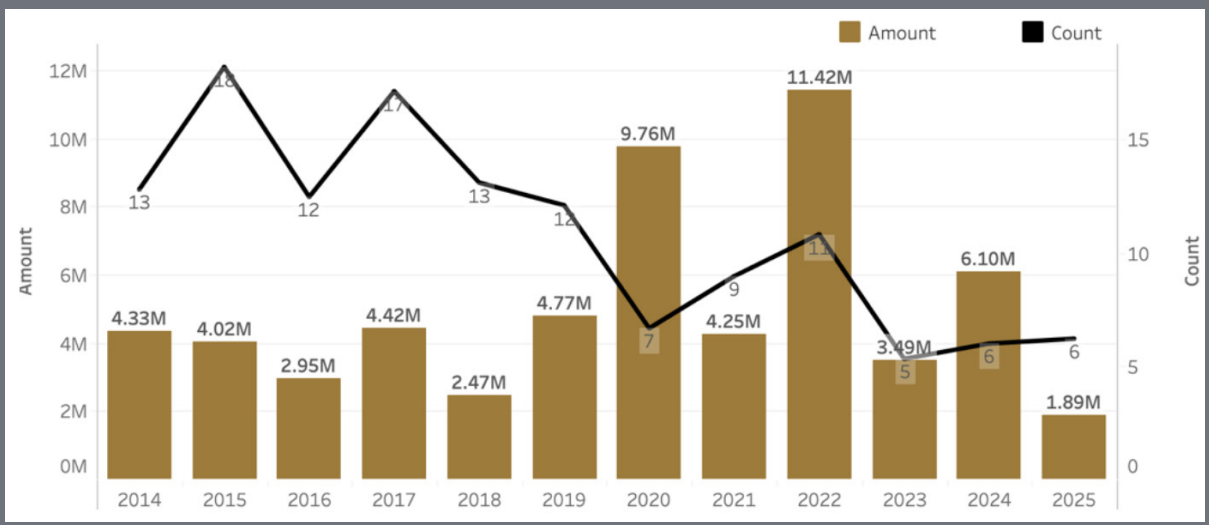

Chipshub and Anvil union is a proven success for semiconductor workforce development

Chipshub, the online platform for everything semiconductors, has finally arrived. After a massive development effort from the nanoHUB team, and with help from the staff at the Rosen Center for Advanced Computing (RCAC), the online platform delivered advanced simulation software to the Purdue Summer Training, Awareness, and Readiness for Semiconductors (STARS) program, a huge step forward for semiconductor workforce development in the U.S.

Chipshub is powered by nanoHUB, the first end-to-end platform for online scientific simulations. Initially supported by the National Science Foundation (NSF), nanoHUB hosts an ever-expanding list of simulation and data tools. Examples include easy-to-use apps that connect to powerful simulations for education in subjects like semiconductor devices, materials, and machine learning, and open courses and free textbooks on topics like nanotransistors, current flow, multiscale materials modeling, and nano-thermal flow theory. The best part? All that's needed to access these tools is a web browser and an internet connection.

"Chipshub extends nanoHUB's success to deliver both open-source and commercial software that supports a semiconductor community through workforce development at scale," said Gerhard Klimeck, Chipshub co-director, Elmore Professor of Electrical and Computer Engineering and Riley Director of the Center for Predictive Devices and Materials and the Network for Computational Nanotechnology. This expansion is strongly aligned with the CHIPS and Science Act and has been supported by NSWC Crane, Purdue Applied Research Institution, and the Silicon Crossroads Microelectronics Commons.

An often overlooked aspect of any successful online platform is the "at scale" piece. In order to support advanced simulations for large user groups, Chipshub partnered with RCAC and utilized the Anvil supercomputer. Anvil, also funded by the NSF, is Purdue's most powerful high-performance computing (HPC) system. Using Anvil's Composable Subsystem, a Kubernetes-based private cloud platform, Chipshub can dynamically scale simulation tools and semiconductor chip design software based on user demand. Tool sessions can be accessed anywhere via a web browser, allowing users to remotely utilize Anvil's high-performance computing resources from their laptops or workstations. On a laptop or personal computer, simulation tools such as the ones provided by Chipshub can take hours to days—sometimes even weeks—to run a single simulation. But, by hosting these tools on Anvil and providing access via the cloud, Chipshub can deliver the power of HPC to hundreds of users at once and significantly cut down time spent waiting for results. This ability to compute at scale allows Chipshub to drive semiconductor workforce development throughout the nation without having to limit classroom size.

"The Anvil Composable subsystem is one of the innovative features of the Anvil supercomputer," says RCAC Chief Scientist Carol Song, principal investigator and project director for Anvil. "It was designed to complement the HPC batch computing system to support a wide range of modern scientific applications. The Chipshub is a perfect example of how the Anvil composable subsystem can provide on-demand computing power at scale through a seamless user experience."

For its inaugural real-world use case in the field of electronic design automation (EDA), Chipshub was utilized by the 2024 STARS summer program. STARS is an eight-week, on-site program offered by Purdue University's College of Engineering. The program is designed to teach undergraduate students deep-tech skills in integrated circuit design, fabrication, packaging, and semiconductor device and materials characterization. One of the two specialized tracks offered by the STARS program was semiconductor chip design. Led by Dr. Mark C Johnson, the chips design track required extensive use of advanced software packages to give students experience in designing, verifying, and submitting a digital integrated circuit for fabrication. Nearly 60 students, including 13 students from the SCALE (Scalable Asymmetric Lifecycle Engagement) workforce development program, were enrolled in the chips design course. It would be nearly impossible for each student to get hands-on experience of semiconductor simulation software using traditional computing resources, but thanks to the support provided by the Chipshub-Anvil collaboration, the course was a huge success. Per Dr. Johnson, "Chipshub proved itself as each member of the STARS cohort were concurrently running simulations, producing chip layouts, and running physical verification continuously for the final three weeks of STARS. Collectively, 12 teams of four to five students each produced a chip design that has been combined into a single layout and will be submitted in September, 2024 for fabrication." Chipshub and Anvil powered 1,800 simulation sessions and 6,000 interactive hours.

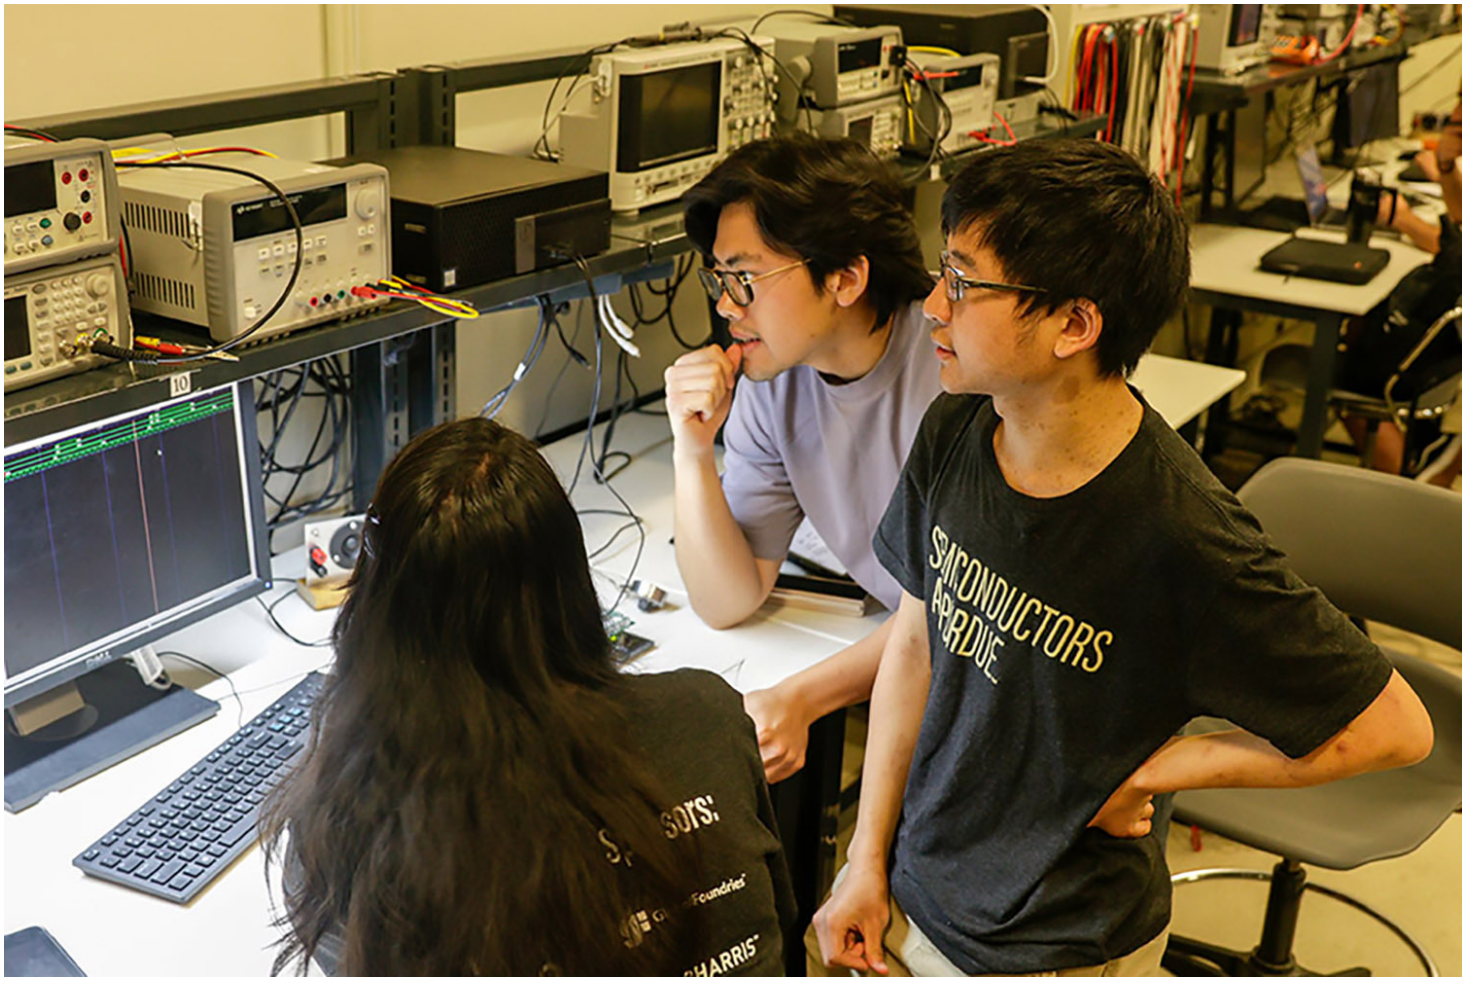

Over the next five years, Chipshub is expected to impact more than 200,000 U.S. engineering students and about 50,000 designers. Now that it has been successfully used with the STARS program, the next steps involve expanding Chipshub access beyond Purdue and making the platform available throughout the country. This would be a boon for the nation, allowing other institutions to utilize these state-of-the-art tools without having to reinvest in the development time and resources already taken on by the Chipshub team.

"There's currently a national effort to reshore a significant chunk of the semiconductor industry—we want to make more chips in the US." Alejandro (Ale) Strachan, Chipshub co-director and Reilly Professor of Materials Engineering. "We want to fabricate them and design them. So there's a need for workforce development in this area. We need to have more engineers, more technicians, more PHDs that are able to work in this field. And so students need to be trained with cutting edge tools, and this is what Chipshub can actually do, at scale."

Notable Software Developments

Award-winning HyperShell software developed by RCAC helps simplify "many task" computing

RCAC staff members have developed a piece of software, known as HyperShell, that helps divide and conquer large volumes of discrete tasks to assist with scheduling and managing what's known as "many task computing." HyperShell has been dubbed "the ultimate workflow automation tool."

In traditional high-performance computing, there's just one big job for the supercomputer to tackle. But more often these days, researchers are doing what's known as "high-throughput computing" or "many-task computing" where the supercomputer is performing a very large number of small, independent tasks, such as analyzing tens of thousands of pieces of data in a dataset.

Traditional HPC job schedulers like Slurm are not well-suited to managing this kind of computing, so RCAC lead research data scientist Geoffrey Lentner saw the need for a piece of software to help facilitate this kind of workflow, and HyperShell was born.

HyperShell, which is written in Python, is useful for researchers doing large volumes of data analysis and processing, such as in fields like bioinformatics, climate modeling, math and statistics optimization and agriculture.

It's also useful for training AI models since the powerful RCAC GPUs are too large for many jobs and Slurm doesn't allow for subdividing them. HyperShell allows a user to stack multiple models on a GPU and fill it up with multiple smaller jobs.

Several tools offer similar functionality to HyperShell but not all together in a single tool with the user ergonomics provided. Novel design elements include but are not limited to (1) cross-platform, (2) client-server design, (3) staggered launch for large scales, (4) persistent hosting of the server, and optionally (5) a database in-the-loop for persisting task metadata and automated retries.

Lentner and former RCAC staff member Lev Gorenstein presented an early version of HyperShell at the 2022 Practice and Experience in Advanced Research Computing (PEARC) conference, and won the best poster award for their poster about it.

Powerful New LLM Services

In collaboration with the Institute for Physical AI (IPAI), Purdue University's Rosen Center for Advanced Computing (RCAC) now offers two powerful new features in its artificial intelligence software services—AnvilGPT and Purdue GenAI Studio. Purdue GenAI Studio and Anvil GPT are large language model (LLM) services that make open-source LLM models like LLaMA accessible to anyone at Purdue (for the former) or with an Anvil allocation (for the latter). Unlike other LLM services, RCAC's newest developments are hosted entirely on-premises using resources within Purdue's supercomputers. This means researchers have more democratized access to LLMs, as well as more control.

AnvilGPT was the first of the two services to be developed. The intent behind the tool was to provide an LLM service that researchers could easily access and use, and which negates the concern of leaking intellectual property or proprietary data. AnvilGPT was made available to anyone nationwide with an Anvil allocation. After its launch earlier in the year, AnvilGPT proved to be a tremendous success, garnering lots of attention and eliciting a higher-than-expected utilization rate. Due to this overwhelmingly positive response, the team at RCAC decided to create a similar tool that could be used by anyone at Purdue, not just ACCESS or NAIRR researchers. Thus, Purdue GenAI Studio was born.

Both AnvilGPT and Purdue GenAI Studio are hosted entirely on-premises, meaning no documents or contexts are uploaded into commercial cloud-hosted AI services. Chats, documents, and models are also not shared between users or used for training. Not only will users feel more secure in their work, but researchers will also be able to use them for grant-funded work that specifically limits the transfer of data to third-party services.

The user interface (UI) for the services is intuitive and easy to use. There are two modalities for interacting with Purdue GenAI Studio or AnvilGPT—both the UI and API,—and additional functionality for both is under active development. The chat interface allows researchers to select from a list of models, including both base models (available to all users) as well as any custom models the individual researcher may have created. Users may also select multiple models to compare the output for any prompt. Additional open-sourced models that are not already listed can be added to either LLM service upon request. Users can also take advantage of the tool's Retrieval-Augmented Generation (RAG) functionality. RAG functionality allows researchers to upload their own data into the model, providing it with more context. This adds to the model's base knowledge, which helps it to provide more accurate responses in particular domains. However, GenAI Studio should not be used with University data classified as sensitive or restricted.

These powerful LLM services highlight the innovation stemming from Purdue's persistent pursuit of excellence within the field of artificial intelligence. Their release helps to bolster the Purdue Computes initiative as well as support national AI research objectives through programs such as National Artificial Intelligence Research Resource (NAIRR) Pilot.

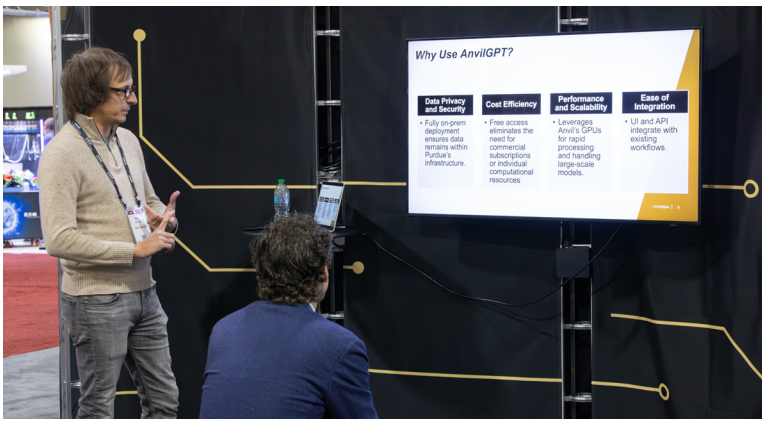

External Funding

- **NAIRR/Anvil AI**
  - The National Science Foundation (NSF) awarded a $4.9 million supplement for Anvil, one of Purdue University's most powerful supercomputers. This funding was given in part to support the National Artificial Intelligence Research Resource (NAIRR) Pilot. This NSF supplement has enabled Purdue to further expand Anvil's graphics processing unit (GPU) compute capabilities to support AI, support high-performance object storage, and provide the operational support needed for users of the NAIRR Pilot.

- **CC* x2**
  - The NSF's Office of Advanced Cyberinfrastructure awarded Purdue University two new Campus Cyberinfrastructure (CC*) grants. The two awards, totaling nearly $1.5 million, will directly address and enhance the cyberinfrastructure at Purdue University.

- **Nasa Grant**
  - RCAC is part of a major NASA grant awarded to develop a new Cyberinfrastructure (CI) tool for post-fire water management and decision-making. This tool, named HydroFlame, will allow researchers to predict the amount of damage wildfires will cause to local watersheds, giving government officials the data they need to protect freshwater supplies. The Principal Investigator (PI) is Adnan Rajib, an assistant professor in the Department of Civil Engineering at the University of Texas at Arlington (UTA). Other team members include individuals from the U.S. Geological Survey and several government and non-government agencies across the western United States. NASA awarded the group $824,020 over three years to develop and refine HydroFlame.

- **StreamCI**
  - RCAC is leading a major NSF grant awarded to create an artificial intelligence (AI)-ready streaming data platform for researchers across domains. This new platform, known as StreamCI, will significantly lower technical barriers associated with harnessing massive data streams, empowering experts from a wide-range of scientific fields to build intelligent and responsive applications that will be more efficient and effective than ever before. The NSF awarded Purdue researchers $4 million over five years to develop and refine StreamCI.

- **MATH Renovations**
  - In August of 2024, the Purdue Board of Trustees gave approval to finance, construct, and award construction contracts for a $16M renovation to the Mathematical Sciences Building Data Center. This long-awaited renovation will result in a 32% increase in floor space, a 60% increase in usable power, and a two-fold increase in cooling capacity, all of which are crucial facility ingredients needed to continue to provide world-class high-performance computing resources to Purdue faculty and the nation.

CI-XP Program Student News

The purpose of the CI-XP (Cyber Infrastructure - eXPerience) Program is to provide work opportunities and real workplace experience that enhance the student's education through the development of professional skills, responsibilities, habits, attitudes, self-confidence, and self-development. CI-XP includes all students in RCAC, Envision Center, the Scientific Solutions Group, and Anvil REU Programs. Over the past year, our students have not only learned a lot, but have been incredibly impressive regarding their work and accomplishments.

RCAC student improves Anvil supercomputer while obtaining Master's degree

Vivek Karunai Kiri Ragavan, a Graduate Research Assistant at the Rosen Center for Advanced Computing (RCAC), recently obtained his Master of Science (MS) degree in Computer Information and Technology. His MS thesis was related to the work he conducted at RCAC, which focused on enhancing cloud-native clusters through custom scheduling.

For his Master's thesis, Ragavan focused on two significant projects, both related to developing custom scheduling capabilities within the Kubernetes Control Plane. His goal was to develop and implement a network-aware scheduler into Kubernetes that considers factors such as latency and bandwidth. Ragavan's solution turned out to be extremely successful. He found that his new network-aware scheduler reduced average latency by nearly 53% and minimized maximum latency spikes by 85% compared to the default Kubernetes scheduler. In fact, Ragavan's work was substantial enough to be accepted for the 2025 IEEE International Conference on Communications in Montreal, Canada, where he presented alongside his academic advisor, Dr. Deepak Nadig, on June 9.

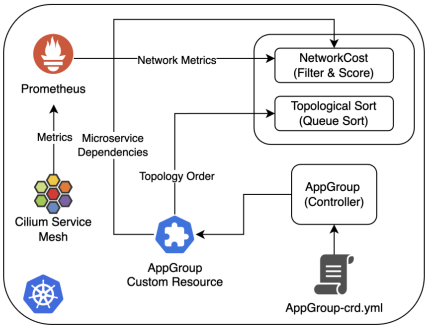

Figure 3.1. NAS Solution Architecture

The second part of Ragavan's MS thesis focused on developing GPU (graphics processing unit) autoscaling mechanisms based on real-time demand. This project was targeted at improving the performance and efficiency of AnvilGPT, RCAC's large language model (LLM) service that makes open-source LLM models like LLaMA accessible worldwide to ACCESS researchers. In the end, Ragavan's customized autoscaler achieved a 70% reduction in TTFT for AnvilGPT, sustaining low latency and high throughput while also dynamically adjusting resources to improve resource allocation.

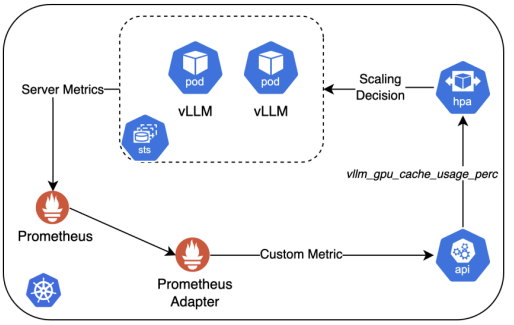

Figure 3.2. GPU Pod Autoscaling Solution Architecture

"So in two years," says Ragavan, "I am really indebted to everything RCAC has done for me. Working there was a really great experience. And really big thanks to Sarah [Rodenbeck], Erik [Gough], and Laura [Theademan], who provided me with these opportunities and guided me on my journey. Because of the experience and the work I was able to accomplish, I have ended up in a really good place right now."

CI-XP Student Presentations at the Envision Center

Earlier this year, RCAC hosted its annual CI-XP Student Program Lightning Talks, giving RCAC students a chance to showcase their projects and accomplishments to peers and staff within the organization. During the Lightning Talks event, 36 students gave two- to four-minute presentations highlighting their projects and what they've learned throughout the school year. Each student gave a brief overview of themselves: their background, course of study, likes and dislikes, and plans after graduation. They then moved on to their accomplishments and takeaways from working at RCAC.

The range of projects the CI-XP students worked on was impressive, with very few seeing any overlap. Some students helped to develop and improve different aspects of the Envision Center's virtual learning programs, including the newly created Collab XR platform. Others worked alongside our Research Software Engineers and high-performance computing (HPC) experts on a variety of hardware- and software-based HPC projects. A few of the students' projects even related to Anvil, our nationally-allocated HPC resource and one of Purdue's most powerful supercomputers. But a unique characteristic that all the CI-XP projects have in common is that they are not just for show, as is often the case with "practice projects" given to students. The work that the CI-XP students conduct is implemented into RCAC systems and utilized by our users. No matter which domain the students work in—VR, HPC, or even communications—their efforts help to advance the organization as a whole, an aspect that truly offers them an authentic workplace experience.

> It's so great to see all of the different projects from students across the CI-XP program. The annual Lightning Talks give everyone an opportunity to learn about the amazing work these students do and helps to reaffirm why this program is special. It is definitely one of my favorite events we host each year.

"It's so great to see all of the different projects from students across the CI-XP program," says Laura Theademan, the Director of Center Operations and Visualization at RCAC. "The annual Lightning Talks give everyone an opportunity to learn about the amazing work these students do and helps to reaffirm why this program is special. It is definitely one of my favorite events we host each year."

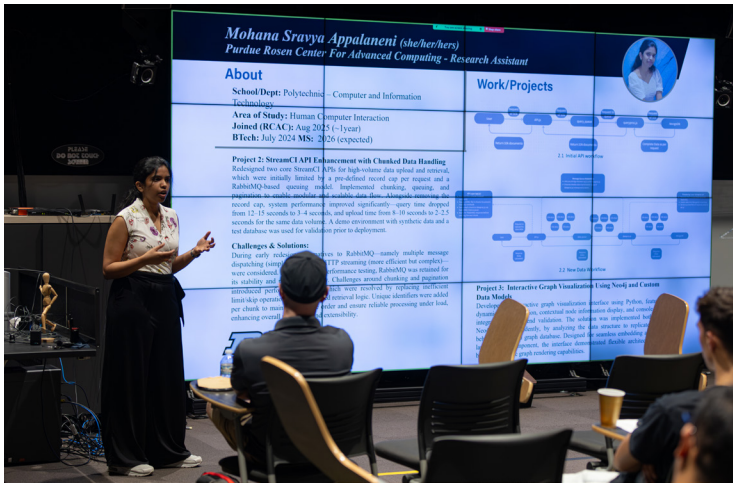

Anvil REU Student presents at GOOD 2025

Richie Tan, one of the participants of the Rosen Center for Advanced Computing's (RCAC) 2024 Anvil Research Experience for Undergraduates (REU) Summer program, recently presented at the Global Open OnDemand 2025 Conference (GOOD 2025). Tan highlighted the web dashboard that he and another REU student, Anjali Rajesh, developed during the summer program. The web dashboard went live on the Anvil supercomputer in January of 2025 and has been an enormous hit with Anvil users since going into production.

The Anvil REU Summer Program ended in August of 2024, but due to the success of the project, Tan was hired to continue improving the dashboard. He worked throughout the Fall semester fixing bugs for version 1 of the dashboard and making optimizations in the backend. This version went into production on Anvil in January, and included features such as:

- **Homepage widgets** showing service units, disk usage, queued jobs, etc.
- **My Jobs** page for a comprehensive view of recent jobs on Anvil.
- **Performance Metrics** page for job performance summary over specific periods of time.
- **In-memory caching** for API requests.

Tan immediately began working on version 2, which will go live on Anvil later this year. Version 2 will see enhancements such as:

- Redesigned homepage for cleaner look with more color-coding.
- Cluster Status page for viewing node statuses.
- Job Overview page for information about a specific job.
- Node Overview page for information about a specific node.
- Custom Time Frame selection option in the Performance Metrics feature.

Industrial Engineering students present their RCAC Capstone Project

Six students from the Edwardson School of Industrial Engineering recently completed their senior capstone project, which focused on helping to predict and reduce downtimes for RCAC's computing systems. The group worked toward creating a predictive-maintenance Artificial Intelligence (AI) tool for monitoring the high-performance computing (HPC) infrastructure, and presented their work at an Industrial Engineering student poster session.

The six students who took on the project were Zechen Wei, Hongchen Liu, Zachary Ramirez, Nicolai Cronin, Carlos Cordova, and Justin Ha. The team worked under the supervision of RCAC staff members Kyle Purple, Ashish, and Samuel Weekly. Throughout the semester, the team analyzed historical failure data on the Anvil system in order to identify problem areas and pain points. The group found that the most common problem on Anvil stemmed from a lack of communication between the nodes and Zabbix, the open-sourced software used to monitor and track the performance of the computer. They then utilized four different machine learning approaches to create their predictive model. Once completed, the group assessed the accuracy of their model, determined the model limitations, and even created a Grafana visualization showcasing the model data.

The six students ended the semester by showcasing their work at the Industrial Engineering poster event. RCAC staff members stopped by to view their poster presentation and were thrilled with how well the group presented. Overall, the team's efforts were a resounding success, resulting in a tool that RCAC can use and build upon in the future.

RCAC successfully hosts third Anvil REU Summer program

Over the summer, the Rosen Center for Advanced Computing (RCAC) hosted its third annual 11-week hands-on internship, the Anvil Research Experience for Undergraduates (REU) Summer 2024 program.

Eight students from across the nation gathered at Purdue's campus in West Lafayette, Indiana, for this year's Anvil REU program. The students took on this internship program to learn about high-performance computing (HPC) and work on projects related to the operations of the NSF-funded Anvil supercomputer at Purdue. During the program, which is hosted with support from the National Science Foundation (NSF), the students were able to gain the knowledge and skills necessary to build and support advanced research computing systems and scientific applications on these systems. The eight participants of the Anvil REU program were:

- **Jeffrey Winters**, Computer Science and Engineering double major, University of California, Merced; Project #
- **Alex Sieni**, Computer Science and Statistics double major, University of North Carolina at Chapel Hill; Project
- **Richie Tan**, Computer Science major, Purdue University; Project #2
- **Anjali Rajesh**, Computer Science major, Rutgers University; Project #2
- **Nihar Kodkani**, Computer Science and Math double major, Purdue University; Project #3
- **Selina Lin**, Computer Science and Math double major, Purdue University; Project #3
- **Philip Wisniewski**, Computer Science major, Purdue University; Project #4
- **Austin Lovell**, Computer Science major, Purdue University; Project #4

The Anvil REU program had the students join in pairs to complete four separate projects. These projects were chosen with real-world applicability in mind—the students would not only be gaining experience with HPC and learning new skill sets, but would simultaneously be increasing Anvil's capabilities. Each project also had two mentors working with the students to help them achieve their goals.

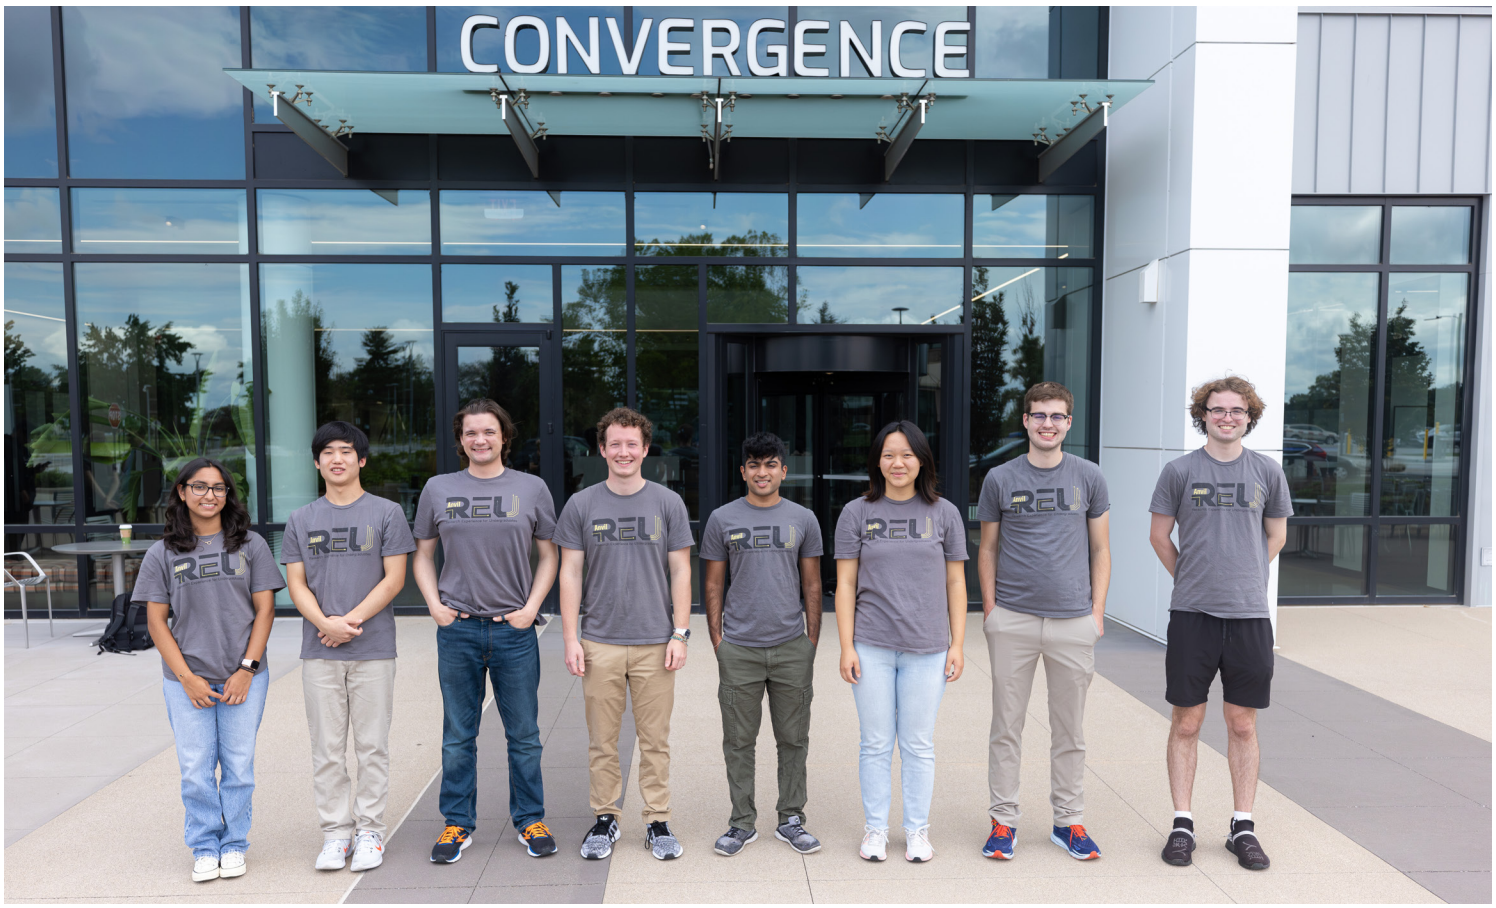

STUDENT HIGHLIGHTS

Nikitha Sivasankaravel, Junior, Psychological Sciences
Position at RCAC: Student Communications Liaison/Student Writer

At the RCAC, I'm currently the Student Writer! As a student writer, my job usually entails helping write science highlights and various articles. Often, I'm reaching out and interviewing various professors and students about their work in HPC/related fields.

Augusto Butkewitsch, Sophomore, Computer Science, Mathematics and Statistics
Position in RCAC: Programmer

As part of my work with RCAC and Envision Center, I am helping develop a 360° interactive simulator of Purdue's nuclear reactor to train students for successful operation. As part of a series of virtual labs, the simulation walks students through the proper steps to achieve criticality in the reactor, analyze real-time data feeds, and understand safety & shut-down procedures.

Koki Motoi, Grad Student, Learning Design and Technology
Position at RCAC: Instructional Designer

The mission of the RCAC Instructional Design project is to integrate instructional design principles into RCAC's educational practices such as in-person workshops or online trainings. For this purpose, I evaluate and propose amendments to workshop trainers so they can make workshops more engaging and interactive. I'm also working on developing a new online learning module on BrightSpace to help beginner users of HPC.

Ankitha Mallekav, Senior, Computer Engineering
Position at RCAC: HPC Research Assistant

Natural Language-Powered Python Code Editor Project description: I developed a Python code editor where users can write code using natural language, and the system automatically translates their descriptions into fully functional Python scripts. The project leverages NLP models to interpret user input, generate executable code, and run it within the editor.

David Piedra, Junior, Cybersecurity
Position in RCAC: Research Support Developer

My favorite project was Automated Unified Reconnaissance, Offensive Research, and Analysis (AURORA) Project description: This project of mine utilizes self-hosted reasoning large language models, which have been quantized for high performance at low cost- in order to automate tasks for cybersecurity such as recon and analysis.

Conferences and Presentations

Purdue makes a splash at SC24

Purdue University made an impressive showing at this year's supercomputing conference, SC24. From offering innovative presentations and workforce development opportunities to receiving major awards, Purdue's presence at SC24 highlighted why the university is a leader in advanced research, technology, and high-performance computing (HPC).

The central theme for the Purdue booth this year was to promote the Purdue Computes initiative, specifically focusing on Purdue's Artificial Intelligence (AI) capabilities. To help achieve this goal, Purdue provided the conference with presentations throughout the week, not only from RCAC staff but also from experts within other departments. Dr. Marisol Koslowski, Associate Dean For Global Engineering Programs And Partnerships, and Professor Of Mechanical Engineering, gave an excellent presentation on large-scale simulations of small-scale material deformation mechanisms, while Dr. Alejandro Strachan, Reilly Professor Of Materials Engineering, discussed how researchers can get the most out of large-scale molecular dynamics simulations with AI tools as well as the newly created online platform for everything semiconductors, Chipshub. Dr. Corey Maley, an Associate Professor Of Philosophy, promoted the College of Liberal Arts's new BA in Artificial Intelligence major and showcased his research on the foundational issues in the philosophy of computation. Rounding out the group was Dr. Aniket Bera, Associate Professor of Computer Science and Director of IDEAS Lab, who gave a presentation titled "From Data to Dynamics: AI in Human Motion Prediction and Analysis." During his talk, he discussed his collaboration with RCAC's Envision Center to obtain motion capture data for advancing AI and robotics, as well as his other AI-centered research projects. RCAC staff members also delivered multiple booth presentations and demonstrations throughout the week, covering a variety of AI and HPC projects, including HyperShell, AnvilGPT, and Purdue GenAI Studio.

Talks and presentations were not the only things Purdue had to offer at SC24. Thanks to the efforts of hardworking—and mostly volunteer—staff members, Purdue supported multiple committees, organizations, and events at the conference, including SCinet, Women in HPC, PEARC, the Early Career Program, the Student Cluster Competition, and the SC24 Job Fair.

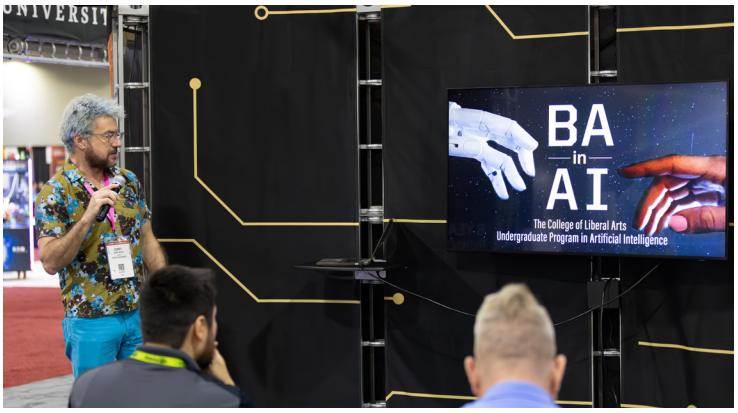

RCAC staff presents at PEARC24

RCAC staff presented papers, led workshops and participated in panels at the 2024 Practice and Experience in Advanced Research Computing (PEARC) conference, themed "HPC: Human Powered Computing." The conference was held in Providence, R.I. from July 21-25, 2024.

Erik Gough, senior research scientist, presented his paper "Evaluation of Kubernetes Schedulers for a Community Cloud Computing Model" as part of the Systems and System Software track.

Guangzhen Jin, Eric Adams, Ruyi Li, Haniye Kashgarani, and Nannan Shan presented their paper "Behind the Scenes: How We Support User Support at an HPC Center" as part of the Abstract Submissions track.

Kashgarani presented her paper "Combinatorial Optimization by Automated Selection of Parallel Portfolios" as part of the Abstract Submissions track.

I Luk Kim, Lan Zhao, and Carol Song presented their paper "Japper: A Comprehensive Framework for Streamlining Jupyter-Based Scientific Web Application Development" as part of the Applications and Software track.

Jaewoo Shin, Zhao, and Song presented their paper "Integrating ML/AI Workflows in a Streaming Data Management and Processing Platform for Building Energy Research" as part of the Applications and Software track.

RCAC staff members also led the following workshops at PEARC24:

Gough and Rajesh Kalyanam, along with collaborators from the University of Illinois Urbana-Champaign, Indiana University, and the University of California San Diego, co-lead the BoF "Best Practices for Kubernetes Security Policy on Composable Clouds."

Carol Song co-led the BoF "Fantastic ACCESS Cyberinfrastructure Resources and Where to Find Them."

Laura Theademan and Suzanna Gardner hosted a panel "The Vital and Diverse Roles and Expertise in Cyberinfrastructure," along with collaborators from other universities.

Many RCAC staff members were also involved in conference leadership. **Eric Adams** was the technical program co-chair, **Geoffrey Lentner** was the proceedings chair, **Theademan** was the exhibitors co-chair, **Nannan Shan** was the co-chair of the panels committee, and **Gardner and Amanda Warren-Glowe** were on the student program committee. **Katy Gunderson** was the graphic designer for the conference.

Minority Serving – Cyberinfrastructure Consortium Annual Meeting

The Rosen Center for Advanced Computing (RCAC) sponsored the 2025 MS-CC Annual Meeting, underscoring the organization's commitment to fostering inclusive research computing partnerships. In addition to providing sponsorship support, Suzanna Gardner, Sr. Research Ops for Outreach and Engagement at RCAC, served on the GNC Committee with MS-CC. During the meeting, Nannan Shan, Senior Computational Scientist, presented on Anvil and the National AI Research Resource (NAIRR), highlighting how these resources can support cutting-edge research and broader access to high-performance computing.

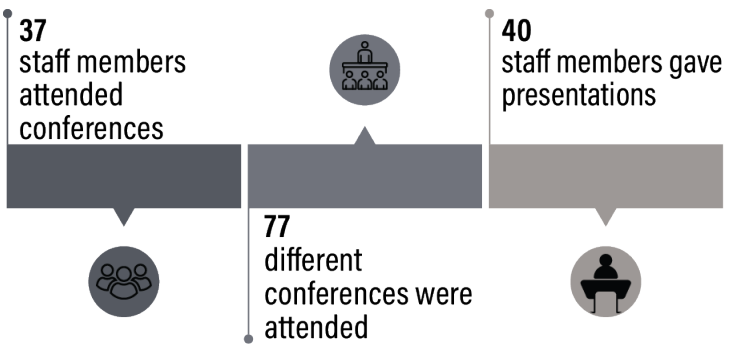

Education and Outreach

Third annual Cyberinfrastructure Symposium hosted by RCAC

In the Fall of 2024, the Rosen Center for Advanced Computing (RCAC) successfully hosted its third annual cyberinfrastructure (CI) symposium. This year's symposium focused on the role that high-performance computing (HPC), artificial intelligence (AI), and semiconductors play in scientific advancements both at Purdue and abroad.

The day began with an opening presentation from Preston Smith, executive director of RCAC, who highlighted the state of Purdue's campus cyberinfrastructure. This was followed by a presentation by Rizki Oktavian, an AI engineer and reactor physicist for Blue Wave AI Labs, one of the Anvil supercomputer Industry Partners. Next on the list of presenters was David Ryglicki, a Tropical Weather Expert/Senior AI/ML Engineer for MyRadar, another of Anvil's Industry Partners. Karen Plaut, Executive Vice President for Research at Purdue, followed Ryglick's presentation by highlighting the variety of ways in which HPC and AI are driving research and innovation across campus.

Michela Taufer, an AAAS Fellow and ACM distinguished scientist, delivered the keynote address. She spoke about the significant challenges scientists face across various domains, institutions, and generations in managing and utilizing large-scale data from experimental facilities, remote sensing, and simulations at major national laboratories.

After Taufer delivered her address, the CI Symposium broke for lunch, during which a student poster session was held in the Envision Center. The winner of "Best Poster Award" was Haoteng Yin, a graduate student in the Department of Computer Science. His poster was titled "Privately Learning from Graphs with Applications in Fine-tuning Large Pretrained Models."

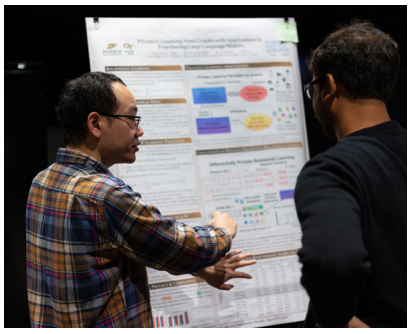

Following the poster session, the CI Symposium continued with six more presentations:

- **Matteo Ruggeri**, Aeronautical and Astronautical Engineering, Purdue University.
- **Erica Carlson**, Professor of Physics and Astronomy at Purdue University.
- **Jeff Terstriep**, Senior Program Manager, The National Center for Supercomputing Applications.
- **Eamon Duede**, Assistant Professor of Philosophy at Purdue University, joint appointment Argonne National Laboratory.
- **Nathaniel Husted**, Chief Scientist for Cyber and Electromagnetic Warfare (EW) technologies at Naval Surface Warfare Center-Crane Division.
- **Carol Song**, Chief Scientist, Rosen Center for Advanced Computing at Purdue University.

RCAC hosts successful 2025 Research Computing CI Symposium in Indianapolis

In the Spring of 2025, the Rosen Center for Advanced Computing (RCAC) hosted the 2025 Research Computing Cyberinfrastructure (CI) Symposium. This event marked the fourth CI Symposium put on by RCAC, and the first to take advantage of the beautiful space at High Alpha in Indianapolis. The theme for the day was "The Greatest Spectacle in Simulation." Presentations were delivered by researchers, staff, and faculty from both the West Lafayette and Purdue Indianapolis campuses and covered a range of topics related to high-performance computing (HPC).

The 2025 CI Symposium began with opening remarks from Karen Plaut, executive vice president for research at Purdue University, and Dave Umulis, senior vice provost and chief academic officer for Purdue University in Indianapolis. After the opening remarks, Preston Smith, executive director of RCAC and director of computing infrastructure for Purdue's Institute for Physical Artificial Intelligence, gave the first presentation. Smith discussed the current state of research computing at Purdue.

The keynote presentation for the day was delivered by Dr. Robbee Wedow. Wedow is an assistant professor of Sociology and Data Science at Purdue, and an adjunct assistant professor of Medical and Molecular Genetics at the Indiana University School of Medicine. His presentation, titled "Genes and Algorithms: The Rise of Computational and Statistical Genomics," discussed his lab's research into genetic and environmental determinants of human health and behavior and the statistical methods and data resources needed to analyze the human genome. Over the years, Wedow's lab has been increasingly reliant on HPC resources to help conduct research into the human genome.

Other speakers from the event were:

- Dr. Huidan (Whitney) Yu, Professor of Mechanical Engineering, Purdue University in Indianapolis, and Adjunct Research Professor of Vascular Surgery, Department of Vascular Surgery, Indiana University School of Medicine
- Dr. Shreyas Sundaram, Marie Gordon Professor, Elmore Family School of Electrical and Computer Engineering Co-Director, Institute for Control, Optimization and Networks, Purdue University
- Dr. Kenichi Yoshida, Professor of Biomedical Engineering, Purdue University
- Andre Gardinalli, HPC and AI Sales Executive, Lenovo

To cap off the day, a community reception was held for attendees and speakers on the High Alpha outdoor Sky Deck, allowing people to network and further discuss the topics of the symposium.

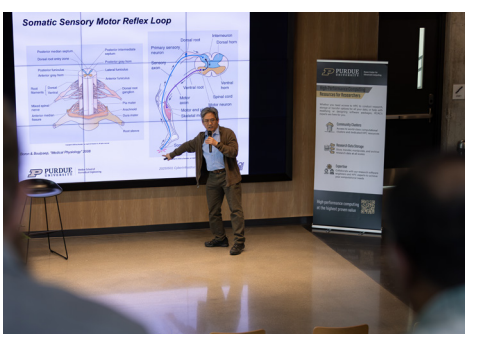

Second annual RCAC summer camps are a huge success

The Rosen Center for Advanced Computing (RCAC) recently hosted two summer camps for high schoolers with the goal of introducing and developing cybersecurity and coding skills. With these camps, the students not only learned new skills, but also had the opportunity to earn college credits while experiencing campus life.

Both camps took place at Purdue's West Lafayette campus from June 22nd through the 27th and were part of the Summer College for High School Students program at Purdue. RCAC hosted each of these camps last year as well, and they were well-received by both students and organizers alike. The demand for the programs this year was so high that each camp filled to max capacity within a week of being announced.

Code Explorers

The first camp was "Code Explorers: Coding, Data, and Environmental Stories." This camp was an immersive STEM program where the students engaged with coding, data collection, and environmental science. The 2025 coding class focused on analyzing and visualizing data related to Earth and Planetary Science. The students learned to code in Python for Data Science using Purdue RCAC's JupyterHub and the Department of Computer Science labs. They gained hands-on experience in data collection, preparation, analysis, and visualization using datasets from both local and national Earth and planetary science resources.

CyberSafe Heroes

The second camp was "CyberSafe Heroes: Empowering High School Students for a Secure Digital Future." This camp focused on cybersecurity, with the goal of introducing students to key concepts in ethical hacking, encryption, and digital safety. Through hands-on labs and engaging activities, including escape rooms, hacking simulations, and meme contests,the students explored fundamental cybersecurity topics such as the CIA Triad, vulnerability management, and network security. Participants used industry-standard tools, including Kali Linux and Metasploit, and engaged with cybersecurity professionals in career panels, forging tangible connections to real-world roles in digital defense. As with the Code Explorers camp, the CyberSafe Heroes camp was open to students with any background, no prior-experience required. George Bailey is the director of cyberTAP at Purdue, and was one of the instructors for the CyberSafe camp.

Making an Impact

Overall the 2025 RCAC summer camps were a great success, receiving overwhelmingly positive feedback from both the students and the instructors. Suzanna Gardner, the Senior Research Operations Administrator for Outreach and Engagement at RCAC, serves as the primary organizer and director of the RCAC summer camps. She is thrilled with what the camps have achieved and the impact they have made in the lives of the students.

"It's incredible to watch students light up when they realize that computing can help solve real problems—whether it's protecting digital spaces or understanding our environment," says Gardner. "With Code Explorers and Cyber Safe Heroes, we wanted to create experiences that felt meaningful and fun, where students could see themselves in these fields and start imagining what their future in tech might look like. I'm happy to say that over the past two years, we've done just that."
and was one of the instructors for the CyberSafe camp.

Staff News

- **Arun Seetharam**, Lead Bioinformatics Scientist, July 1, 2024
- **Primus Chimdia Kabuo,** Senior Research Software Engineer, July 22, 2024
- **Elham Jebalbarezi Sarbijan,** Research Scientist, August 12, 2024
- **Ibrahem Alshybani,** Senior Computational Scientist, August 19, 2024
- **Xiao Liu**, Senior Computational Scientist, August 26, 2024
- **Mihir Ahlawat,** Senior AI Scientist, September 30, 2024
- **Jacob Roberts**, Lead Research Solutions Engineer, November 4, 2024
- **Joseph Levell**, Senior Manager, Scientific Applications, February 10, 2025
- **Bithi De**, Senior Computational Scientist, February 17, 2025
- **Rose Wilfong,** Senior Research Data Facilitator, March 10, 2025
- **Joshua Remender,** HPC Support Specialist, March 24, 2025
- **Vineeth Gundeti,** Senior Research Software Engineer, April 22, 2025
- **Jacob Verburgt,** Senior AI Applications Scientist, May 19, 2025
- **Douglas Schultz**, Senior Manager, Research Computing Support Services, June 16, 2025

RCAC computational scientists delivered 25 in-classroom or online training events in AY 24-25, reaching 443 students - a total of 1,906 contact hours. The AI topic specifically saw 8 training events for 943 contact hours.

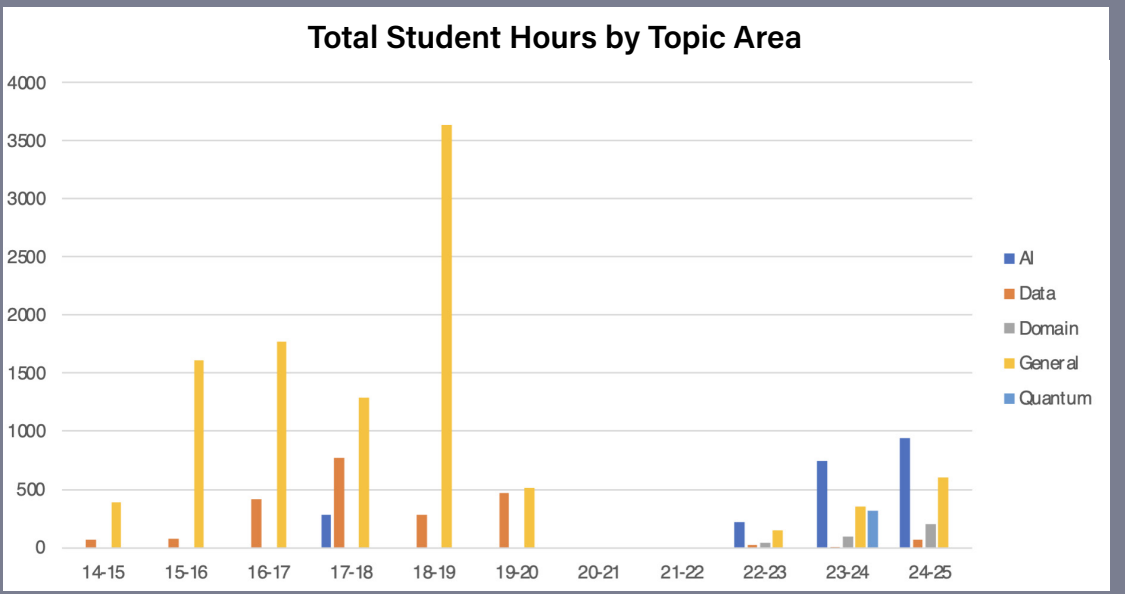

WHPC Events

Purdue co-hosts "Building and Maintaining Supportive Communities for Women in HPC" event

Purdue's Women in High-Performance Computing (WHPC) group, which is part of a broader engagement initiative by the Rosen Center for Advanced Computing (RCAC), co-hosted a workshop on Sept. 3 organized by the Virginia WHPC chapter about current challenges and opportunities for fostering a more diverse and inclusive HPC community.

The event began with introductions to the Virginia, Purdue and Northeast WHPC chapters, which jointly co-hosted the event. Laura Theademan, director of center operations and visualization for RCAC, introduced the Purdue chapter.

After that, the group split into breakout rooms to discuss specific issues facing women in HPC including:

- The role of mentors and engaging allies
- Topics/types of events to be covering within chapters
- Gaps in programming efforts or outreach activities and how they can be filled going forward
- What participants would like to see from other WHPC chapters
- Current challenges in gender diversity and inclusivity in HPC

"The meeting was informative and has strengthened my enthusiasm to continue participating in Women in HPC. In fact, after speaking with one of the representatives, we are working on organizing an event to introduce ECE graduate students to HPC."

Purdue WHPC is led by women staffers affiliated with RCAC. WHPC is a diverse community encompassing undergraduate, graduate, staff and faculty men and women who are interested in exposing women to high-performance computing and encouraging their pursuit of research and careers in HPC and other technology fields.

Purdue WHPC organizes a range of activities including a scholarship program that supports travel for women students to industry conferences, regular meetings to discuss technical HPC-related issues of interest, opportunities to network with the WHPC community, a mentorship program, workshops, and exposure to external resources and opportunities. All WHPC activities are generously supported via private donations.

Success Stories & Testimonials

Sruthi Dasika

"Attending the Purdue WHPC workshop was an enlightening experience that broadened my perspective on the nuances of gender in both professional and personal spheres."

George Lyle

"I came to the workshop because I'd previously read the book 'Invisible Women,' so I thought it would be interesting to hear other people's interpretations of it. Each of the presenters sparked a new and instructive conversation about the challenges of a world built around one single default identity, and the struggles of people who don't have access to that default."

Durga Madarapu

"For a long time at Purdue, I could not find an HPC community. That changed when I discovered RCAC and the Women in HPC group. I felt more connected as a researcher and as a Purdue student. I am deeply grateful for the support and scholarship from Purdue WHPC."

Travel Scholarships

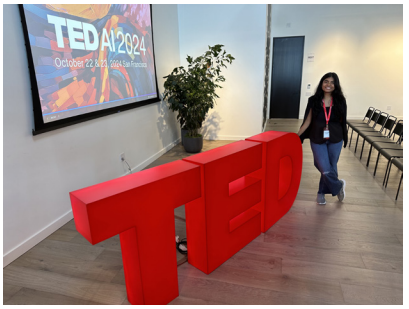

Tanvi Kadari

"It was such an enriching experience to get an opportunity to learn from industry leaders and academics about the latest advancements in AI in fields from marine biology to theoretical physics and gain perspective on where the industry is heading,"

Megha Bisht

"The energy, creativity and innovation in every session and booth were contagious." Being surrounded by passionate individuals who are redefining what's possible in the gaming world was truly motivating."

Elisa Chen

"Grace Hopper was an incredible experience, filled with uplifting moments and valuable professional development opportunities

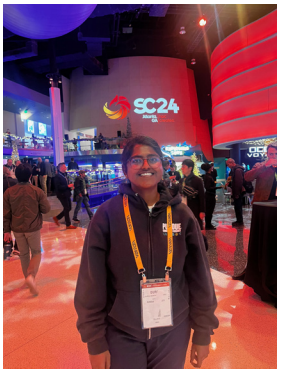

Durga Mandarapu

"The scale of HPC research and its transformative impact on real-world applications was far beyond anything I had previously imagined,"

Ankitha Mallekav

"It was so fun and motivating to see so many other women in tech with similar goals. Seeing their achievements and passion in the field encouraged me even more to pursue my own path in technology."

Ishwarya Samavedhi

"My biggest takeaway from GHC is that I realized my values and ideas matter just as much as anyone else's. These incredible women reminded me that all bring something to the table."

Jesal Movani

Every session, every conversation felt like it was expanding my world, offering me new perspectives and fresh ideas. I left with a renewed commitment to what's possible in my own career."

Sponsorship

Aside from providing access to leading-edge computational and data storage systems as well as expertise and support, the Rosen Center for Advanced Computing also engages in sponsorship opportunities for both individuals and events. RCAC is dedicated to advancing the learning and career outcomes for students, promoting the use of HPC, and developing the next generation of HPC professionals. By taking part in strategic sponsorship opportunities, RCAC further bolsters its efforts in achieving these goals.

InnovateHer is a 36-hour hackathon at Purdue University created specifically for and by students of underrepresented identities in technology.

"I am thrilled that we [RCAC] were able to sponsor InnovateHer for a second year in a row," says Graham. "It is an amazing event that brings students from multiple disciplines together to work on often neglected areas in computing and technology. Supporting students to develop skills in computing and HPC is something that we as an organization strongly believe in, and the more involved we can be, the broader our impact. Personally, I had a blast at the hackathon, and can't wait to be back at next year's event."

BoilerMake XII. BoilerMake is Purdue University's premier hackathon, inviting undergraduate students from across the nation to join together for a weekend of computational and technological development.

"I love the creative energy of events like this," says Geoffrey Lentner, a Lead Research Data Scientist for RCAC, and the main driver behind RCAC's involvement in BoilerMake. Lentner had volunteered as an individual for the past three BoilerMake events and decided it was time to get the whole organization involved. "It's always a lot of fun. I'm glad we [RCAC] were able to sponsor the hackathon. Offering some of Anvil's GPUs to the students was a great idea, and was really helpful with both fine-tuning models as well as hosting inferencing prototypes."

The Catapult Hackathon, is an event put on by the ML@Purdue student club.

RCAC's sponsorship included an allocation of GPU hours on the Gaustschi community cluster, with the winning team getting to keep the balance of hours that weren't used during the hackathon.

"RCAC's support has been integral to the success of Catapult," says Setpal. "Teams spent 24 hours developing innovative, AI-driven solutions to various challenges and the 8x H100 GPU nodes available to each of them through Gautschi meant teams were virtually unrestricted by computing capabilities. The creativity it enabled was palpable throughout the hackathon, and we look forward to collaborating further with RCAC on future events that ML@Purdue organizes."

With support from the Rosen Center for Advanced Computing (RCAC), Purdue doctoral candidate Haoteng Yin recently attended NeurIPS 2024, one of the world's leading conferences in artificial intelligence and machine learning.

"My experience at the RCAC CI Symposium was fantastic," says Yin. "I had insightful conversations about high-performance computing and AI for science, and I gained a deeper understanding of the computational tools and resources available through RCAC. The event gave me new ideas on how to better integrate these resources into my own research on designing efficient and scalable learning frameworks for massive graph data."

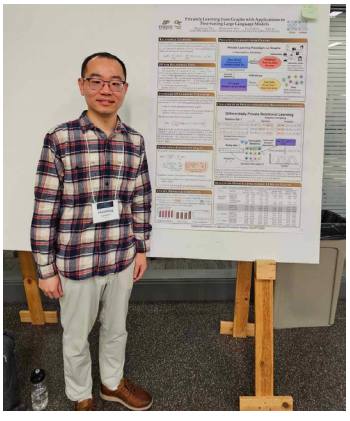

"The RCAC support was instrumental in making my participation at NeurIPS possible," Yin notes. "I am deeply grateful for this opportunity, and it has reinforced my commitment to contributing to the community and making my research results accessible."

PURDUE UNIVERSITY

Rosen Center for Advanced Computing

1 Billion
CPU hours provided

2.59 Million
GPU hours provided

299
Community cluster PIs

64%
% of Purdue research expenditures to RCAC users

4000+
Community cluster users

81,178
Community cluster OOD sessions

27%
% of doctoral graduates using community clusters

169
RCAC Outreach events

1045
GenAI Users

81
Students supported by CI-XP internships

1,906
RCAC Student hours

Sponsored research expenditures within RCAC **9.11 Million**

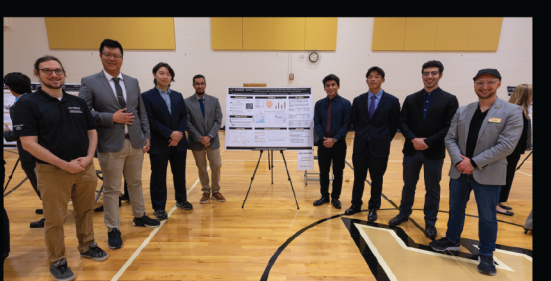

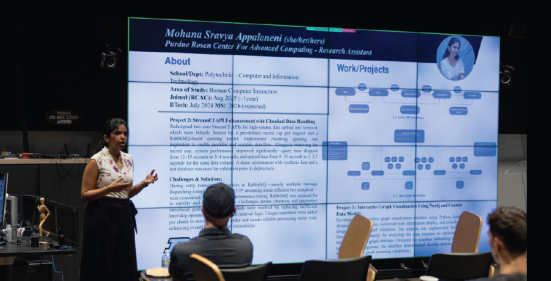

PURDUE UNIVERSITY | Rosen Center for Advanced Computing

High-performance computing at
the highest proven value.