

BIGHPC TEXAS/PORTUGAL PARTNERSHIP FORGED TO TRANSFORM DATA SCIENCE

TACC, computer science experts to create new R&D for academia and the commercial sector

By 2025, the [global high performance computing \(HPC\) market](#) is forecasted to account for \$50 billion, up \$15 billion from 2019. Fueled by the growth of big data, artificial intelligence, and 3-D imaging, HPC is increasingly used to drive commercial, industrial, and personal usage.

Scientists and users in industry alike are increasingly applying HPC infrastructures such as compute clusters, networking, and storage to support big data applications. These applications utilize larger, more complex data sets that can overwhelm traditional data processing resources. As the proportion of supercomputer cycles devoted to big data fields such as machine learning, fMRI, and experimental particle physics analyses grows, the workloads of HPC and big data are converging and require hardware and software resources capable of supporting each on equal footing.

To address this challenge in today's HPC data centers, a new partnership has been forged among the [Portuguese Foundation of Science and Technology](#) and The University of Texas at Austin (UT Austin), specifically the [UT Austin Portugal program](#), the [Texas Advanced Computing Center \(TACC\)](#), and the [Department of Computer Science](#). [Wavecom](#), a private wireless company in Portugal, is leading the effort and plans to market the outcomes to the computing community in Portugal and beyond.

The project — BigHPC: A Management Framework for Consolidated Big Data and HPC — started in April 2020 and includes \$2 million in funding over the next three years.

"It's an honor that we were selected to work with Portugal on this important project," said Vijay Chidambaram, one of two principal investigators on the grant and a professor in the Department of Computer Science at UT Austin.

The UT Austin Portugal program dates back to 2007, and is one of several partnerships between the Portuguese government and research institutions. The program's goal is to elevate science and technology in Portugal while fostering strong partnerships to help universities continue to innovate. The partnership with UT Austin was extended in 2018, continuing the alliance until at least 2030.

"The idea of the BigHPC grant is to focus on practical R&D that will lead to real applications for scientists whether they're in industry, academia, or a national lab, and transfer that knowledge into the commercial sector," said Todd Evans, who is a member of TACC's HPC Performance & Architectures Group and the second principal investigator on the grant. "The project will also contribute to building the much needed capacity to implement and manage HPC and big data services in academia and industry," Evans said.

The result will be a novel solution for monitoring and managing the infrastructure, data, and applications of current and next-generation HPC data centers, while ensuring the best performance and resource usage for applications and infrastructures. It will also enable both big data analytics applications and traditional HPC to be deployed using heterogeneous hardware and run in the same environment.

Experts from TACC and the [Minho Advanced Computing Centre \(MACC\)](#) in Portugal will help integrate the tools developed into a single software bundle that will be validated through real use-cases and a pilot deployed in both the TACC and MACC data centers.

The project will also leverage the expertise in HPC and big data management of researchers from the [Institute for Systems and Computer Engineering, Technology and Science](#), a Portuguese non-profit that acts as an interface between academia and industry; and the [Laboratory of Instrumentation and Experimental Particle Physics](#), a state-run Portuguese research laboratory founded in 1986.

When a scientist conducts computational research, good results depend on the efficiency of the underlying infrastructure. This was highlighted when researchers at UT Austin recently mapped the spread of COVID-19.

"The COVID-19 research translated into tasks that were run on TACC's supercomputers," Chidambaram said. "If we can make the systems in TACC's data center more efficient, we will directly speed up how long it takes to run these tasks. This helps researchers conducting science decrease time to solution."

MODERN-DAY HPC CHALLENGES

Academic and industry data centers are rising to meet new challenges in the HPC environment. Chidambaram says the two primary challenges are managing new, heterogeneous storage and virtualization technologies in data centers in an efficient manner; and the fact that traditional software needs to be modified to work with new storage technologies.

"HPC has been focused on very traditional techniques that have been around for more than 30 years," Chidambaram said. "These techniques are often not well-suited to new hardware."

For example, persistent memory is a new capability that provides very low latency access to data, but also has very unique characteristics. "If you used persistent memory like you would a magnetic hard drive, you're not going to get all the benefit you can from that device," Chidambaram added. "Similarly, in terms of virtualization technologies back when HPC really took off, there wasn't a notion of running containers. This is widespread in the field now."

Containers are a method of packaging an application with many software components along with complex dependencies to run high-end applications. They have altered the development of software today because they hide the complexity and make the deployment easier.

"Coping with heterogeneous hardware, large-scale infrastructure, and the different HPC application requirements is challenging," Evans said. "Even if an application exists and is vetted, it can be extremely labor intensive to modify it to run in a particular environment or on a particular system. Secondly, it's challenging to build an application such that it's taking advantage of the hardware capabilities that are available to it."

Chidambaram's lab — the [UT Systems and Storage Lab](#) — is working to build the next generation of storage technologies. To do this effectively, he needs partners.

"It's one thing to generally know what HPC requires. It's quite another thing to talk directly to the people who are managing the clusters and data centers," Chidambaram said.

"TACC is excited to be involved," Evans said. "It gives us an opportunity to think in a deeper way about virtualization, specifically container technologies that are becoming more prevalent."

"You can't practically build some of these packages from scratch anymore, particularly in relatively restrictive HPC environments, so you need containers to support applications," Evans continued. "This project is going a step further by sharing container-supported applications that are able to run in an optimal way for a given hardware configuration."

Many areas of the overall scientific community will benefit from the BigHPC grant — academic supercomputing centers, industrial data centers, national labs, and nonprofits.

Chidambaram concluded: "Typically, these proposals are only academic. In this instance, we have TACC and MACC to provide on-the-ground experience running supercomputers in an educational setting for research. We have Wavecom to talk with us about industrial production problems. And we have non-profits experienced in bridging academic and industrial efforts. We're all hoping to learn how best to manage data centers now and in the future, and how best we can use storage and virtualization technologies. Our collaborators will tell us whether our ideas make sense and work in the real world."

###