HYPERION RESEARCH

HPC User Forum Update

Reflections on Datacenter Sustainability Perspectives from TACC

Jaclyn Ludema and Mark Nossokoff
March 2024

IN THIS UPDATE

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. In September 2023, the 83rd HPC User Forum took place in Tucson, Arizona. This update summarizes a presentation from that conference given by Dan Stanzione, Executive Director of Texas Advanced Computing Center (TACC) and Associate Vice President for Research at University of Texas-Austin. Stanzione presented TACC perspectives on datacenter sustainability including infrastructure improvements and incentivizing users to improve software.

Source: TACC, 2023

PRESENTATION: TACC SUSTAINABILITY, DAN STANZIONE, EXECUTIVE DIRECTOR

Stanzione gave an overview of the TACC facility, which serves as host to various HPC systems and platforms, both for the National Science Foundation (NSF) and other organizations. The NSF systems run at TACC include:

- Frontera: Fastest supercomputer on a university campus in the US
- Stampede-2: Recently decommissioned, served as the flagship system for the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) and provided HPC capabilities to thousands of researchers across the US
- Jetstream2: First user-friendly, scalable cloud computing environment for ACCESS, giving researchers the ability to create customized virtual machines and computing architectures
- Chameleon: Large-scale testbed system for Computer Science experimentation, designed to be deeply reconfigurable and give users bare metal system control

Lonestar6 was also highlighted as the newest addition to the TACC resources. Lonestar is funded through the University of Texas Research Cyberinfrastructure and provides Texas Researchers access to cutting-edge CPU and GPU technology. The four NSF platforms and five non-NSF platforms at TACC aggregate together to approximately 20,000 servers on the main floor of the datacenter, about a million cores, and just under 1,000 GPUs. A typical day at TACC uses 6MW of power, with a maximum capacity for 9.5 MW. Plans are in place to add another 20 MW of datacenter capacity in the near future with the construction of the Leadership-Class Computing Facility (LCCF), for a total of 30 megawatts by 2025.

FIGURE 1

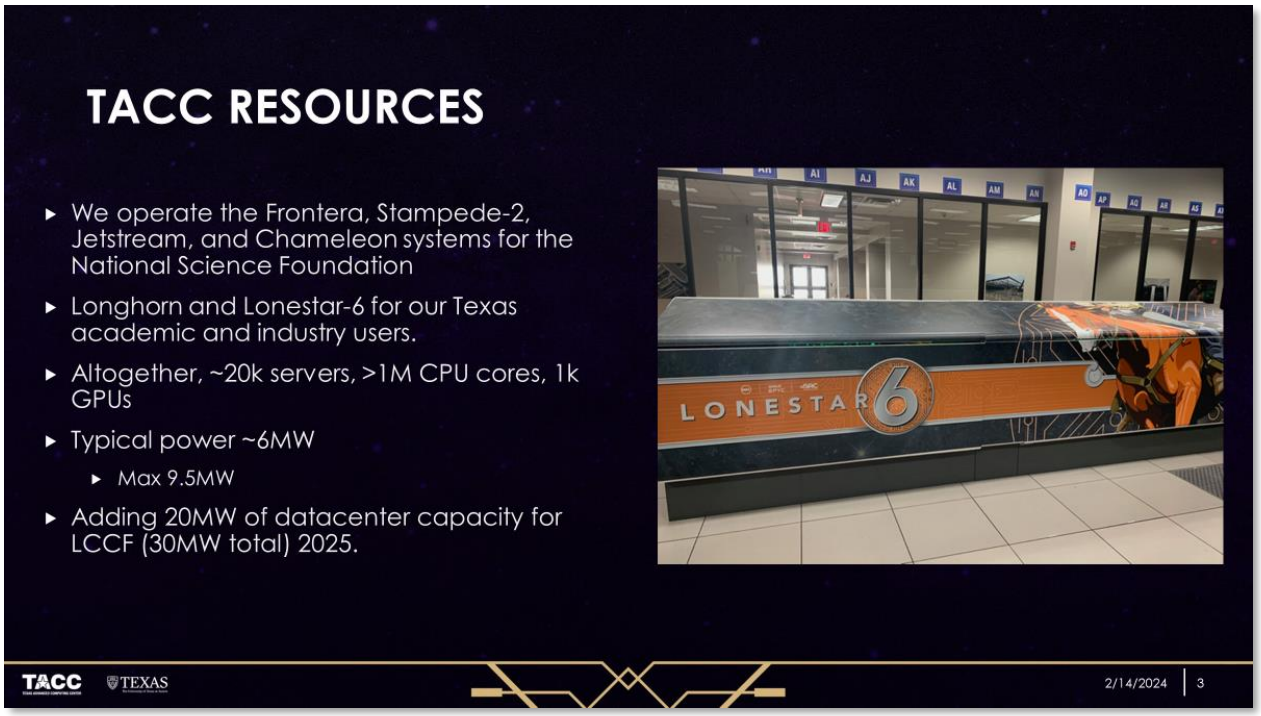

Source: TACC, 2023

Stanzione discussed the evolution of cooling strategies at TACC. In 2012, TACC employed in-row chillers and hot-out containment for systems like Stampede 1. Later, with Frontera in 2019, TACC transitioned to direct liquid cooling (DLC), using cold plates mounted to processors with liquid pumped through them. Stampede 3, slated to enter full production during the first half of 2024, will also use DLC. Stanzione mentioned that while many DLC options exist in the market, only rear-mounted radiators are necessary for these systems since they resemble standard COLO datacenter racks, and each use about 10kW. Full immersion cooling is the latest evolution for cooling at TACC, used for one subsystem of Frontera that has PCI-card GPUs and the new Lonestar6 system.

Cooling strategies at TACC have had to accommodate the increases in system density, with Stampede 1 using 30 kW per rack, Frontera using 60 kW per rack, and now Lonestar 6 using 70 kW per rack. Stanzione stated that the decision-making process for cooling strategies prioritizes system density and management over sustainability considerations, although each generation of cooling solutions aims to improve efficiency.

FIGURE 2

Source: TACC, 2023

Stanzione continued to discuss the sustainability strategies TACC employs within the broader datacenter infrastructure. Solar panels have been mounted to the covered staff parking lot, generating 200 kW of solar power. He highlighted the challenges posed by high temperatures in Texas and discussed strategies like using compressors for chilling.

TACC also purchases 20% of its power as wind credits from their Austin area provider. Stanzione mentioned that expanding this particular sustainability strategy at the TACC facility would be costly due to the area provider's monopoly on the power grid. However, plans for the additional computing capacity at LCCF will be located 10 miles north of TACC at the Switch facility in Round Rock. The Switch facility is strategically located where the power grid is unregulated, and for a 7% upcharge, the LCCF can receive 100% wind power.

FIGURE 3

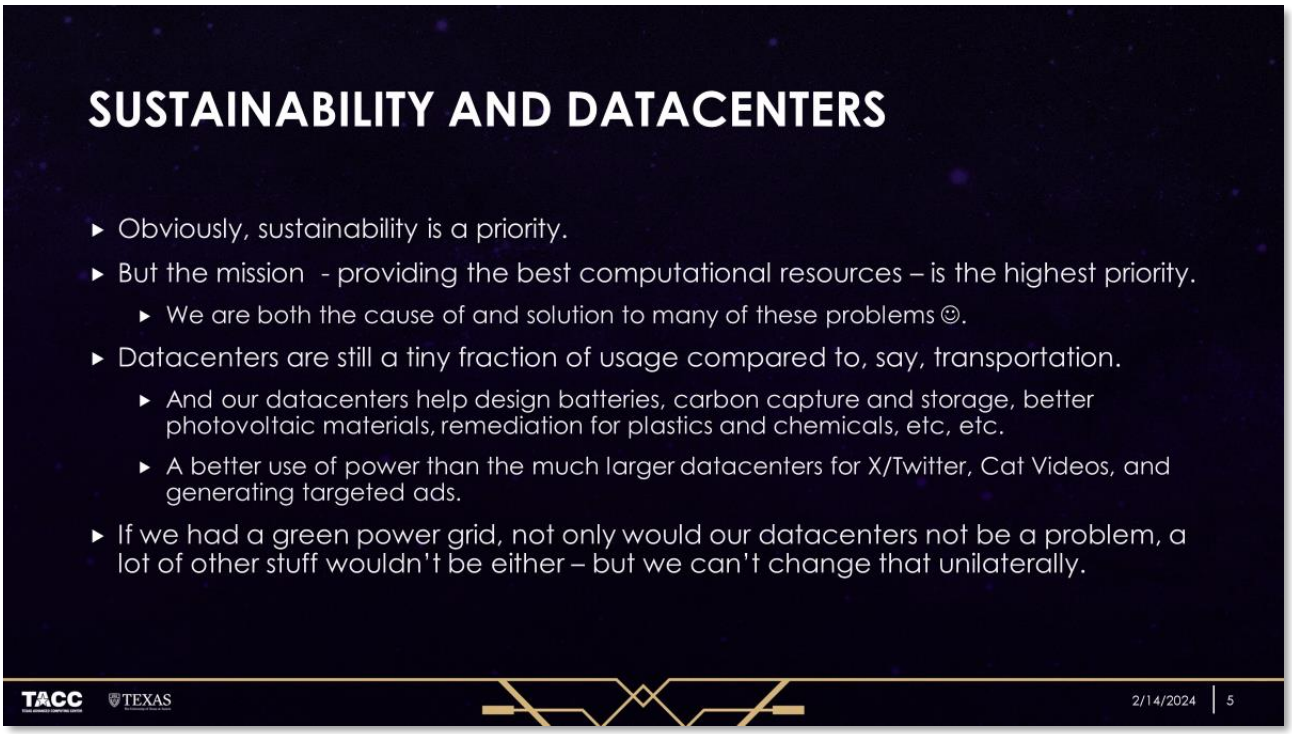

Source: TACC, 2023

Stanzione pointed out the dichotomy of HPC datacenter sustainability efforts, as they are both the cause and the solution to various factors of global climate change. With each new datacenter scaling and performance goal comes additional challenges of finding reliable and sustainably sourced power supplies. On the other hand, datacenters devoted to science and research such as TACC are used to solve some of the most complex sustainability challenges in a variety of industries. He cited projects performed at TACC, including the design of better batteries, carbon capture technology, new photovoltaic materials, and a recent project on the selective evolution of bacteria that will better break down plastics and forever chemicals.

He emphasized the importance of perspective when considering datacenter sustainability efforts. In comparison to industries such as transportation, datacenters' power consumption remains a tiny fraction of global power concerns. Stanzione thinks that the positive impacts TACC projects have on sustainability efforts have to be weighed against any power consumption concerns. He also pointed out that datacenter power consumption concerns would simply disappear if the power grid unilaterally

made the switch to renewable energy sources. In his opinion, more emphasis should be placed on improving the power grid writ large.

FIGURE 4

Source: TACC, 2023

Stanzione described several experiments being run at TACC looking to find ways to improve the efficiency of its datacenter operations. These experiments include an experimental hydrogen fuel cell in the datacenter's power loop, chilled water storage for improved chiller efficiency during the hottest times of the day, shifting loads to different times of the day when the power grid is not maxed out, and some novel cooling technologies projects.

TACC is also looking for ways to improve how the power within the systems is managed. One method currently in place is capping power at the module level for cards such as the Grace-Hopper instead of at the server level. This method substantially reduces the need to build out the datacenter power infrastructure. For example, if built this way, the upcoming 2025 Horizon project could be below 9MW instead of the projected 10MW.

Stanzione emphasized the importance of focusing on software optimization rather than solely on data center efficiency. He says that most datacenters are already fairly efficient, and at most, efficiency efforts will result in 10-15% improvements. He suggested that real change and significant improvements can be achieved by optimizing software to make better use of computational resources.

FIGURE 5

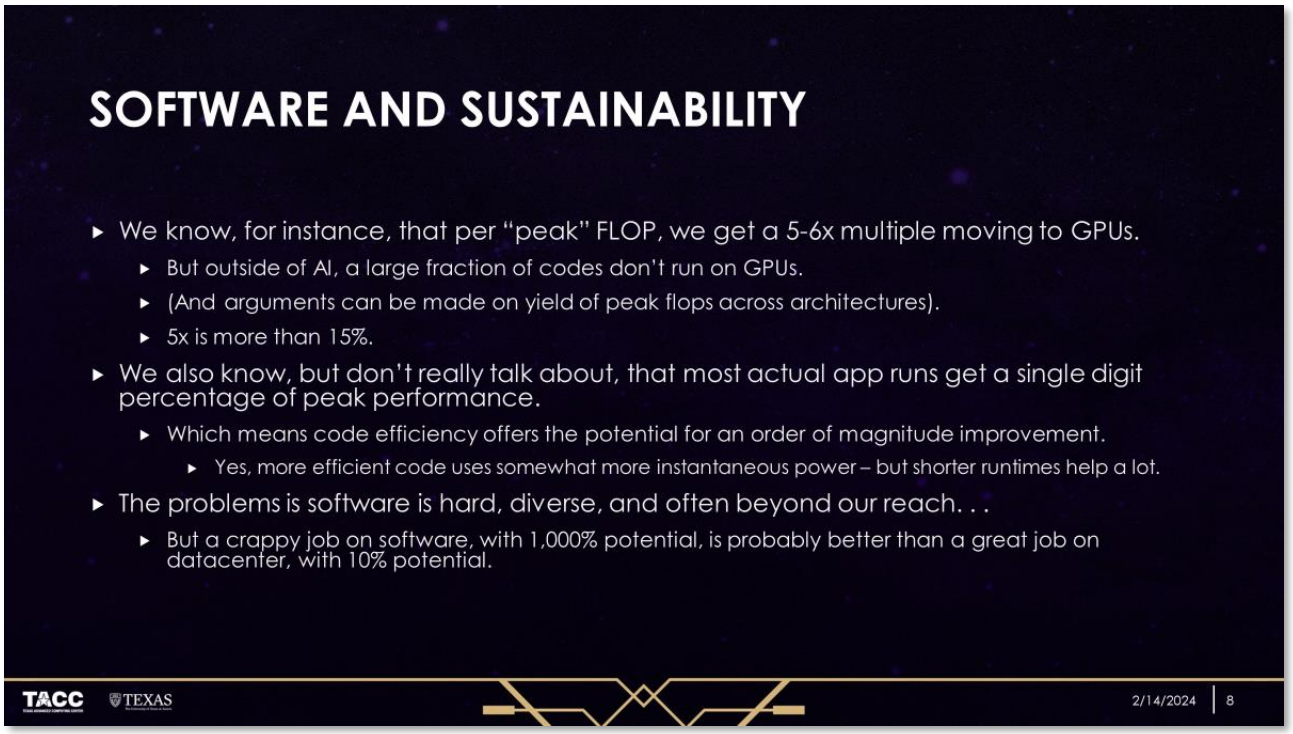

Source: TACC, 2023

Stanzione stressed the critical importance of addressing software challenges for achieving real change in computational performance. Despite advancements like GPUs offering significant peak flop efficiency, there remains vast untapped potential in software optimization. Stanzione suggests that software improvements could potentially yield an order of magnitude enhancement, surpassing the relatively smaller gains achievable through data center efficiency enhancements.

He highlighted the difficulty of software optimization due to its complexity and diversity, compounded by the fact that the responsibility for software lies primarily with users rather than datacenter management. He discussed efforts to incentivize software optimization, such as monitoring performance metrics and creating awareness of inefficiencies. He suggested exploring alternative charging mechanisms, like billing based on energy consumption, to incentivize users to prioritize efficiency over speed.

FIGURE 6

Source: TACC, 2023

Stanzione acknowledged the challenges in implementing such changes but suggested that creating awareness and incentivizing users could drive significant improvements in computational efficiency. He proposed starting with small steps, such as providing users with energy consumption data for their jobs, to gradually build awareness and encourage more efficient computing practices. He also discussed the potential of charging users based on energy consumed, rather than wall-clock hours as an incentive for efficient computing practices. He noted the abundance of wind generation in West Texas and the possibility of leveraging excess energy for computational tasks.

Stanzione underscored the critical importance of software optimization in realizing computational efficiency gains and outlined strategies for incentivizing users to write more efficient code. He highlighted ongoing efforts at TACC to monitor job performance and explore alternative charging mechanisms to promote energy-efficient computing practices.

For more information or to view this and other presentations given at HPC User Forums dating back to 2008, visit www.hpcuserforum.com.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology, and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). Hyperion Research provides thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com