Texas A&M High Performance Research Computing (HPRC) Resources to Build Your Digital Twin Faster

Digital Twin Workshop
February 8, 2022

High Performance
Research Computing
DIVISION OF RESEARCH

High Performance Research Computing

Our Mission:

- Provide **computing** resources
- Provide consulting, technical guidance, and training to support users of these resources.
- Collaborate on computational and data-enabled research.

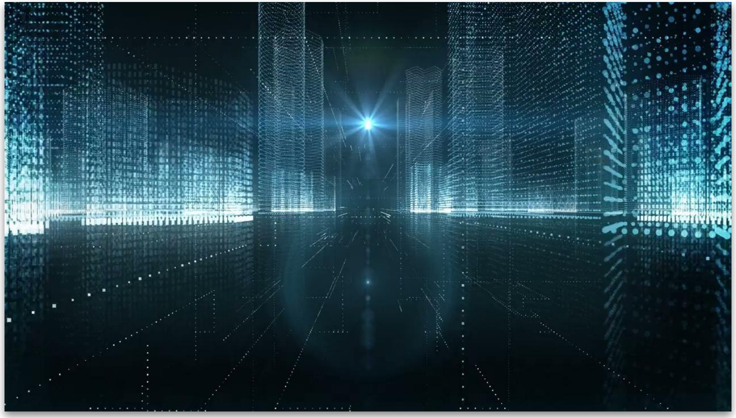

HPRC Services

- **Free of charge** to all faculty, postdocs, research staff, students and external collaborators
- Computing cycles for research and university course purposes
- **Application is required for access**
- User Services
  - Helpdesk: New user start-up assistance and general support
  - Training: Short Courses, Workshops, & YouTube videos
  - Advanced Support: Software and research consulting
    - Expertise in many science and engineering research domains
- Access to state and national advanced computing resources

Digital Twins

0. Standalone twin
1. Descriptive twin
2. Informative twin
3. Predictive twin
4. Comprehensive twin
5. Autonomous twin

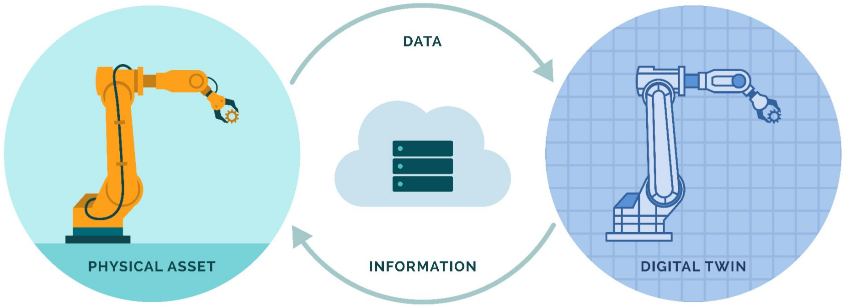

High Performance Research Computing Clusters

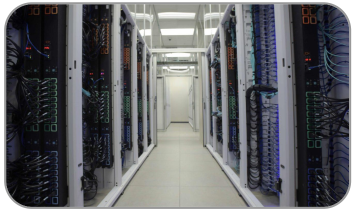

	Grace	Terra	ViDaL
Total Nodes (Cores)	925 (44,656)	307 (8,512)	24 (1,120)
General Nodes	48 cores 384GB	28 cores 64GB	40 cores 192 GB
Features	GPUs (A100, RTX 6000, T4) Large Memory Nodes	GPUs (K80, V100) KNL	**Compliant Computing** GPUs (V100) Large Memory Nodes
Interconnect	HDR100 InfiniBand	Omni-Path	40Gb Ethernet
Global Disk (raw)	8.9 PB	7.4 PB	2 PB

https://hprc.tamu.edu/resources

NSF MRI FASTER

Fostering Accelerated Scientific Transformations, Education, and Research

- **Composable** software-hardware approach
- 184-Intel Ice Lake nodes (11,520-core) with InfiniBand. (64-core, 256GB memory, and 3.84TB NVMe disk per node)
- **NVIDIA GPUs**: 200x T4, 40x A100, 10x A10, 4x A30, and 8x A40 GPUs
- Each node can compose up to 20 GPUs.

This project is supported by NSF award #2019129

High Performance Computing (HPC) Architecture Comparison

Legacy HPC

- Built on Converged HW
- Static Hardware Design
- Fixed GPUs/Accelerators
- Fixed Memory
- Legacy Storage: SATA and SAS

Modern HPC

- Built on Disaggregated HW
- Composable Hardware Platform
- Composable GPUs/Accelerators
- Composable Memory - Optane
- Modern Storage: NVMe-oF

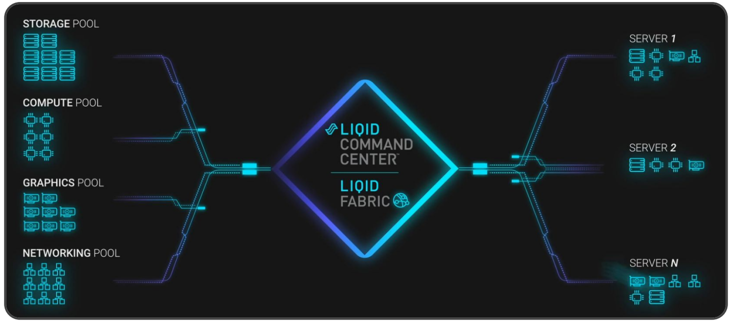

Modern HPC Platforms Support Composable GPUs/Accelerators and Memory

ACES - Accelerating Computing for Emerging Sciences

ACES is an innovative advanced computational prototype to be developed by Texas A&M University partnering with TACC and UIUC.

ACES
ACCELERATING COMPUTING FOR EMERGING SCIENCES

HPCwire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

- Home
- Technologies
- Sectors
- COVID-19
- AI/ML/DL
- Exascale
- Specials
- Resource Library
- Podcast
- Events
- Job Bank
- About
- Our Authors
- Solution Channels
- Subscribe

Three Universities Team for NSF-Funded 'ACES' Reconfigurable Supercomputer Prototype

By Oliver Peckham

September 23, 2021

As Moore's law slows, HPC developers are increasingly looking for speed gains in specialized code and specialized hardware – but this specialization, in turn, can make testing and deploying code trickier than ever. Now, researchers from Texas A&M University, the University of Illinois at Urbana-Champaign and the University of Texas at Austin have teamed, with NSF funding, to build a $5 million prototype supercomputer ("ACES") with a dynamically configurable smörgåsbord of hardware, aiming to support developers as hardware needs grow ever more diverse.

ACES (short for "Accelerating Computing for Emerging Sciences") is presented as an "innovative composable hardware platform." ACES will leverage a PCIe-based composable framework from Liqid to offer access to Intel's high-bandwidth memory Sapphire Rapids processors and more than 20 accelerators: Intel FPGAs; NEC Vector Engines; NextSilicon co-processors; Graphcore IPUs (Intelligence Processing Units); and Intel's forthcoming Ponte Vecchio GPUs. All this hardware will be coupled with Intel Optane memory and DDN Lustre Storage and connected with Mellanox NDR 400Gbps networking.

This project is supported by NSF award #2112356

"ACES will enable applications and workflows to dynamically integrate the different accelerators, memory, and in-network computing protocols to glean new insights by rapidly processing large volumes of data," the NSF grant reads, "and provide researchers with a unique platform to produce complex hybrid programming models that effectively supports calculations that were not feasible before."

https://www.hpcwire.com/2021/09/23/three-universities-team-for-nsf-funded-aces-reconfigurable-supercomputer-prototype/

ACES System Description

ACES
ACCELERATING COMPUTING FOR EMERGING SCIENCES

Component	Quantity	Description
Allocatable resources		Total cores: 11,520
CPU-centric computing with variable memory requirements	120 nodes (11,520 cores)	Dual Intel Sapphire Rapids 2.1 GHz 48 core processors with HBM2e memory 96 cores per node, 512 GB memory, 1.6 TB NVMe storage (PCIe 5.0), NVIDIA Mellanox NDR 200 Gbps Infiniband
Composable infrastructure	120 nodes	Dynamically reconfigurable infrastructure that allows up to 20 PCIe cards (GPU, FPGA, VE, etc.) per compute node
Data transfer nodes	2 nodes	Same as compute nodes, 100 Gbps network adapter

To be deployed in late 2022

ACES - Accelerating Computing for Emerging Sciences (To be deployed in 2022)

Component	Quantity*	Description
Graphcore IPU	16	16 IPUs direct-attached to a server
Intel Agilex FPGA	20	Agilex FPGA with a broad hierarchy of memory including DDR5, HBM2e and Optane Persistent Memory
NextSilicon coprocessor	20	Reconfigurable accelerator with an optimizer continuously evaluating application behavior.
NEC Vector Engine	24	Vector computing card with 8 cores and HBM2 memory
Intel Ponte Vecchio GPU	100	Intel GPU for HPC, DL Training, AI Inference
Liqid Intel Optane PCIe SSDs	6	3 TB PCIe SSD cards addressable as memory using Intel Memory Drive Technology

*Estimated quantities

Graphcore IPUs (Intelligence Processing Unit)

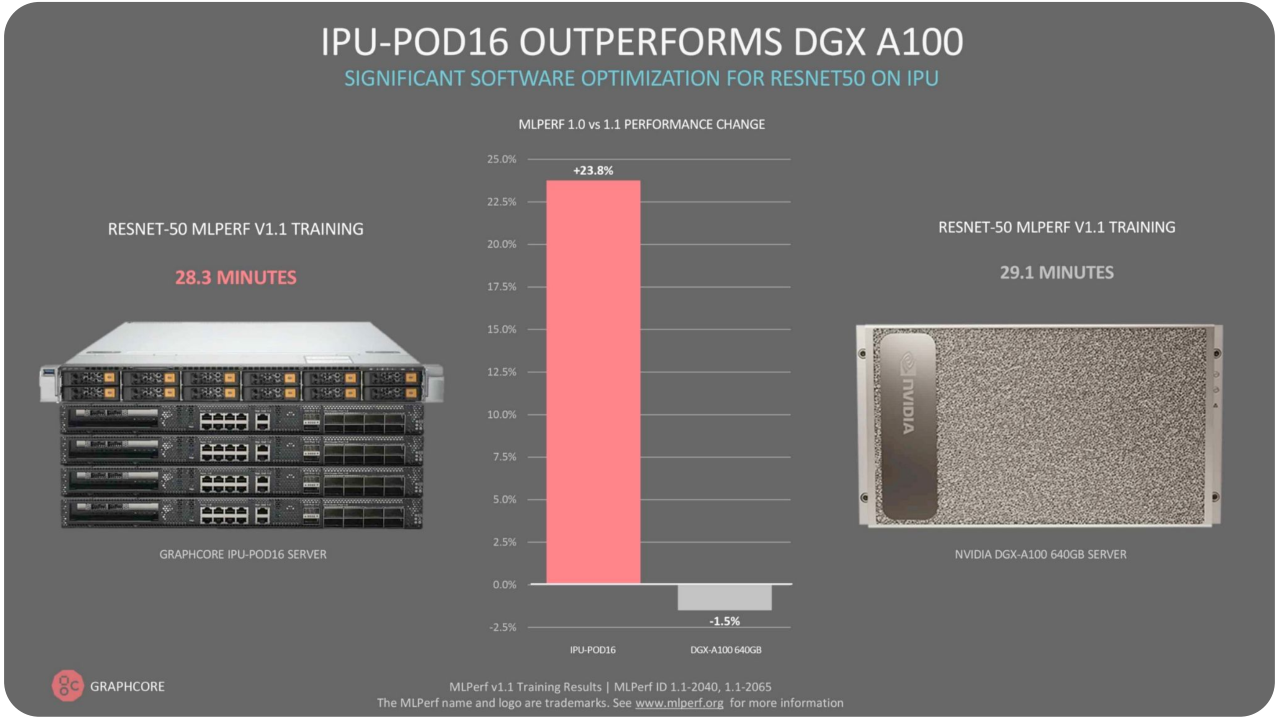

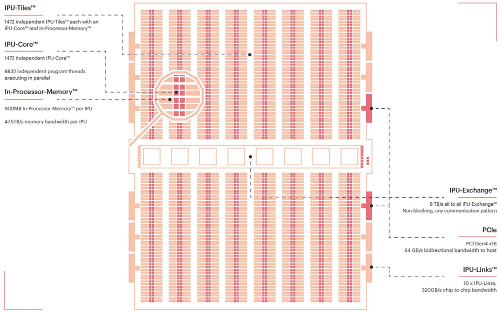

https://www.graphcore.ai/posts/accelerating-resnet50-training-on-the-ipu-behind-our-mlperf-benchmark

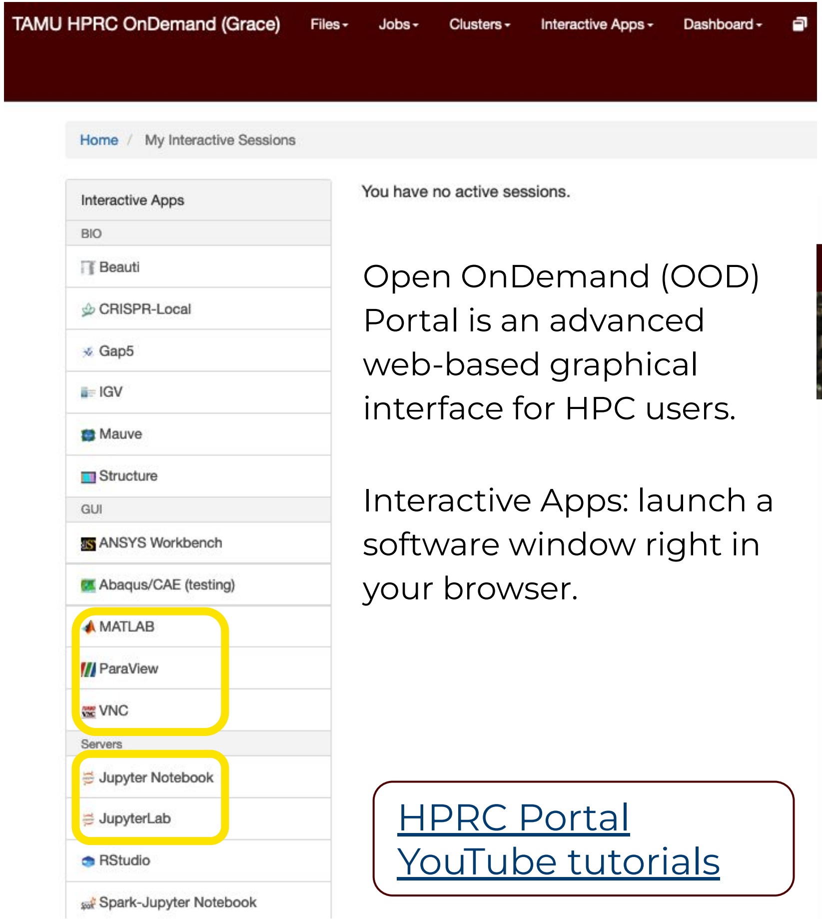

Open OnDemand (OOD) Portal is an advanced web-based graphical interface for HPC users.

Interactive Apps: launch a software window right in your browser.

[HPRC Portal YouTube tutorials]

HPRC Portal

https://portal.hprc.tamu.edu

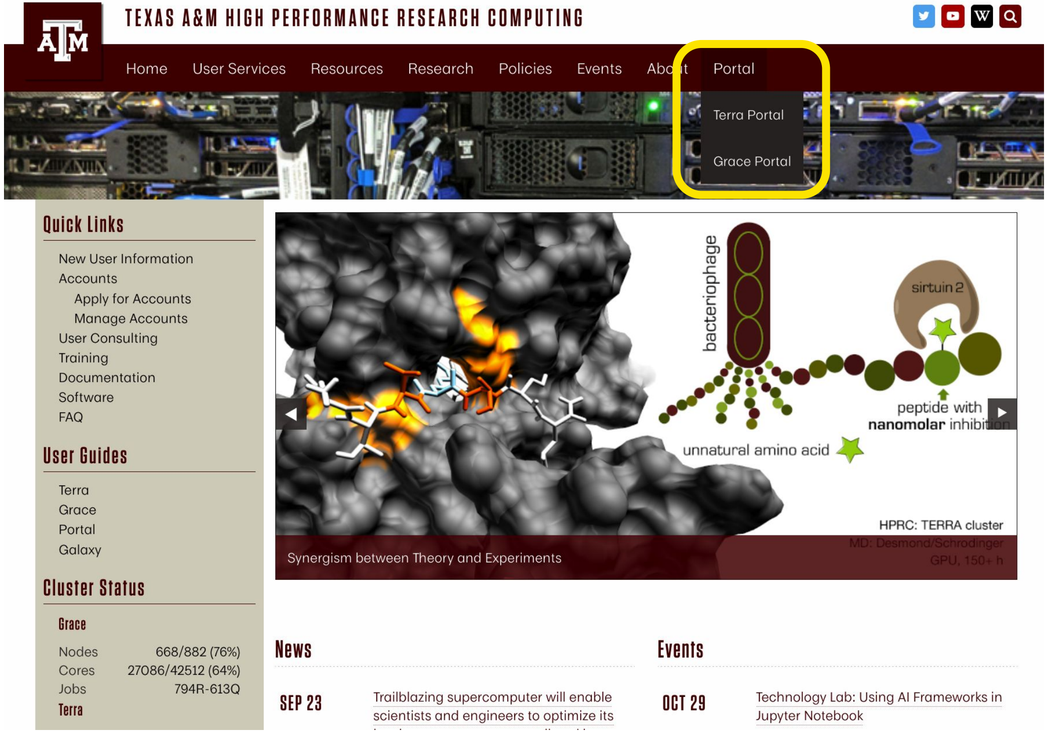

Available Software Modules

https://hprc.tamu.edu/wiki/SW:Modules

mla command to quickly search for installed software:

```
mouse@terra2 ~]$ mla scikit-learn
Using /home/mouse/module.avail.terra
scikit-learn/
scikit-learn/0.18.1-intel-2017A-Python-2.7.12
scikit-learn/0.19.1-foss-2017b-Python-2.7.14
scikit-learn/0.19.1-foss-2017b-Python-3.6.3
scikit-learn/0.19.1-foss-2018a-Python-3.6.4
scikit-learn/0.19.1-fosscuda-2017b-Python-3.6.3
......
scikit-learn/0.21.3-fosscuda-2019b-Python-3.7.4
scikit-learn/0.21.3-intel-2019b-Python-3.7.4
scikit-learn/0.22.1-intel-2019b-Python-3.7.4
scikit-learn/0.23.1-foss-2020a-Python-3.8.2
scikit-learn/0.23.1-fosscuda-2020a-Python-3.8.2
scikit-learn/0.23.1-intel-2020a-Python-3.8.2
scikit-learn/0.23.1-intelcuda-2020a-Python-3.8.2
scikit-learn/0.23.2-foss-2020b
scikit-learn/0.23.2-intel-2020b
scikit-learn/0.23.2-intelcuda-2020b
```

Python
Matlab
Keras
PyTorch
scikit-learn
Pandas
NumPy
Matplotlib
Julia
....
Compilers: C++, Fortran, Intel OneAPI, GNU, ...
CUDA, OpenCL
OpenMPI, IntelMPI
...

Advanced Support Program

Collaborations on computational research projects.

HPRC analysts can contribute expertise in:.

- Software development for research workflows
- Developing GUIs and apps for research projects
- Porting applications to HPC clusters
- Code development, optimizing and analysis on serial and parallel platforms
- Leveraging mathematical libraries
- Workflow automation in scientific processes

Please send us an e-mail: help@hprc.tamu.edu

ASP is supported in part by NSF award #1925764, CC* Team: SWEETER -- SouthWest Expertise in Expanding, Training, Education and Research and NSF award #2112356, Category II: ACES - Accelerating Computing for Emerging Sciences

Partnering on Outreach

Leverage our programs to strengthen your broader impacts

- Teach a short course on computing to the TAMU community.
- Become an instructor in the Summer Computing Academy program camps for middle and high school students.
- Join the the NSF SWEETER CyberTeam to explore computing-driven research and educational partnerships with universities in Texas, New Mexico and Arizona.
- Participate in the NSF BRICCs community to support research computing at smaller institutions and community colleges.
- Make your computing products available on the NSF ACES, NSF FASTER, and NSF Frontera machines.
- Mentor our students in international Student Cluster Competitions.

ACES
ACCELERATING COMPUTING FOR EMERGING SCIENCES

BRICCS
BUILDING RESEARCH INNOVATION AT COMMUNITY COLLEGES

CC* CYBERTEAM SWEETER
RESEARCH | ENABLEMENT | TRAINING

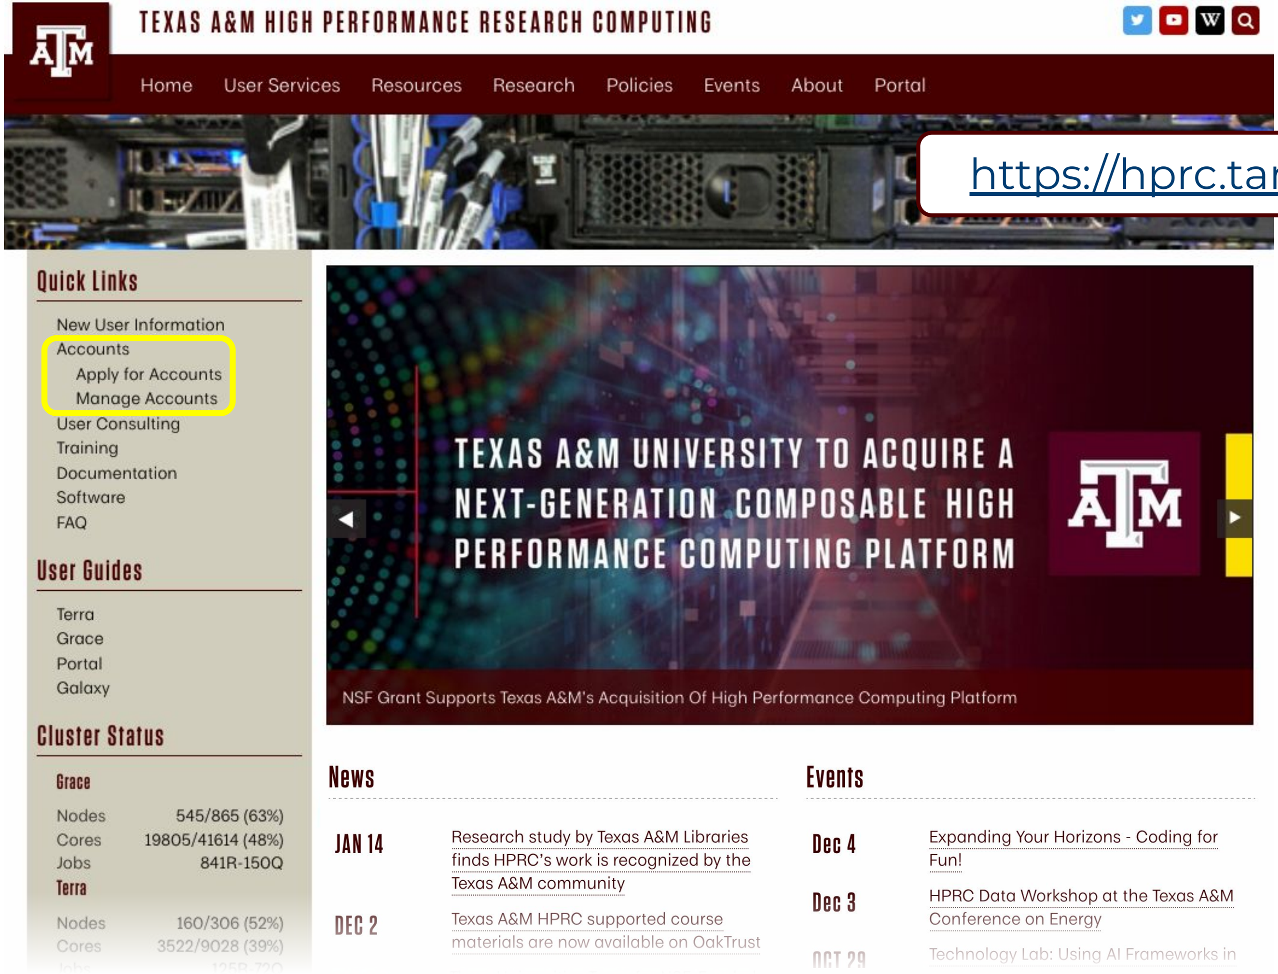

https://hprc.tamu.edu

High Performance
Research Computing
DIVISION OF RESEARCH

hprc.tamu.edu

HPRC Helpdesk:

help@hprc.tamu.edu
Phone: 979-845-0219