Category II: ACES Accelerating Computing for Emerging Sciences

Drs. Honggao Liu, Dhruva Chakravorty, and Lisa M. Perez

High Performance Research Computing
Texas A&M University

October 26, 2021

Goals of ACES

The ACES platform removes significant bottlenecks in advanced computing by introducing the flexibility to compose various components (i.e., processors, accelerators and memory) on an as-needed basis to solve problems that were previously not addressable.

ACES will benefit many research and development projects in the fields of artificial intelligence and machine learning (AI/ML), cybersecurity, health population informatics, genomics and bioinformatics, human and agricultural life sciences, oil & gas simulations, de novo materials design, climate modeling, molecular dynamics, quantum computing architectures, imaging, smart and connected societies, geosciences, and quantum chemistry.

NSF Award # 2112356
https://www.nsf.gov/awardsearch/showAward?AWD_ID=2112356

HPC Architecture Comparison

Legacy HPC

- Built on Converged HW
- Static Hardware Design
- Fixed G P U
- Fixed Memory
- Legacy Storage: SATA and SAS
- Vendor Lock

Next-Generation HPC

- Built on Disaggregated HW
- Composable Hardware Platform
- Composable GPU
- Composable Memory - Optane
- Modern Storage: NVMe-oF
- Open Platform

Next Generation HPC Platform Supports Composable GPU and Composable Memory

Composability Dynamic Research Infrastructure

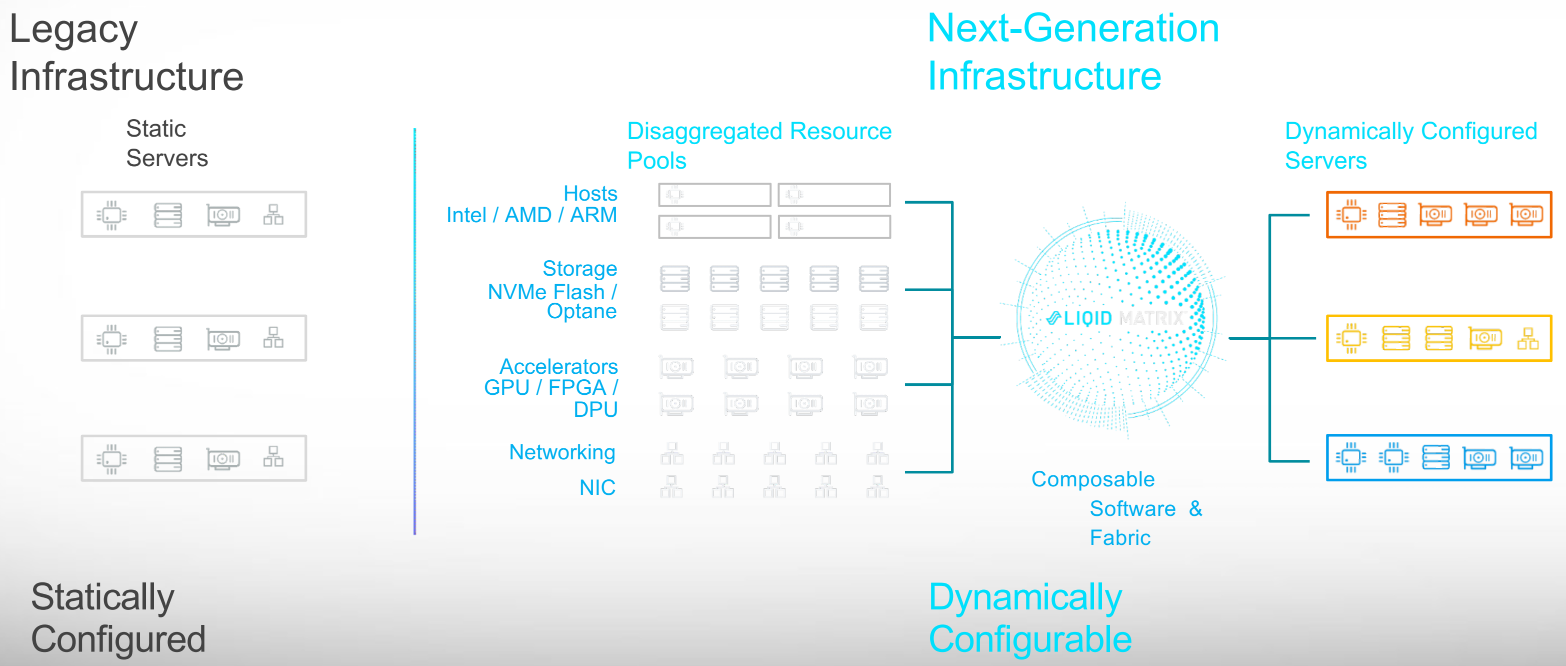

ACES: Composable HPC/AI Platform

User Interface/Portal

Liqid Matrix Software and Fabric

Servers Nodes

Composable Resources

(GPU, FPGA,, VE, Optane, etc.)

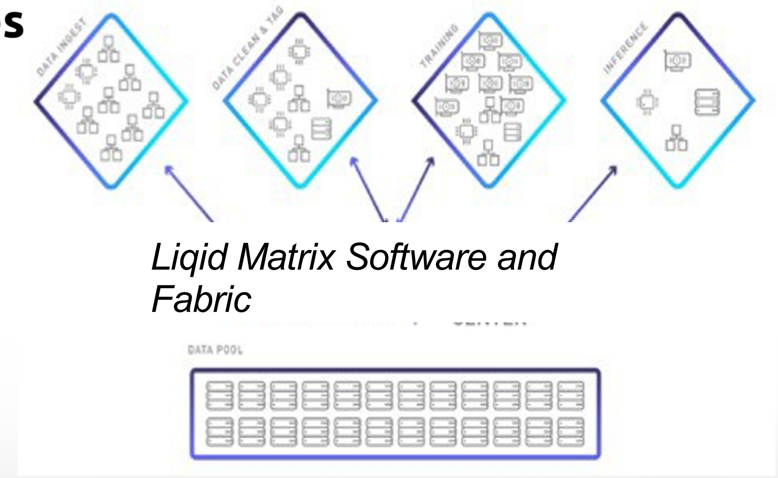

Liqid Matrix Software and Fabric

Dynamic Resource Allocation for Each HPC/AI Workflow

System Description

Component	Quantity	Description
Allocatable resources		Total cores: 11,520 Core hours (SUs) per quarter: 252 M Core hours (SUs) per year: 100 M
CPU-centric computing with variable memory requirements	120 nodes (11,520 cores)	Dual Intel Sapphire Rapids 2.1 GHz 48c processors, 96c per node, 512 GB memory, 1.6 TB NVMe storage (PCIe 5.0), NDR200 InfiniBand adapter
Liqid composable infrastructure	120 nodes	Dynamically reconfigurable infrastructure that allows up to 20 PCIe cards (GPU, FPGA, VE, etc.) per compute node

System Description - Accelerators

Component	Quantity	Description
Graphcore IPU	16	16 IPUs direct-attached to a server.
Future Intel Agilex FPGA replacement	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

System Description - Infrastructure

Component	Quantity	Description
Interconnect		NVIDIA Mellanox NDR 400 Gbps Infiniband
Storage for active accounts		2 PB usable of DDN Lustre divided into high speed scratch (90%) and general purpose/home (10%) with DDN Data Services for Lustre
Login nodes	2 nodes (192 cores)	Dual Intel Sapphire Rapids 2.1 GHz 48c processors, 512 GB memory, 1.6 TB NVMe storage (PCIe 5.0), NDR200 InfiniBand adapter
Management nodes	3 nodes	Same as compute nodes
Data transfer nodes	2 nodes	Same as compute nodes, 100 Gbps network adapter

Research Workflows - CPU / GPU Centric

Hardware Profile	Application examples
Primarily distributed memory parallel	WRF, CHARMM, STAR, OpenFOAM, CESM, HMMER, STAR, METIS, Gaussian, Bioconda, Molecular Dynamics (MD, NAMD etc.)
Double precision GPU intensive	MD (LAMMPS, GROMACS, NAMD) VASP, Quantum Espresso, WRF
Half/Single precision GPU intensive	LAMMPS, TensorFlow, GROMACS, Python, AMBER
Large memory, variable core count	CESM, CWPSU, Bowtie, R, Orca, GAMESS, Cufflink
Large memory, GPU intensive	ParaView, TensorFlow (Keras), Pytorch, Caffe, Orca, WRF

Research Workflows - Accelerators

Hardware Profile	Application examples
Vector Engine	Plasma Simulation, Weather / Climate Simulation, NumPy acceleration, Earth Sciences, Oil and Gas - Seismic Imaging and Reservoir Simulation, Chemistry – VASP and Quantum ESPRESSO, AI / ML – Statistical Machine Learning and Data Frame (Apache Spark / Big Data)
Graphcore IPU	LSTM Neural Networks, Markov Chain Monte Carlo, Graph data, Natural Language Processing (Deep Learning)
FPGA	Big Data, Deep Learning Inference, Streaming data analysis, Genomics, AI Models for embedded use cases, CXL memory interface, Microcontroller emulation for Autonomy simulations, MD codes
Optane PCIe SSD	Bioinformatics, Computational Fluid Dynamics, R, WRF, MD codes
NextSilicon	Computational Fluid Dynamics (OpenFOAM), Molecular Dynamics (NAMD, AMBER, LAMMPS), Weather/Environment modeling (WRF), Biosciences (BLAST), Graph Search (Pathfinder), Cosmology (HACC), Quantum ChromoDynamics (MILC)

Project Deliverables - Phase 1

1. Project Management Board Meeting - **Taking place**
2. Form external Advisory Board - **Completed**
3. Meet and update NSF Program Director biweekly - **Organized**
4. Biweekly project discussions with industry partners - **Organized, taking place daily**
5. Create ACES project website, google group and slack channel - **Completed**
6. Develop portal, and documentation - **Active work**
7. Phase 1 composable computing platform that includes 1 Graphcore node and 4 IPUs, 2 Intel FPGAs, 2 Next Silicon Coprocessors, and 8 NEC Vector Engine cards, and 8 composable Liqid Intel Optane cards in the FASTER machine - testing completed by **January 2022**
8. Phase 1 prototype available to the NSF national research community via XSEDE - **Q1 2022**
9. Offer users training on the ACES environments - **Q1 2022, Active work**

Project Deliverables - Phase 2

1. Configure, acquire and deploy the ACES phase 2
    a. Phase 2 equipment arrives in West Campus Data Center - **June 30, 2022**
    b. Installation, testing and acceptance of Phase 2. Phase 1 equipment will be integrated into Phase 2 - **September 30, 2022**
    c. User testing begins - **October 1, 2022**
    d. Operations for the ACES system - **November 1, 2022**
    e. Allocations for users through XRAC or its replacement - **Q4 2022**
2. Assist national users in early accessing ACES and testing novel technologies, and effectively using the ACES instrument - **Q1 2022 onward**
3. Informal CI-training courses geared toward training middle- and high-school students and teachers, and college students and researchers - **Q1 2022 onward**

ACES Leadership Team

Name	Title	Organization
Honggao Liu	Executive Director	Texas A&M HPRC
Dhruva Chakravorty	Associate Director	Texas A&M HPRC
Lisa Perez	Associate Director	Texas A&M HPRC
Timothy Cockerill	Director of User Services	TACC at UT Austin
Shaowen Wang	Professor and Department Head	University of Illinois at Urbana Champaign
Francis Dang	Associate Director	Texas A&M HPRC
Ed Pierson	V.P. for IT and Chief Information Officer	Texas A&M Division of Information Technology
Costas Georghiades	Sr. Assoc. V.P. for Research	Texas A&M Division of Research

Project Leadership Roles

Name	Project Role	Organization	Responsibilities
Honggao Liu	PI, Director	TAMU HPRC	Project oversight, ACCESS-CO & NSF coordination
Dhruv Chakravorty	co-PI, Program Manager	TAMU HPRC	User support, XSEDE liaison, and research enablement
Lisa Perez	co-PI	TAMU HPRC	Education, Outreach and Training, Scientific Software application testing and evaluation.
Tim Cockerill	co-PI	TACC at UT Austin	Allocations, manage TACC staff contributions, and Frontera work
Shaowen Wang	co-PI	University of Illinois at Urbana Champaign	Application based testing and evaluation.
Francis Dang	Senior Investigator	TAMU HPRC	Operations, networking and infrastructure management
Ed Pierson	Senior Investigator	TAMU Division of Information Technology	Integration of this project with IT facilities and services, security, networking
Costas Georghiades	Senior Investigator	TAMU Division of Research	award management, acquisition and subcontract management

ACES Project Personnel

HPRC Team

Mark Huang
Marinus Pennings
Charles Michael Dickens
Richard Lawrence

TAMU IT

- Edwin Pierson (Senior Investigator)
- Michael Afshin Sardaryzadeh (Senior Investigator)
- William Deigaard (Senior Investigator)

Software Applications Team

- Abishek Gopal (Senior Investigator)
- Jian Tao (Senior Investigator)
- Xinyue Ye (Senior Investigator)
- Zhe Zhang (Senior Investigator)

External Advisory Council

- Will meet biannually to evaluate the project's progress
  - Gabrielle Allen (University of Wyoming)
  - John Goodhue (MGHPCC)
  - Richard Gerber (NERSC)
  - Victor Hazlewood (University of Tennessee)
  - Dan Katz (NCSA)
  - Anita Nikolich (UIUC)
  - Barry Schneider (NIST)
  - Carol Song (Purdue)
  - Dan Stanzione (TACC at UT Austin)

Contact Information

- Helpdesk: help@hprc.tamu.edu
- Join the ACES google group: aces@lists.tamu.edu
- Website: https://hprc.tamu.edu/aces
- Join nsf-aces-category-ii Slack workspace
- Join #aces channel on the SWEETER Slack workspace
- Twitter: @aces_nsf