

FASTER

Next Generation Composability

March 6, 2022



High Performance
Research Computing
DIVISION OF RESEARCH

Texas A&M University



NSF MRI FASTER

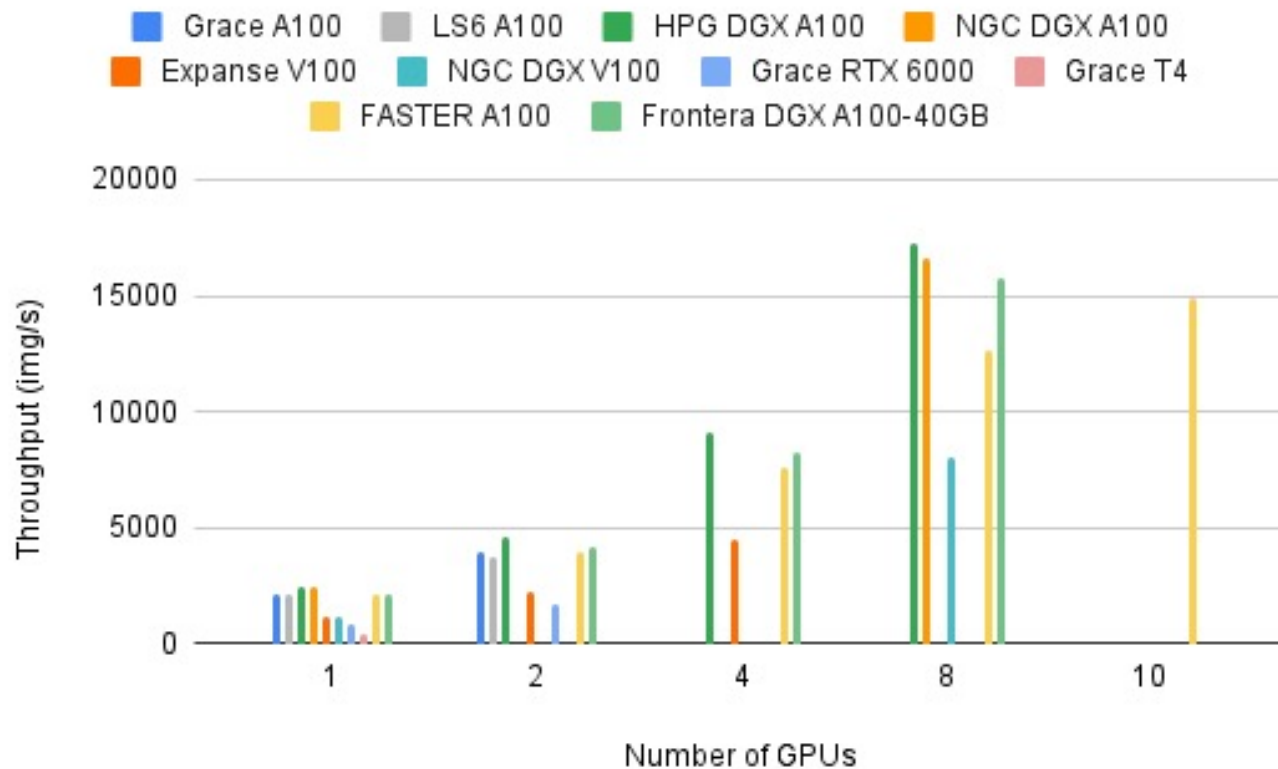
Fostering Accelerated Scientific Transformations, Education, and Research

- **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.

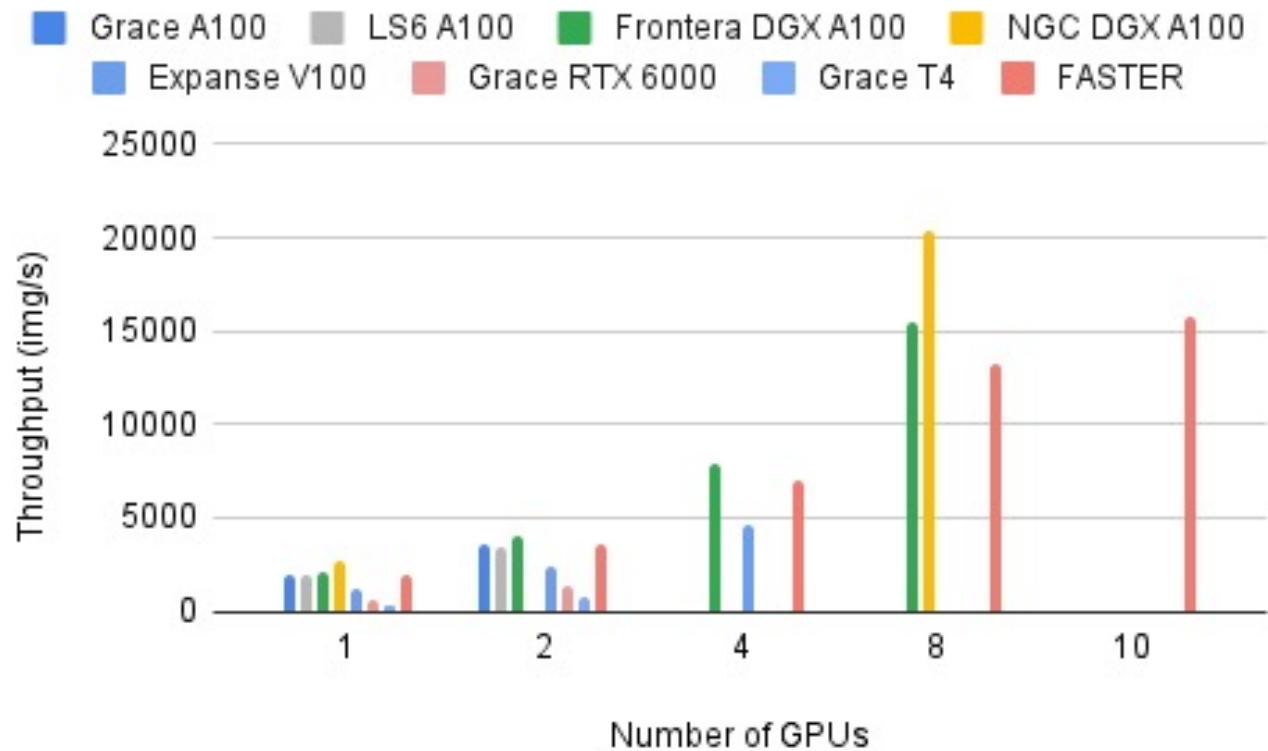


This project is supported by NSF award #[2019129](#)

PyTorch ResNet50 on FASTER

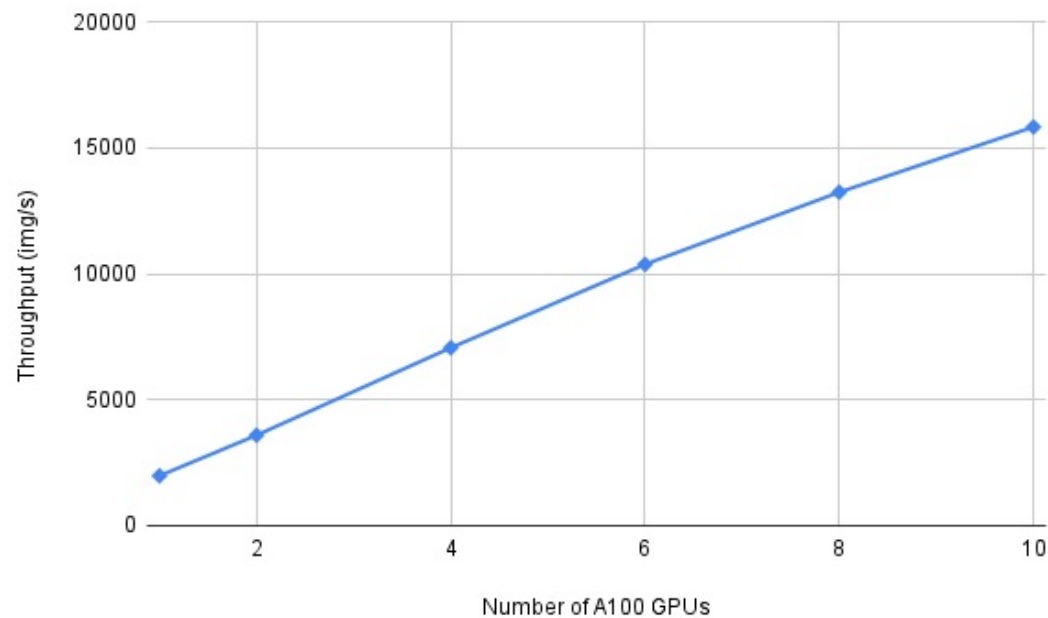


TF ResNet50 on FASTER



This project is supported by NSF award #[2019129](#)

AI/ML Benchmarks on FASTER



This project is supported by NSF award #[2019129](#)



<https://hprc.tamu.edu>

Quick Links

- New User Information
- Accounts
 - Apply for Accounts
 - Manage Accounts
- User Consulting
- Training
- Documentation
- Software
- FAQ

User Guides

- Terra
- Grace
- Portal
- Galaxy

Cluster Status

Grace	
Nodes	545/865 (63%)
Cores	19805/41614 (48%)
Jobs	841R-150Q
Terra	
Nodes	160/306 (52%)
Cores	3522/9028 (39%)

TEXAS A&M UNIVERSITY TO ACQUIRE A NEXT-GENERATION COMPOSABLE HIGH PERFORMANCE COMPUTING PLATFORM

NSF Grant Supports Texas A&M's Acquisition Of High Performance Computing Platform

News

- JAN 14** [Research study by Texas A&M Libraries finds HPRC's work is recognized by the Texas A&M community](#)
- DEC 2** [Texas A&M HPRC supported course materials are now available on OakTrust](#)

Events

- Dec 4** [Expanding Your Horizons - Coding for Fun!](#)
- Dec 3** [HPRC Data Workshop at the Texas A&M Conference on Energy](#)
- DEC 29** [Technology Lab: Using AI Frameworks in](#)



Exhaustive Software Modules

Python
Matlab
Keras
PyTorch
scikit-learn
Pandas
NumPy
Matplotlib

...
Compilers: C++,
Fortran, Intel
OneAPI, GNU, ...
CUDA, OpenCL
OpenMPI,
IntelMPI
...

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



Sample Benchmarking Job Script

```
1  #!/bin/bash
2
3  ##NECESSARY JOB SPECIFICATIONS
4  #SBATCH --job-name=NvidiaDeepLearningBenchmark
5  #SBATCH --time=01:00:00
6  #SBATCH --ntasks=24                #Request 24 tasks
7  #SBATCH --mem=180G                 #Request 180GB per node
8  #SBATCH --output=NvidiaDeepLearningBenchmark.%j #Send stdout/err to "Example40ut.[jobid]"
9  #SBATCH --gres=gpu:a100:1         #Request 1 A100 GPU per node (can be shared)
10 #SBATCH --partition=gpu           #Request the GPU partition/queue
11 ##SBATCH --reservation=benchmarking
12 #SBATCH --account=123456
```

Web Portal For Computing

<https://portal.hprc.tamu.edu>

TEXAS A&M HIGH PERFORMANCE RESEARCH COMPUTING

Home User Services Resources Research Policies Events About **Portal** Terra Portal Grace Portal

TERRA TOOLBOX Request Assistant Request Software

Quick Links

- New User Information
- Accounts
- Apply for Accounts
- Manage Accounts
- User Consulting
- Training
- Documentation
- Software
- FAQ

User Guides

- Terra
- Grace
- Portal
- Galaxy

Cluster Status

Grace	
Nodes	668/882 (76%)
Cores	27086/42512 (64%)
Jobs	794R-613Q
Terra	

Cluster Statistics

Node Utilization

- Allocated (Red)
- Mixed (Blue)
- Idle (Green)

Core Utilization

- Used (Red)
- Free (Green)

Jobs

Running	Pending
505	9351

Accounts

Account	Default	Allocation	Used	Balance
122853910111	Set Default	20000	0	20000
122853910233	Set Default	200000	198148.06	1851.94
122853913205	default	5000	-3735.43	1264.57
122853915531	Set Default	200000	50000	150000

Disk Quotas

Disk	Disk Usage	Limit	File Usage	Limit
/home	2.76 GB (27.59 %)	10 GB	8882 (88.82 %)	10000
/scratch	282.87 GB (27.62 %)	1 TB	100574 (50.29 %)	200000

Your Jobs

Job ID	Name	State	Partition	Log	Error	Kill
7522977	...rd/sys/vnc	RUNNING	gpu	Log	Error	Kill

News

SEP 23 Trailblazing supercomputer will enable scientists and engineers to optimize its

Events

OCT 29 Technology Lab: Usir Jupyter Notebook

Synergism between Theory and Experiments

bacteriophage

unnatural amino acid



HPRC Portal - Interactive Apps

BIO

- Beauti
- DIYABC
- FigTree
- IGV
- JBrowse
- Krait
- Mauve
- Structure
- Tracer
- CRISPR-Local
- Gap5

GUI

- ANSYS Workbench
- Abaqus/CAE
- LS-PREPOST
- MATLAB
- ParaView
- VNC

Servers

- Jupyter Notebook
- JupyterLab
- RStudio
- Spark-Jupyter Notebook

Imaging

- AFNI
- Chimera
- Coot
- Diffusion Toolkit & TrackVis
- FSL
- Fiji
- ICY
- ImageJ
- Vaa3D
- cisTEM

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
Liquid Intel Optane PCIe SSDs	6	3 TB PCIe SSD cards addressable as memory using Intel Memory Drive Technology

*Estimated quantities



High Performance
Research Computing
DIVISION OF RESEARCH

hprc.tamu.edu

HPRC Helpdesk:

help@hprc.tamu.edu

Phone: 979-845-0219