Argonne NATIONAL LABORATORY

Artificial Intelligence Testbeds at Argonne National Laboratory

Venkatram Vishwanath,
Argonne Leadership Computing Facility
venkat@anl.gov

July 21, 2022

Argonne Leadership Computing Facility

The Argonne Leadership Computing Facility provides world-class computing resources to the scientific community.

- Users pursue scientific challenges
- In-house experts to help maximize results
- Resources fully dedicated to open science

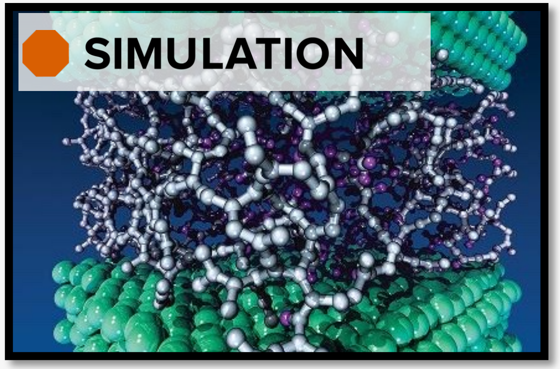

ALCF offers different pipelines based on your computational readiness. Apply to the allocation program that fits your needs.

Architecture supports three types of computing

- Large-scale Simulation (PDEs, traditional HPC)
- Data Intensive Applications (scalable science pipelines)
- Deep Learning and Emerging Science AI (training and inferencing)

ALCF Systems Evolution

5.7 TF

IBM BG/L
2004

557 TF

Intrepid
IBM BG/P
2007

10 PF

Mira
IBM BG/Q
2012

11.7 PF

Theta
Intel-Cray XC40
2017

15.6 PF

ThetaGPU
NVIDIA
DGX A100
2020

44 PF

Crux
HPE-AMD
+
Polaris
HPE
2021

> 2 EF

Aurora
Intel-HPE
2022

JLSE (2013) **AI Testbed** (2020) **Edge Testbed** (2021)

Aurora

Leadership Computing Facility
Exascale Supercomputer

PEAK PERFORMANCE

≧ 2 Exaflops DP

Intel GPU

Ponte Vecchio

Intel Xeon PROCESSOR

Sapphire Rapids wt HBM

PLATFORM

HPE Cray-Ex

Compute Node
2 SPR+HBM processor;
6 PVC; Unified
Memory Architecture; 8 fabric endpoints;

GPU Architecture
Xe arch-based “Ponte Vecchio” GPU
Tile-based chiplets
HBM stack
Foveros 3D integration

System Interconnect
HPE Slingshot 11; Dragonfly topology with adaptive routing

Network Switch
25.6 Tb/s per switch, from 64–200 Gb/s ports (25 GB/s per direction)

Node Performance
>130 TF

System Size
>9,000 nodes

Aggregate System Memory
>10 PB aggregate System Memory

High-Performance Storage
220 PB @ EC16+2**, ≧25 TB/s DAOS**

Programming Models
oneAPI, MPI, OpenMP, C/C++,
Fortran, SYCL/DPC++
Python-based environments
Machine learning and Deep learning fraameworks

AURORA ESP Data and Learning Projects and Methods

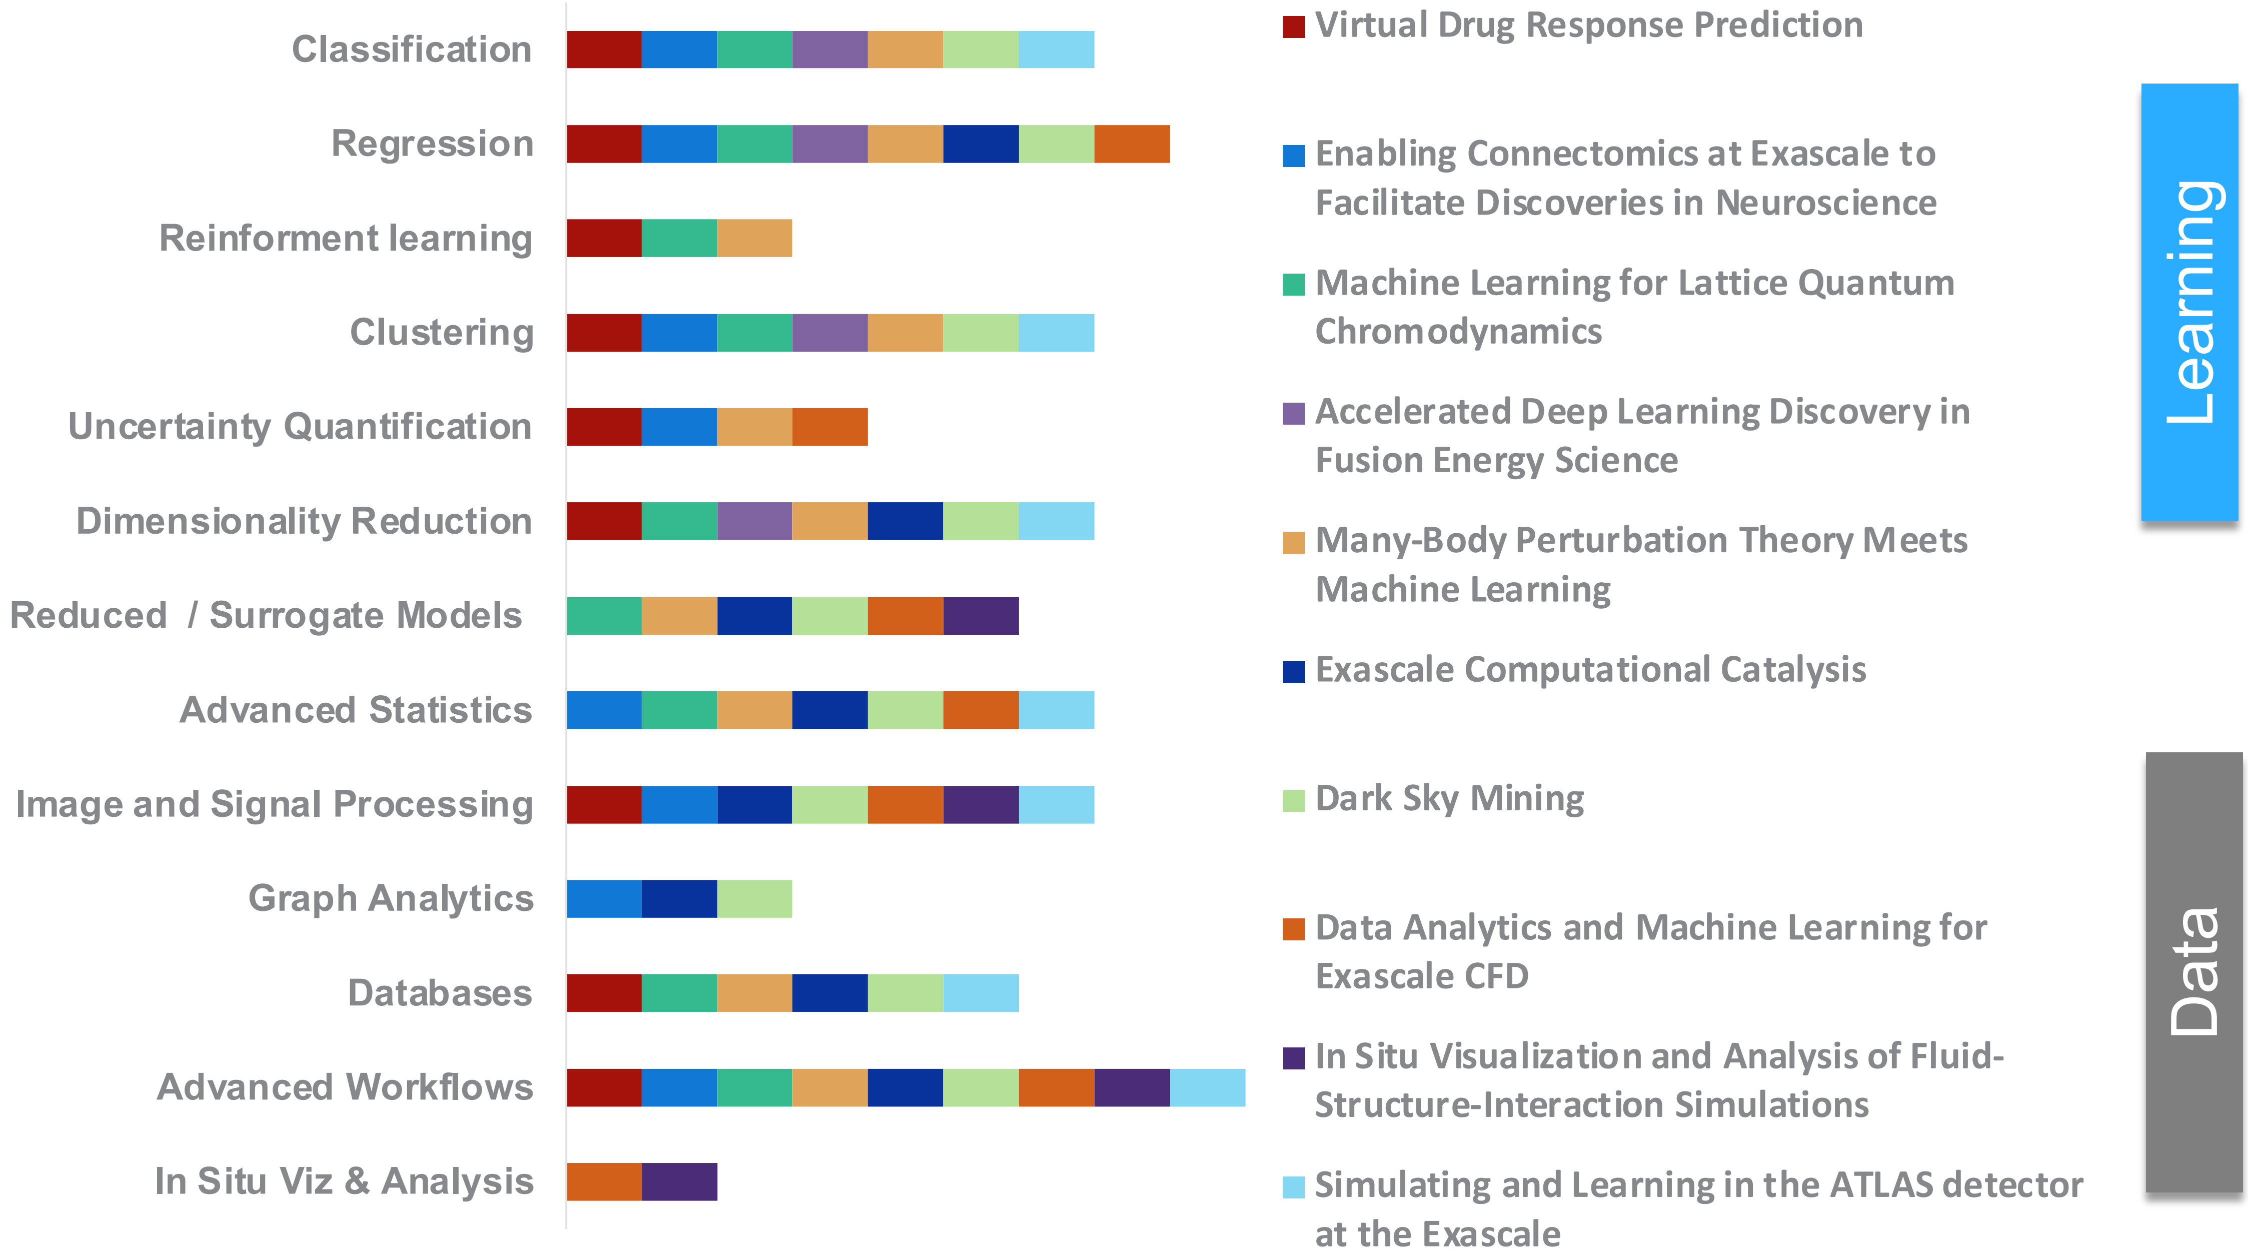

SURGE OF SCIENTIFIC MACHINE LEARNING

- Simulations/ surrogate models
  - Replace, in part, or guide simulations with AI-driven surrogate models
- Data-driven models
  - Use data to build models without simulations
- Co-design of experiments
  - AI-driven experiments

Design infrastructure to facilitate and accelerate AI for Science applications

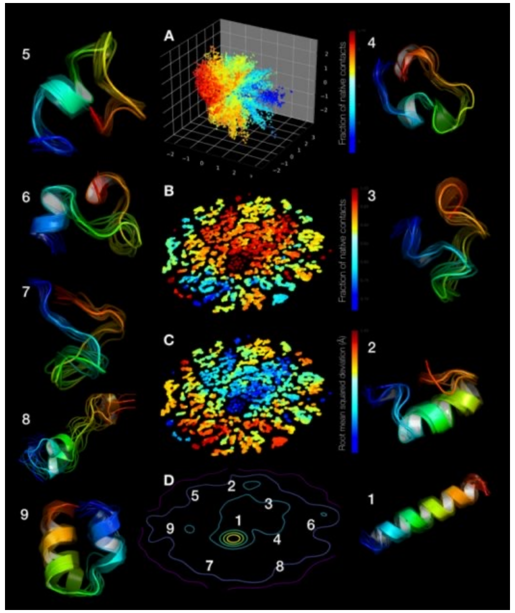

Protein-folding

Braggs Peak

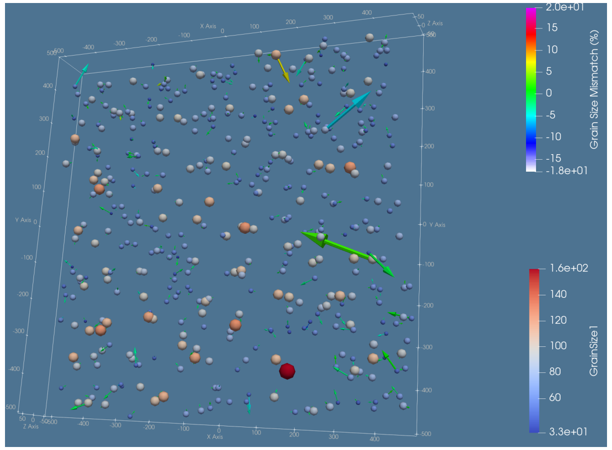

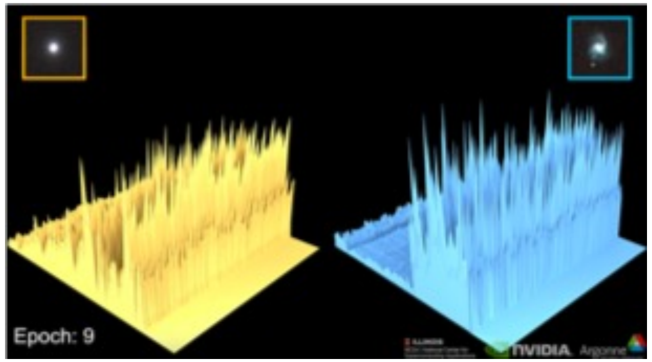

Galaxy Classification

INTEGRATING AI SYSTEMS IN FACILITIES

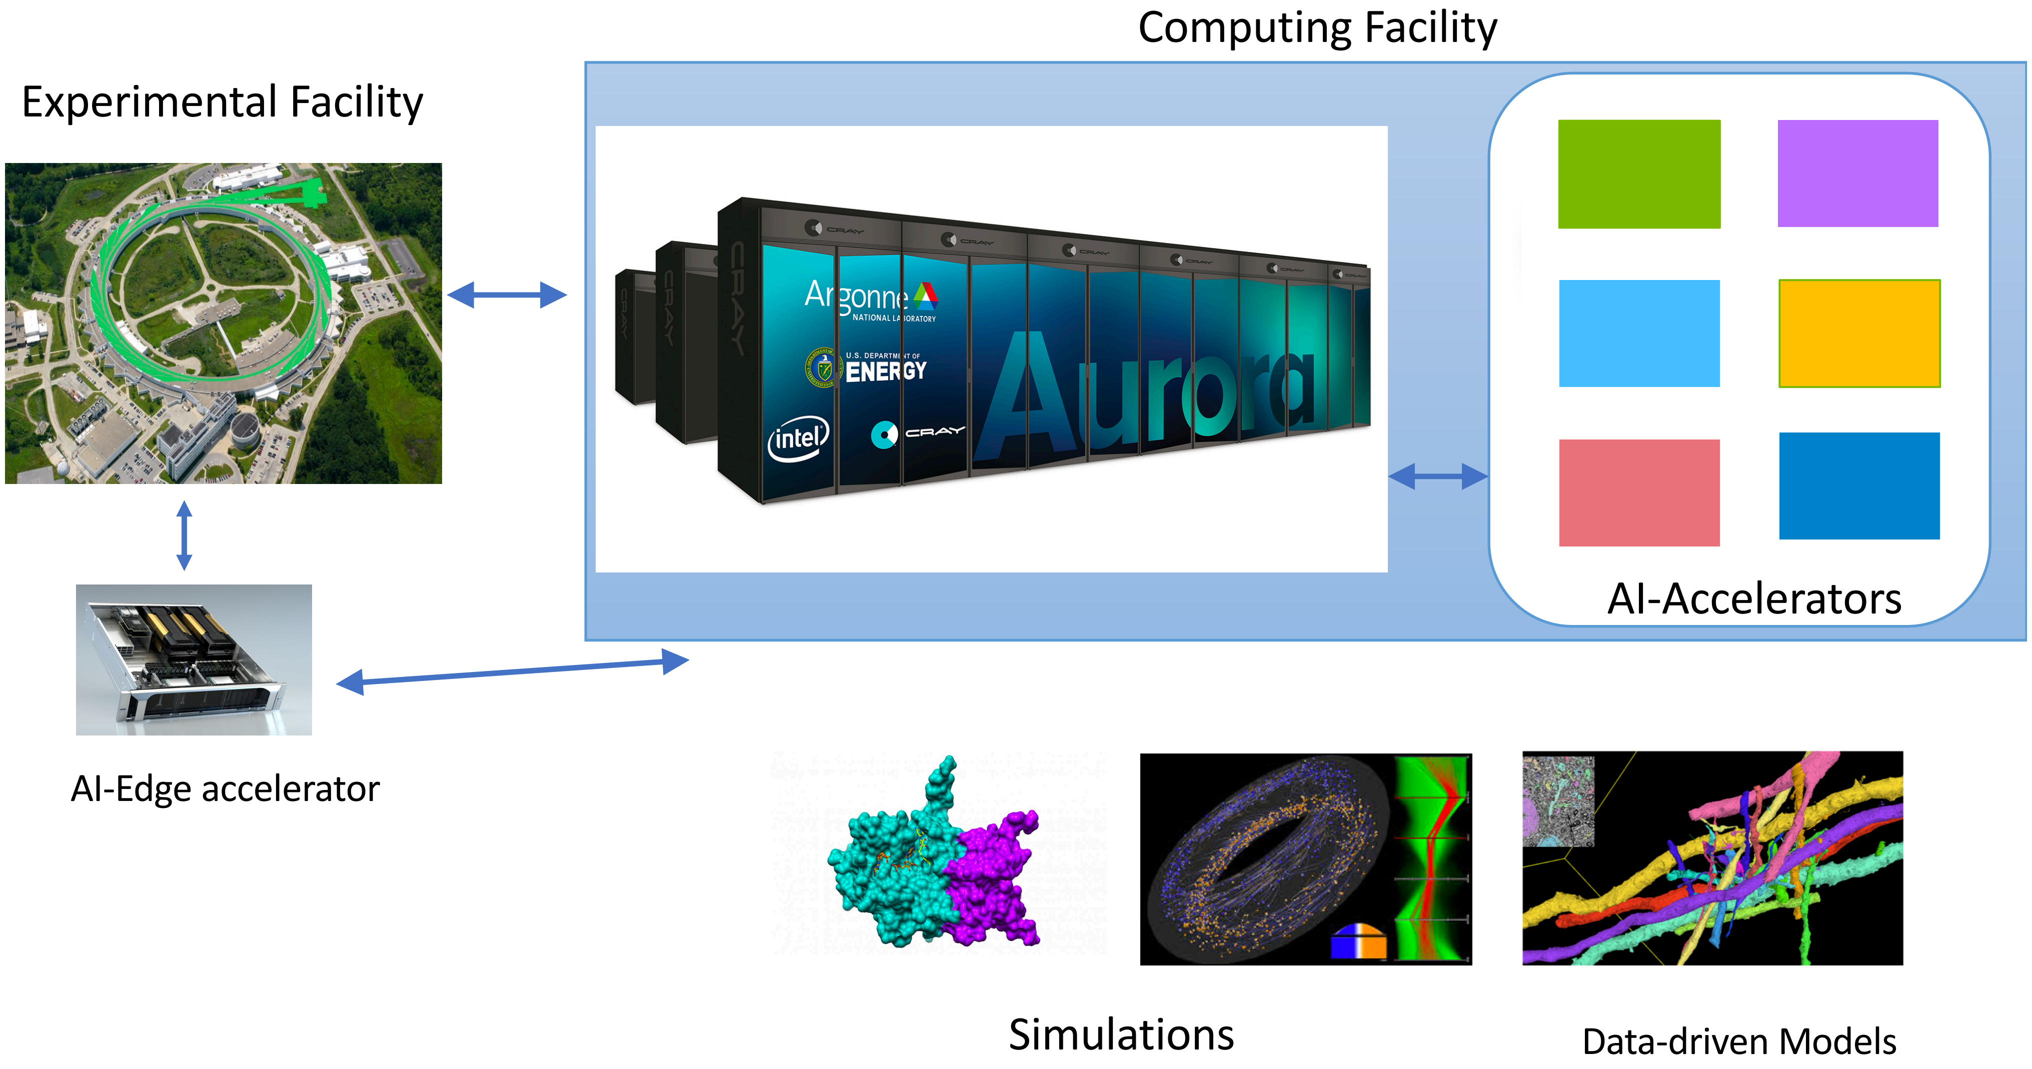

AI PATHFINDING

Goals of ALCF AI Activities at Argonne

Accelerate science by effective coupling of AI-systems, exascale supercomputers and experimental facilities

1. Maturity of software and hardware for science
2. Ability to scale hardware and integrate with facility
3. Application at scale to science

ALCF AI Testbeds

https://www.alcf.anl.gov/alcf-ai-testbed

Cerebras (CS-2)

SambaNova

Graphcore

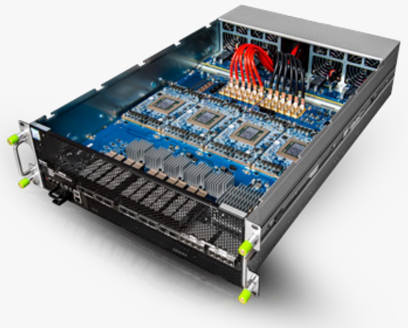

Habana

Groq

- Infrastructure of next-generation machines with hardware accelerators customized for artificial intelligence (AI) applications.
- Provide a platform to evaluate usability and performance of machine learning based HPC applications running on these accelerators.
- The goal is to better understand how to integrate AI accelerators with ALCF's existing and upcoming supercomputers to accelerate science insights

Dataflow Architectures

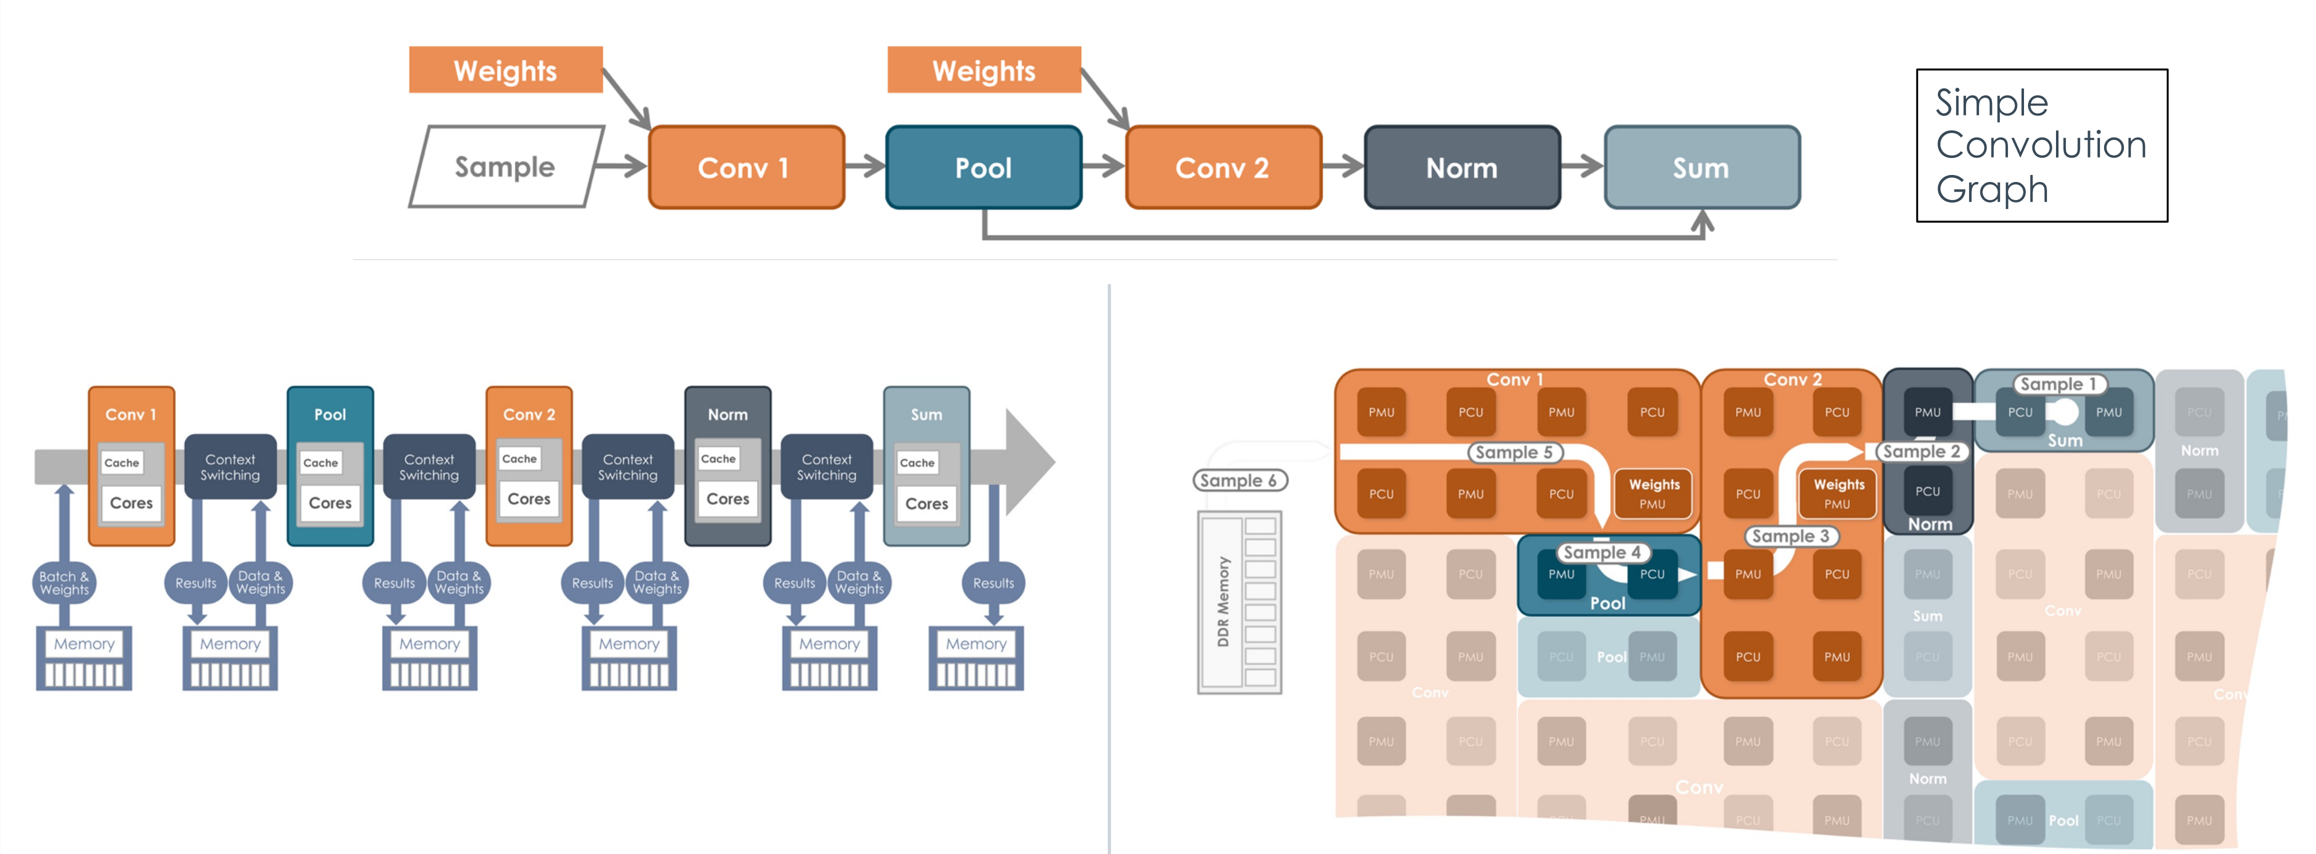

GPU accelerators: Each kernel is launched onto the device and bottlenecks include memory bandwidth and kernel-launch latencies

Dataflow: Kernels are spatially mapped onto the accelerator and data flows on-chip between them reducing memory traffic

Image Courtesy: Sumti Jairath, SambaNova

	Cerebras CS2	SambaNova Cardinal SN10	Groq GroqCard	GraphCore GC200 IPU	Habana Gaudi1	NVIDIA A100
Compute Units	850,000 Cores	640 PCUs	5120 vector ALUs	1472 IPUs	8 TPC + GEMM engine	6912 Cuda Cores
On-Chip Memory	40 GB	>300MB	230MB	900MB	24 MB	192KB L1 40MB L2
Process	7nm	7nm	14nm	7nm	7nm	7nm
System Size	2 Nodes	2 nodes (8 cards per node)	4 nodes (8 cards per node)	1 node (8 cards per node)	2 nodes (8 cards per node)	Several systems
Estimated Performance of a card (TFlops)	>5780 (FP16)	>300 (BF16)	>188 (FP16)	>250 (FP16)	>150 (FP16)	312 (FP16), 156 (FP32)
Software Stack Support	Tensorflow, Pytorch	SambaFlow, Pytorch	GroqAPI, ONNX	Tensorflow, Pytorch, PopArt	Synapse AI, TensorFlow and PyTorch	Tensorflow, Pytorch, etc
Interconnect	Ethernet-based	Infiniband	RealScale TM	IPU Link	Ethernet-based	NVLink

AI FOR SCIENCE APPLICATIONS ON AI TESTBED

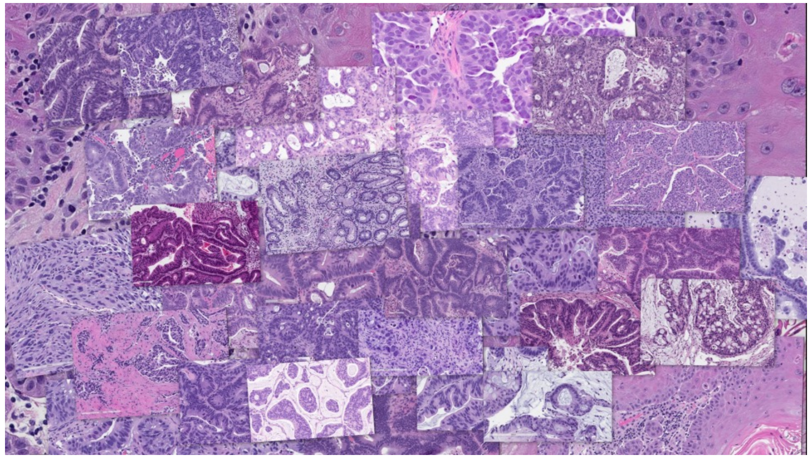

Cancer drug response prediction

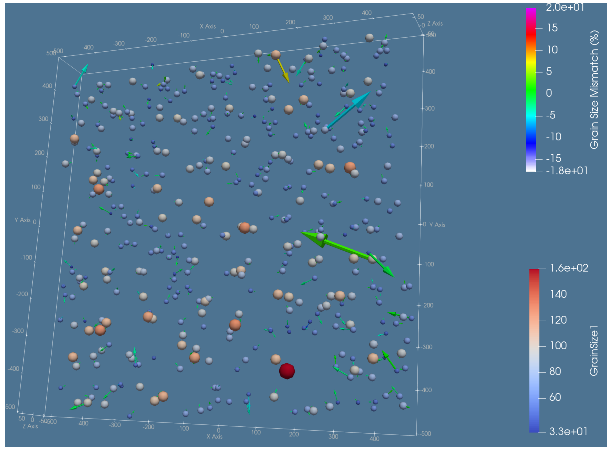

Imaging Sciences-Braggs Peak

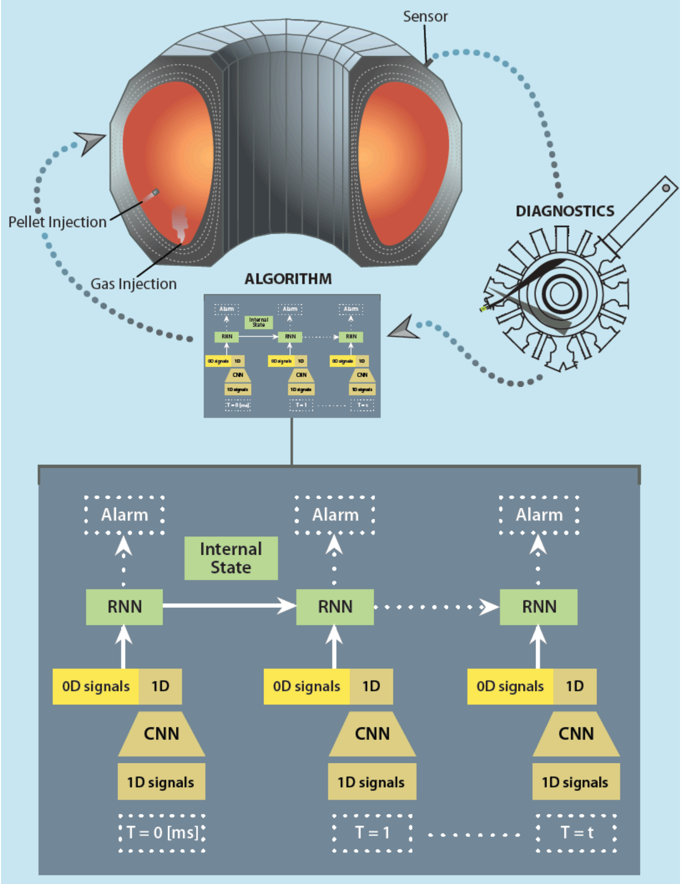

Tokomak Fusion Reactor operations

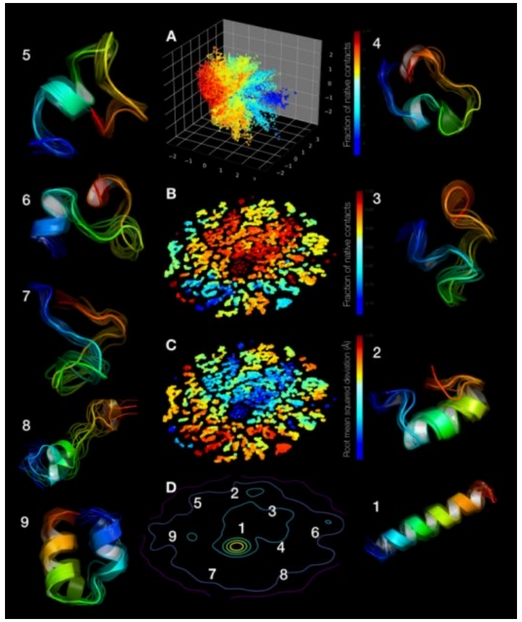

Protein-folding(Image: NCI)

and more..

Argonne NATIONAL LABORATORY

AI-enabled bridging of cryo-EM observables with atomistic fluctuations

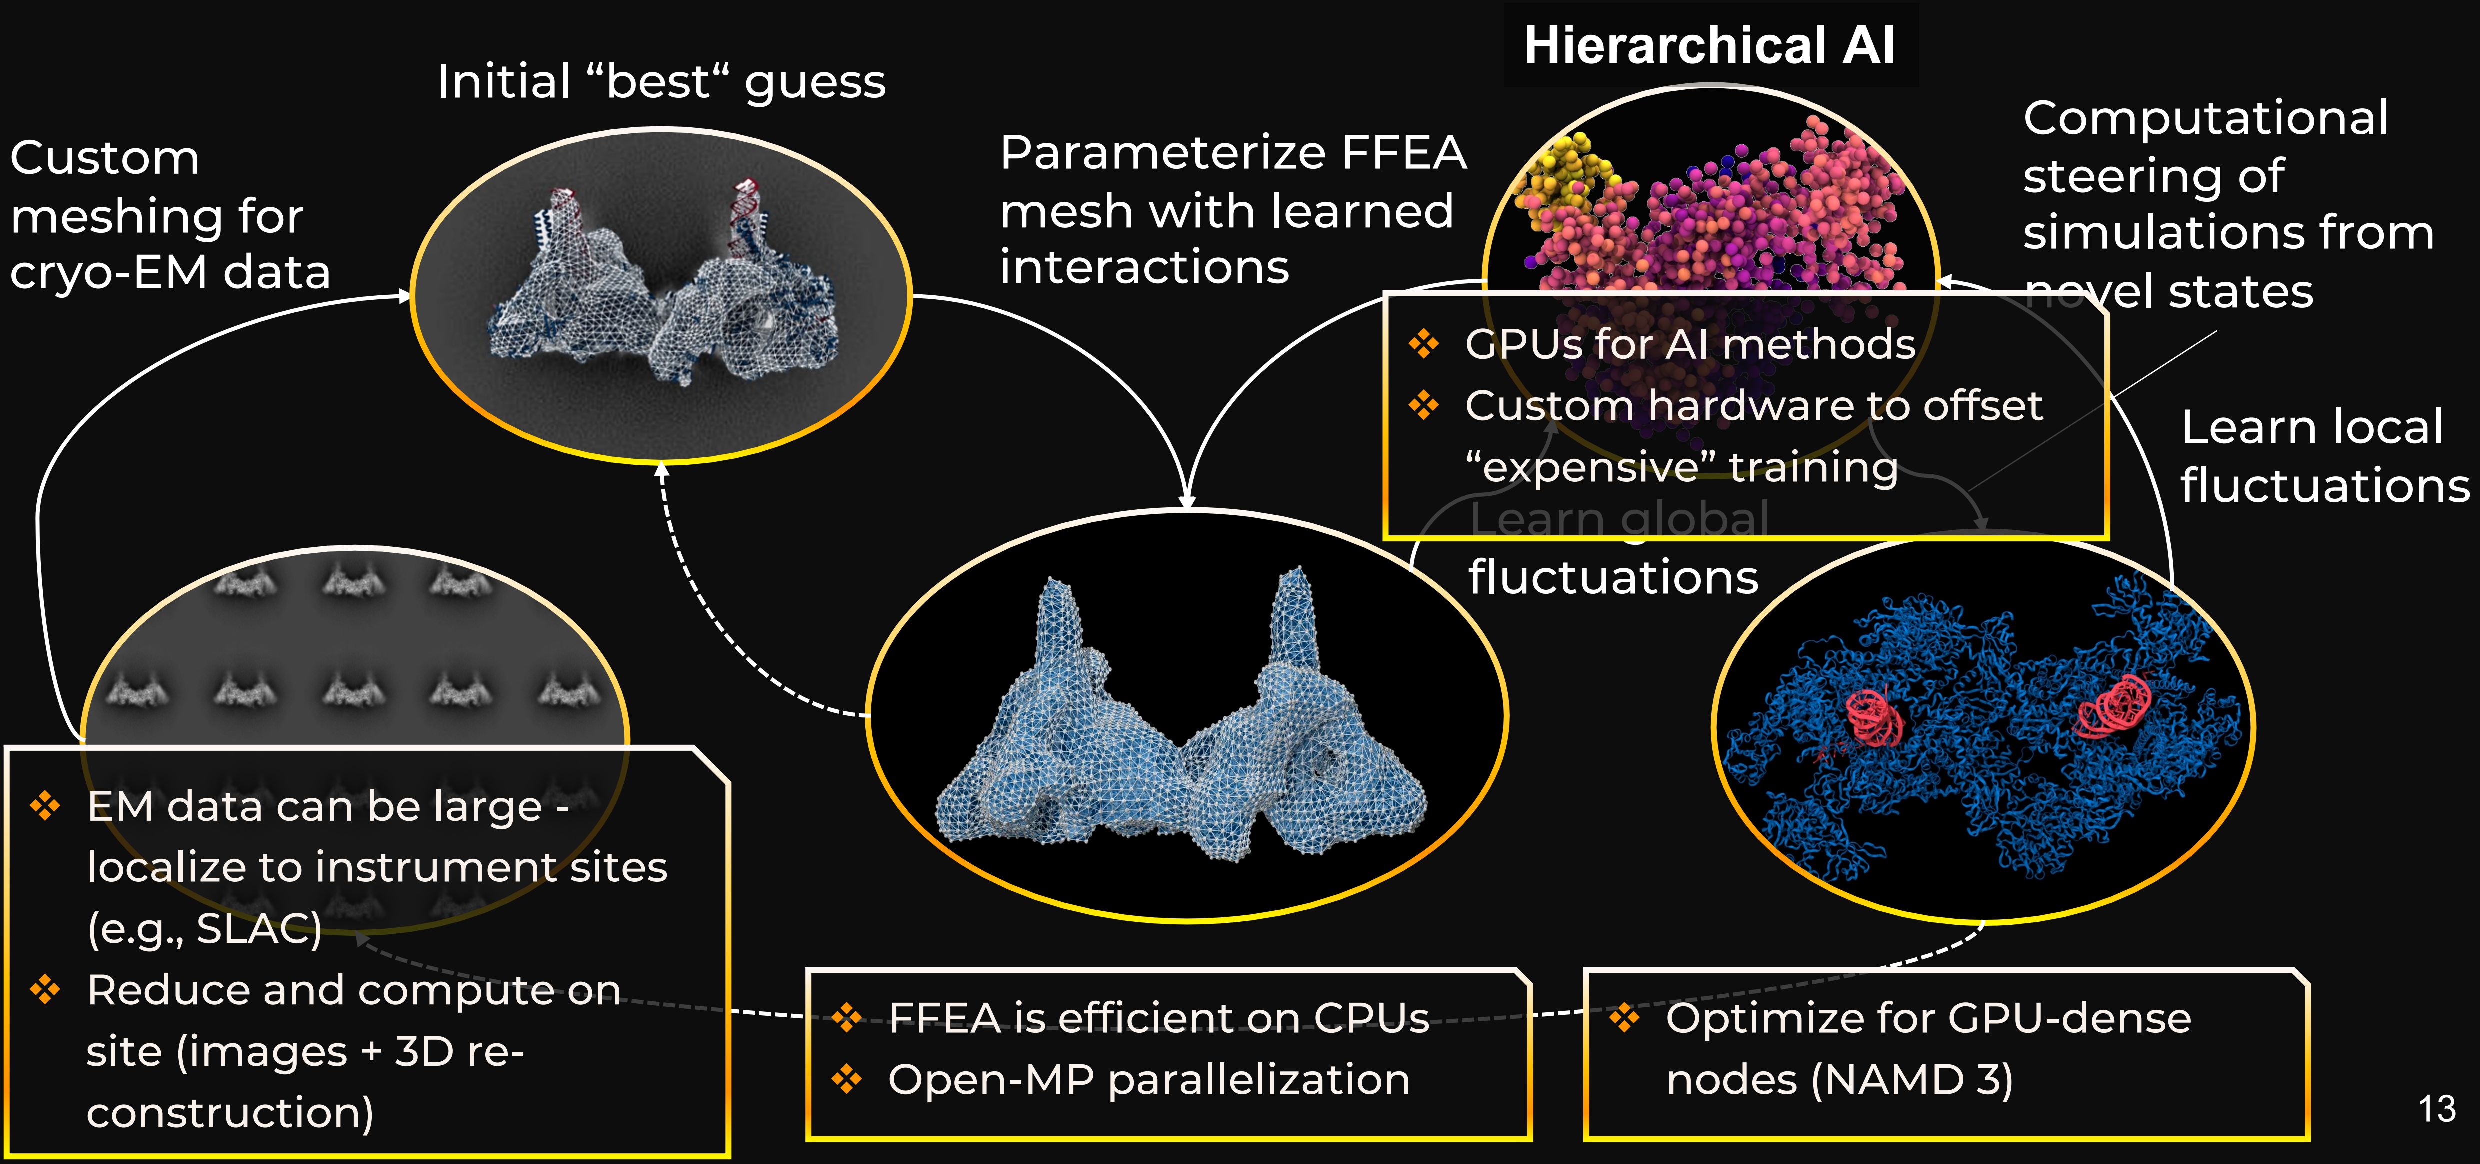

COVID-19 CVAE Training on Summit and Cerebras CS-2

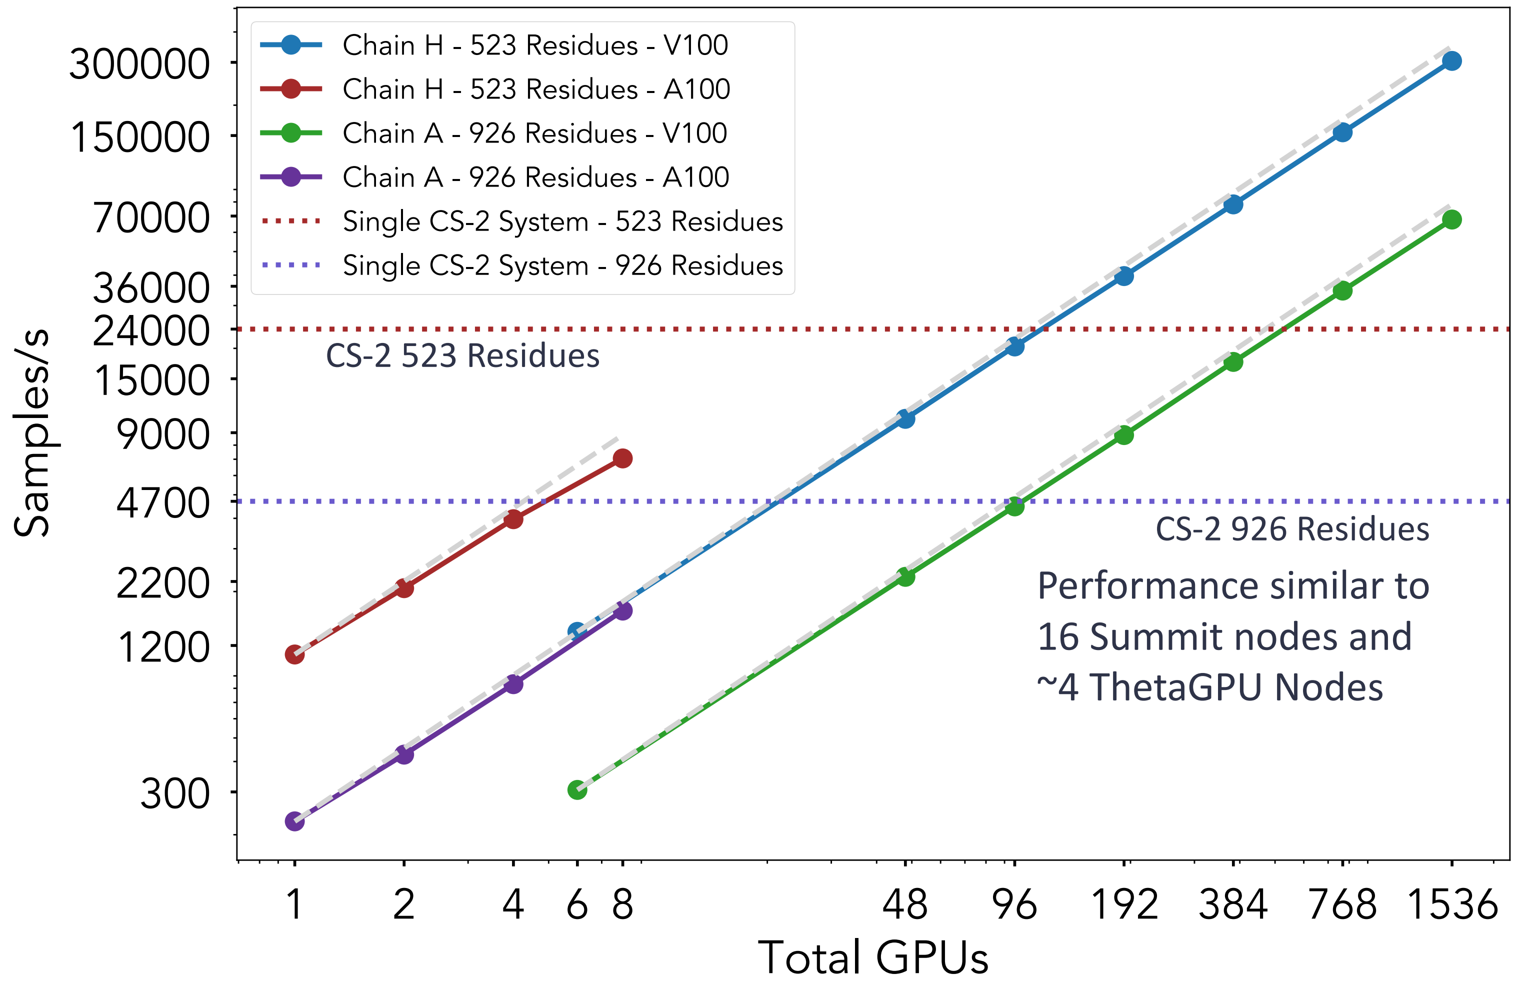

- Single CS-2 delivers performance of over 100 GPUs on CVAE
- Results are for **out-of-the-box performance** based on model config not optimized for CS-2.

Performance	523 X 523	926 X 926
Throughput (samples/sec)		
1x CS-2 System	24,000	4700
1x V100 GPU	228	51
1x A100 GPU	~1100	~150
Speedup (CS2 vs. GPU)		
1 x V100 GPU	113x	101x
1 x A100 GPU	~22X	~32X

Intelligent Resolution: Integrating Cryo-EM with AI-driven Multi-resolution Simulations to Observe the SARS-CoV-2 Replication-Transcription Machinery in Action, SC21 COVID19 Gordon Bell Finalist, In IJHPCA 2022
https://www.biorxiv.org/content/10.1101/2021.10.09.463779v1.full.pdf

COSMIC TAGGER ON SAMBANOVA DATASCALE

Goal:

Image segmentation task for liquid argon time projection chamber (LArTPC) detectors in Neutrino Physics experiments to classify each input pixel into one of three classes – Cosmic, Muon, or Background

Challenges:

Models and acquired images are limited by the size one can fit on current systems. These are expected to grow with future experiments

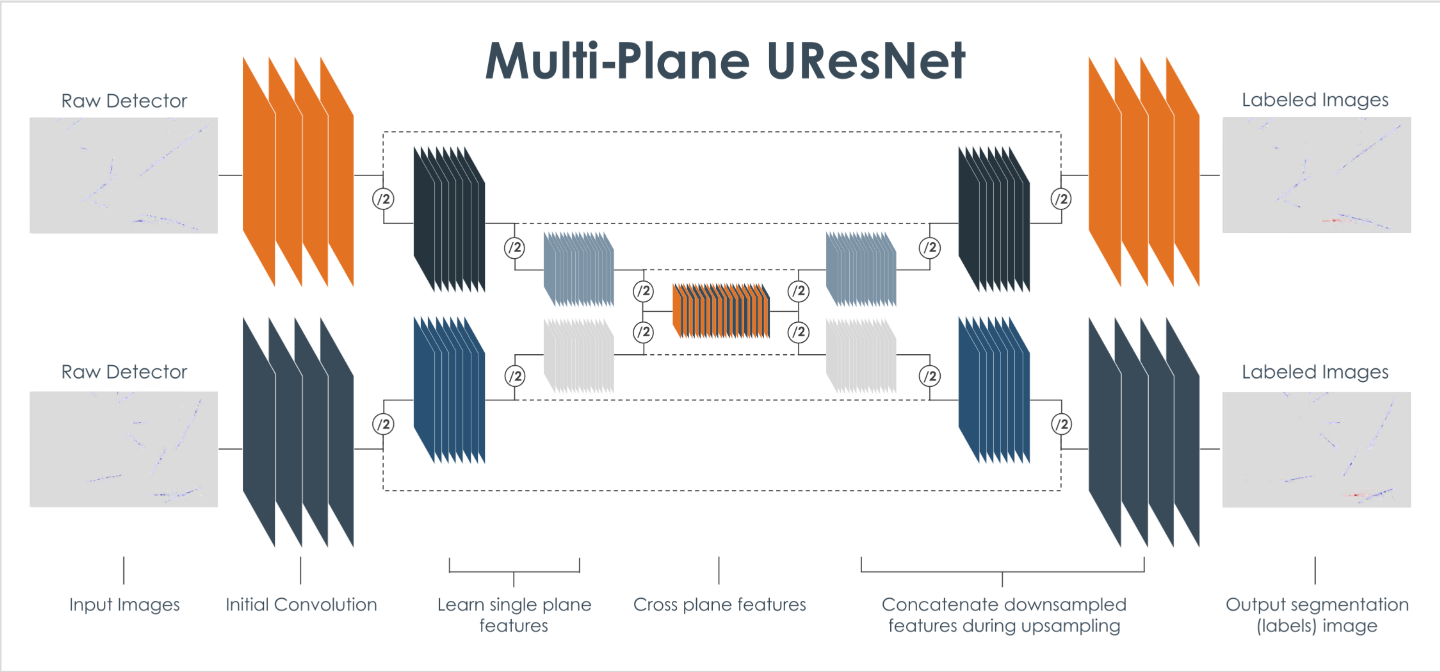

COSMIC TAGGER ON SAMBANOVA DATASCALE

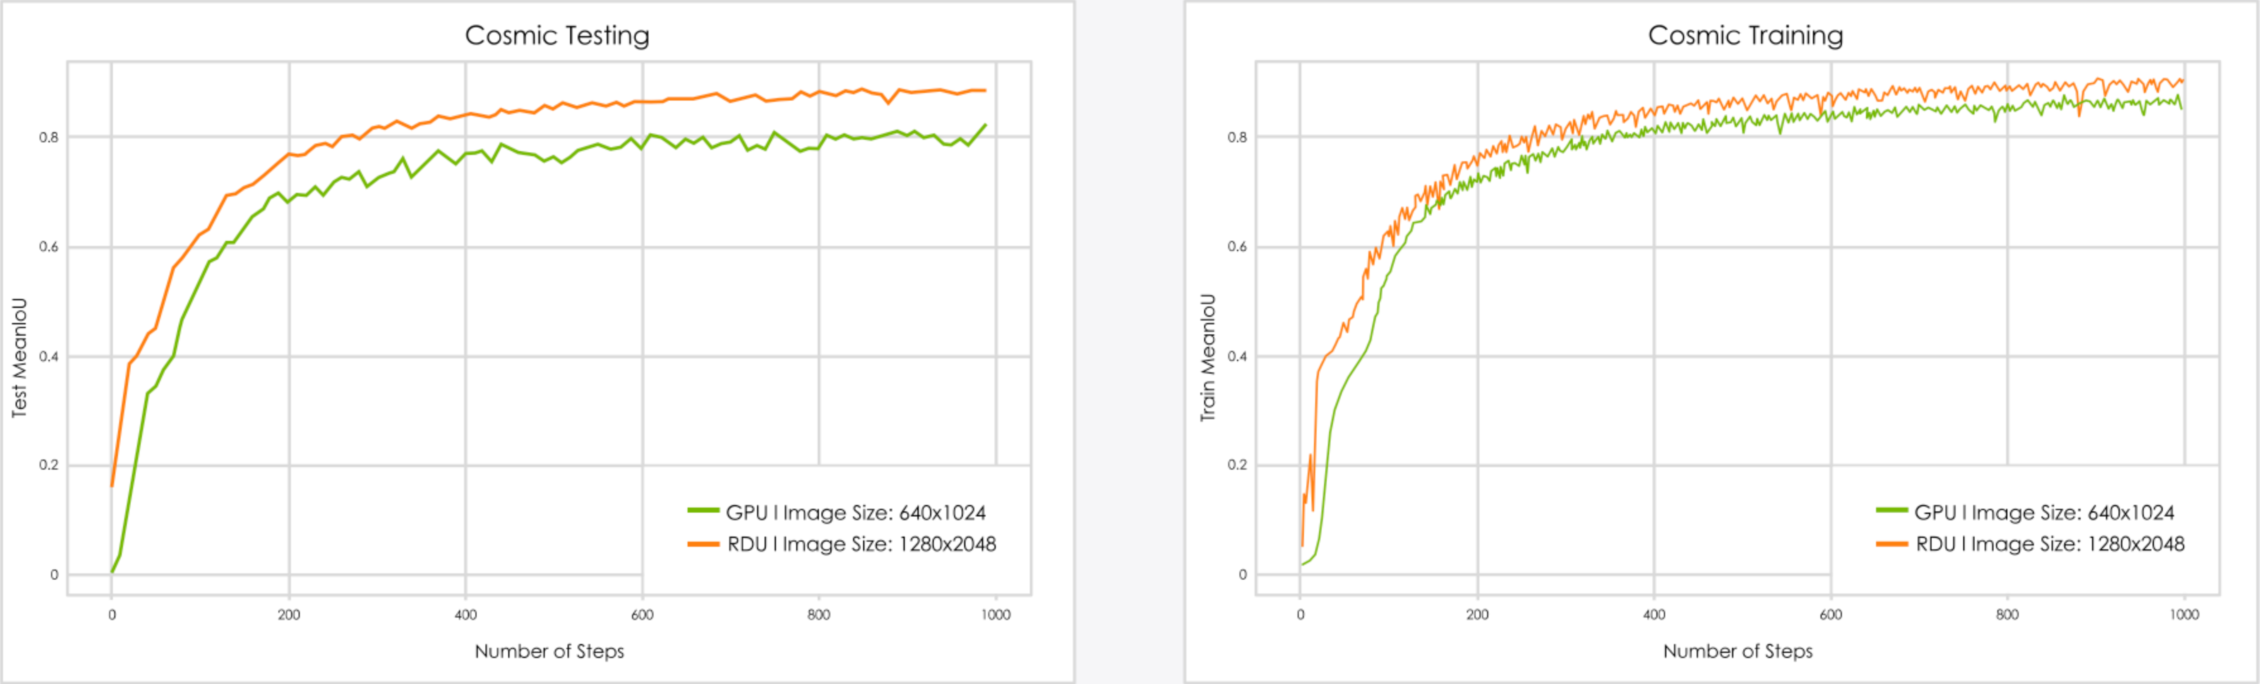

SambaNova RDUs able to accommodate larger image sizes and achieve higher accuracy

M. Emani et al., "Accelerating Scientific Applications With SambaNova Reconfigurable Dataflow Architecture," in Computing in Science & Engineering, vol. 23, no. 2, pp. 114-119, 1 March-April 2021, doi: 10.1109/MCSE.2021.3057203.

Early Experience with Inference on Groq

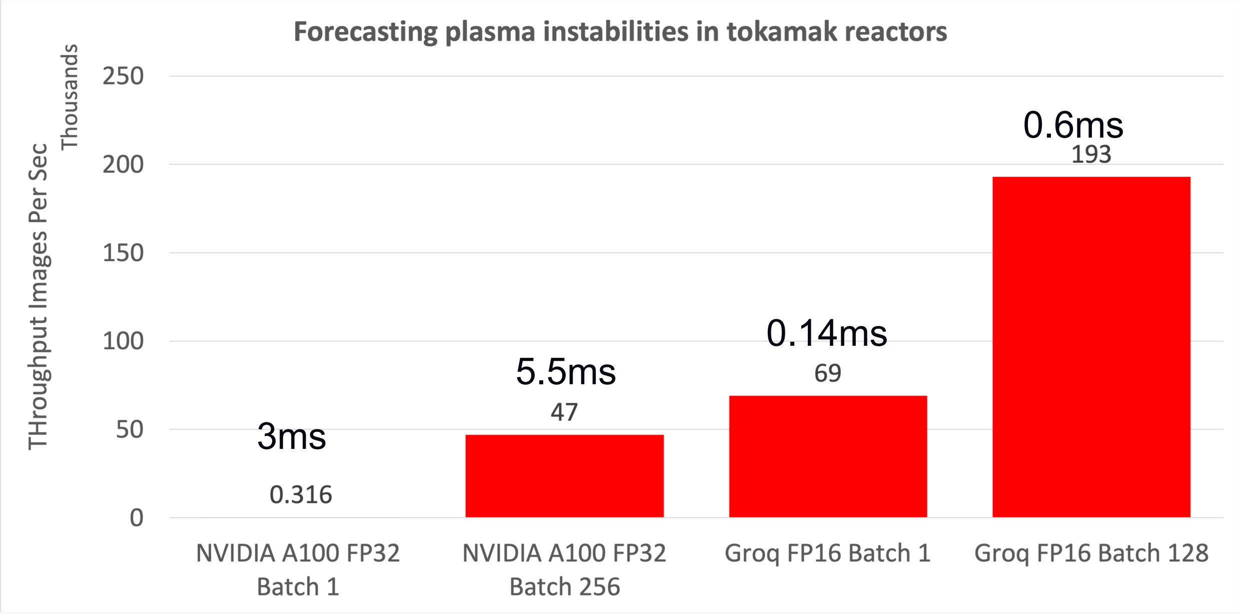

Forecasting Plasma Instability in Tokamak

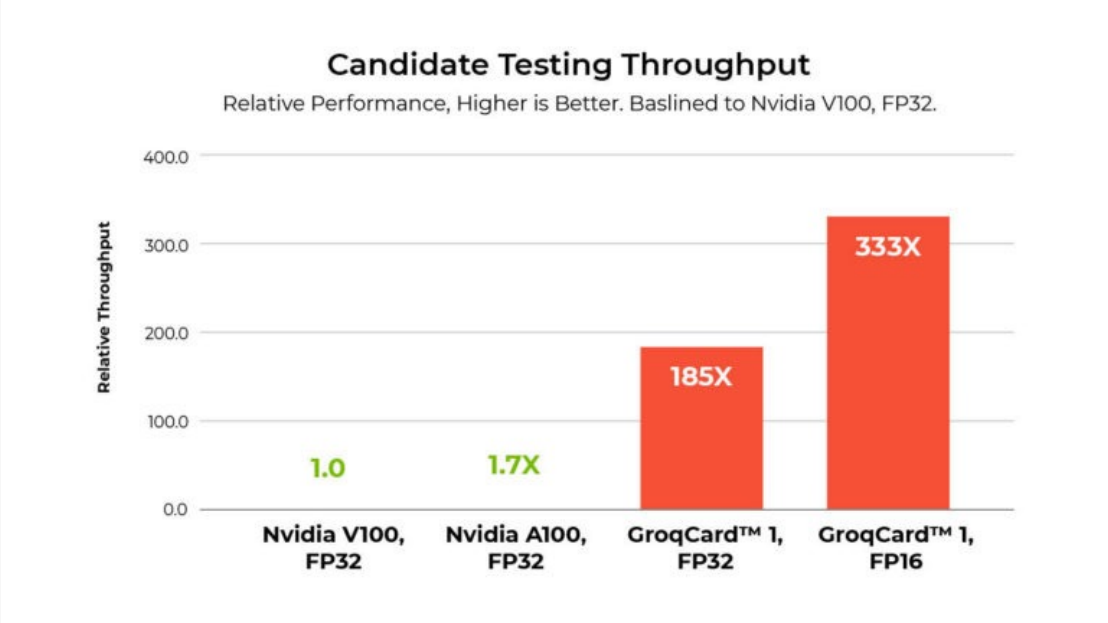

COVID19 Candidate drug molecule screening

Promising results using GroqChip for science Inference use-cases with respect to latency and throughput in comparison to GPUs

Argonne NATIONAL LABORATORY

Fast X-Ray Bragg Peak Analysis

Goal: Enable rapid analysis and real-time feedback during an in-situ experiment with complex detector technologies

Proposed Approach: Deep learning-based method, BragNN, for massive extraction of precise Bragg peak locations from far-field high energy diffraction microscopy data. BragNN has achieved 200X improvement over conventional pseudo-Voight profiling

Challenges: Model training capability is limited by the hardware

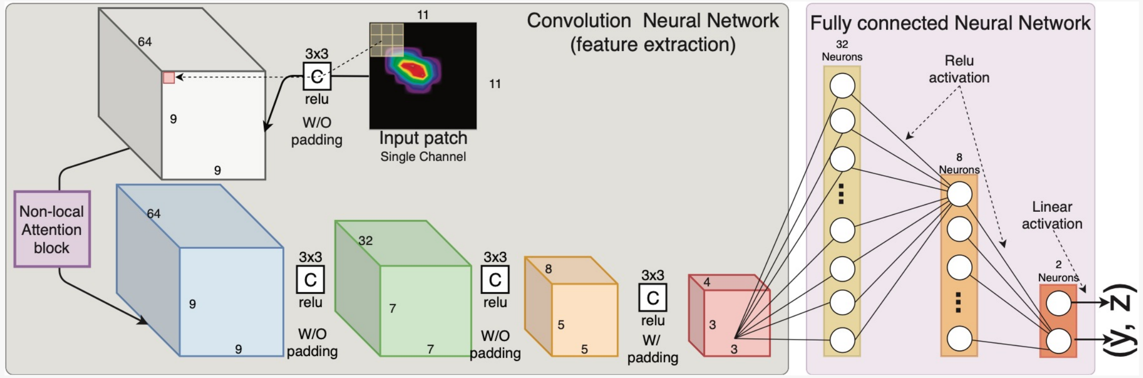

Application of the BraggNN deep neural network to an input patch yields a peak center position (y, z). All convolutions are 2D of size 3 × 3, with rectifier as activation function. Each fully connected layer, except for the output layer, also has a rectifier activation function.

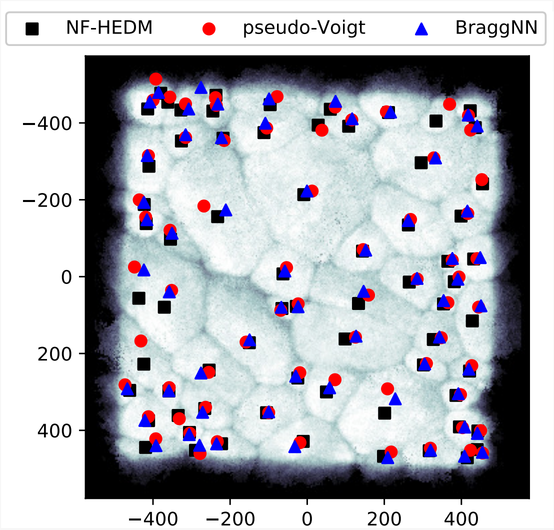

A comparison of BraggNN, pseudo-Voigt FF-HEDM and NF-HEDM. (a) Grain positions from NF-HEDM (black squares), pseudo-Voigt FF-HEDM (red circles) and BraggNN FF-HEDM (blue triangles) overlaid on NF-HEDM confidence map

Courtesy: Z. Liu et al. BraggNN: Fast X-ray Bragg Peak Analysis Using Deep Learning. International Union of Crystallography (IUCrJ), Vol. 9, No. 1, 2022

Argonne NATIONAL LABORATORY

Fast X-Ray Bragg Peak Analysis

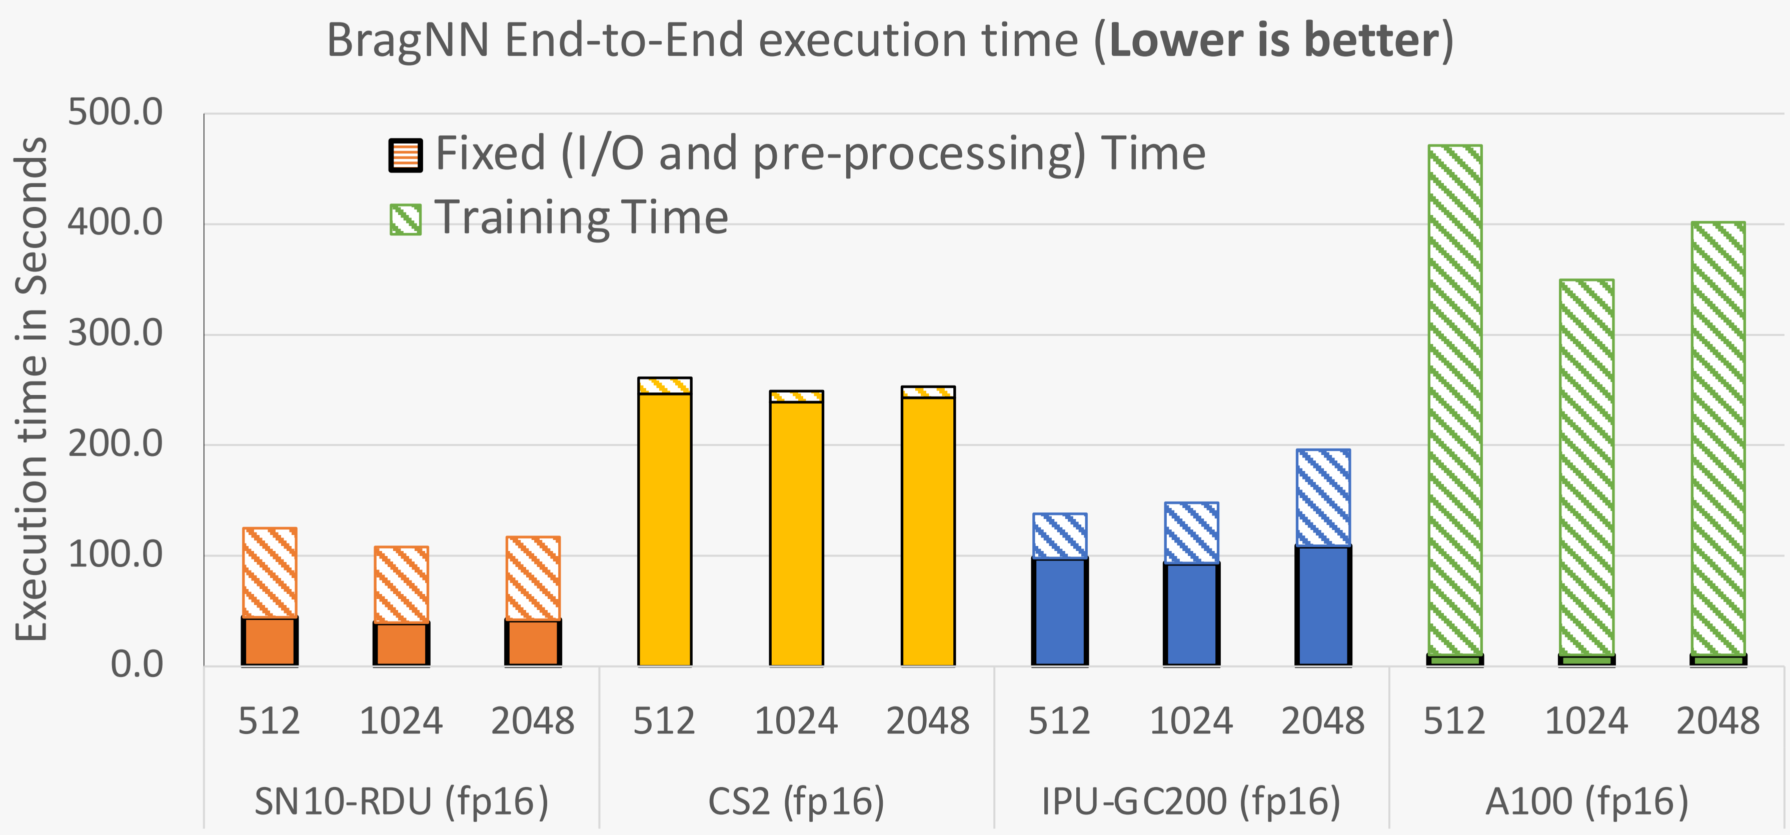

SambaNova and Graphcore achieve lowest time to solution and achieve up to 3.7X to 3.4X speedup in comparison to Nvidia A100 respectively. Cerebras achieves up to 80% improvement over A100

Argonne NATIONAL LABORATORY

Fast X-Ray Bragg Peak Analysis

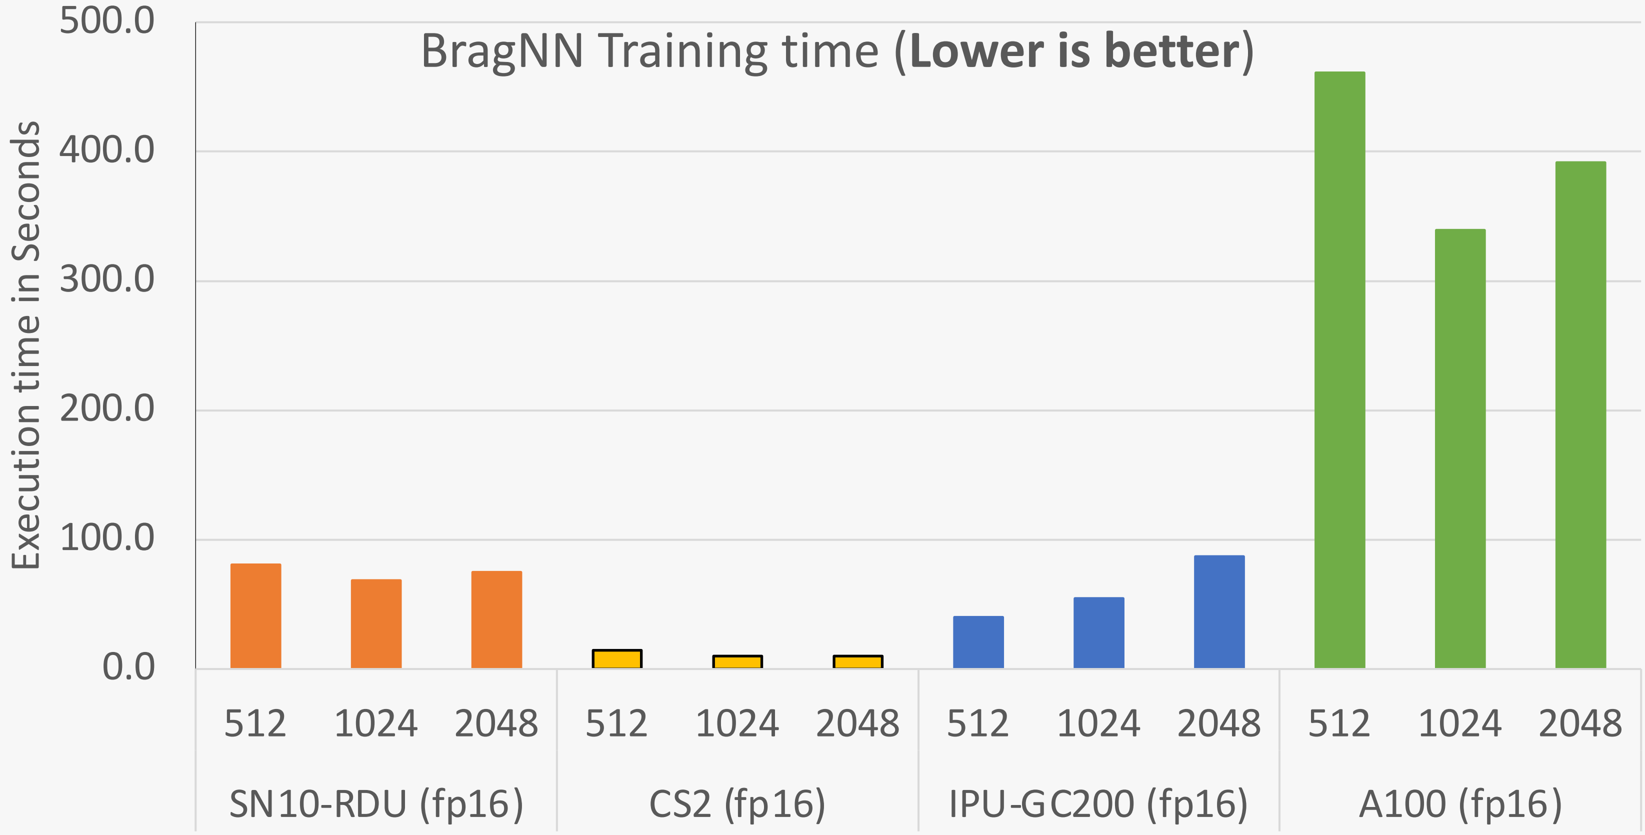

For training time, we ignore the data loading and pre-processing time (Fixed cost time).
Cerebras CS2 achieves up to 33X improvement over A100 while SN and Graphcore achieve up to 6-11X improvement over A100 respectively for training.
Cerebras performance includes use of multi-replica optimization and similar optimizations need to evaluated on other systems

Scaling UNet-2D Training

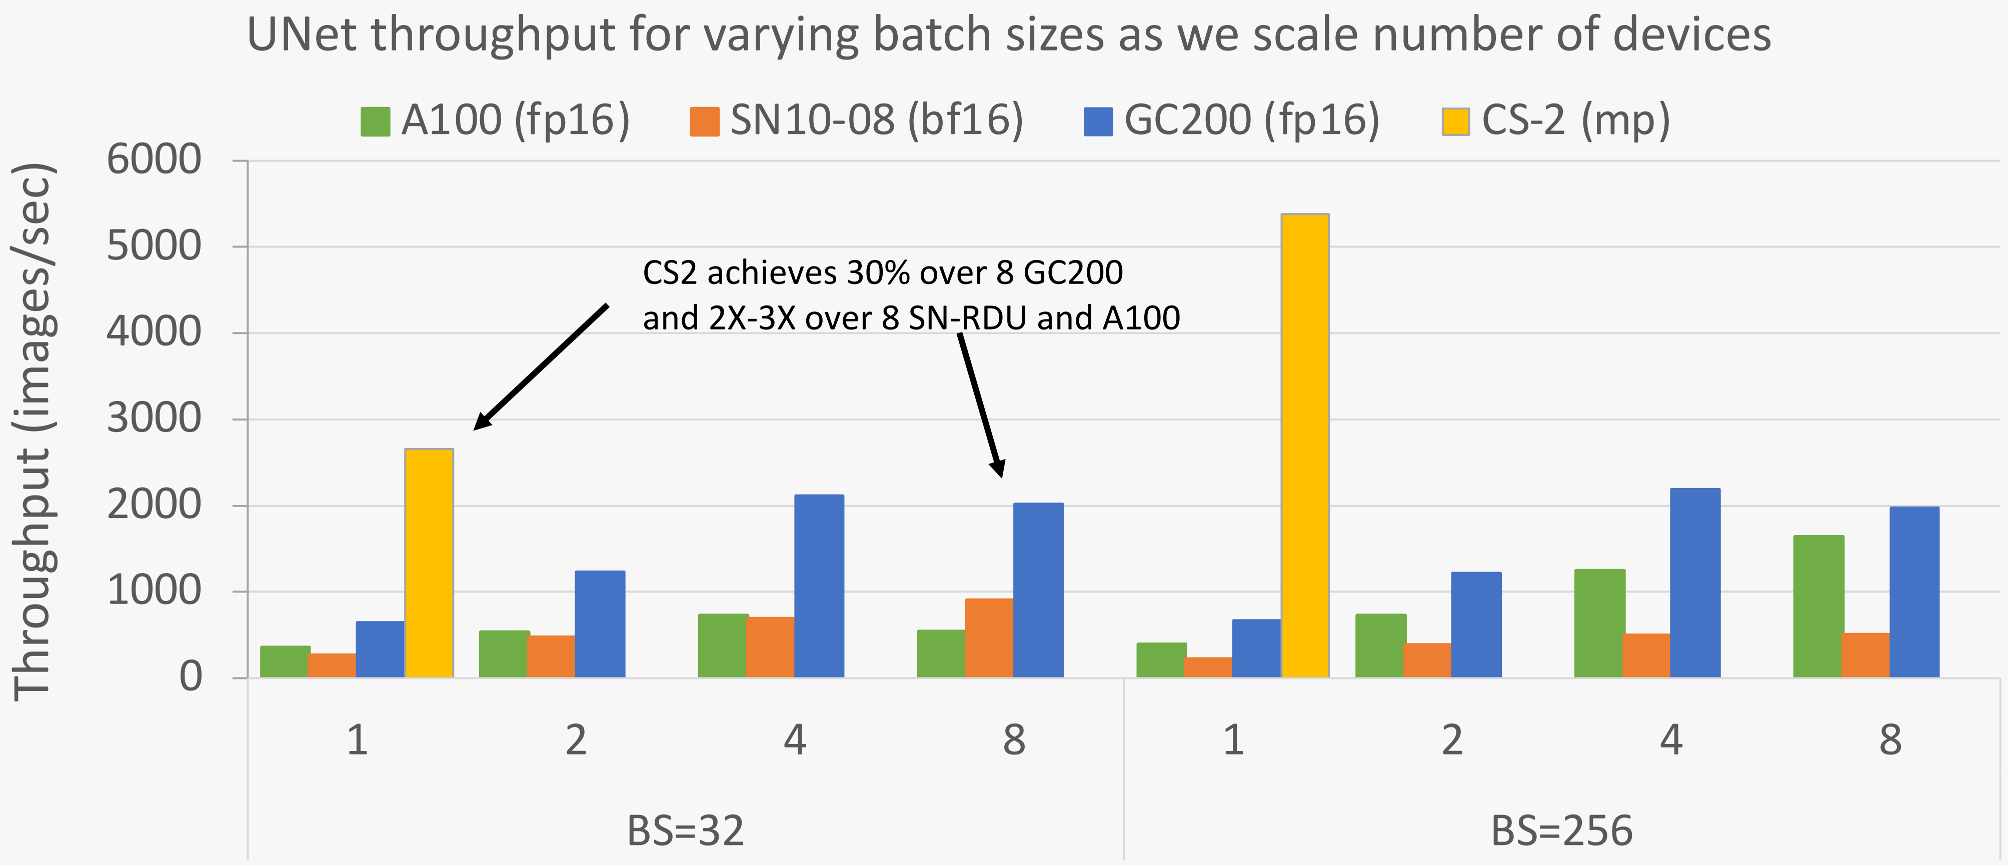

Scale to 1, 2, 4 and 8 devices with batch sizes (BS) 32 and 256, for image size of 256x256

Note: Graphcore performance includes an optimization to prefetch data and work is ongoing to incorporate similar optimizations for other systems

For smaller batch sizes (32), Cerebras CS2 achieve up to a 30% improvement over 8 GC200 devices, over 2X and 3X in comparison to using eight SN 10-RDU and A100. For larger batch sizes, we see similar trends though with improved A100 performance.

Communication Performance - All Reduce

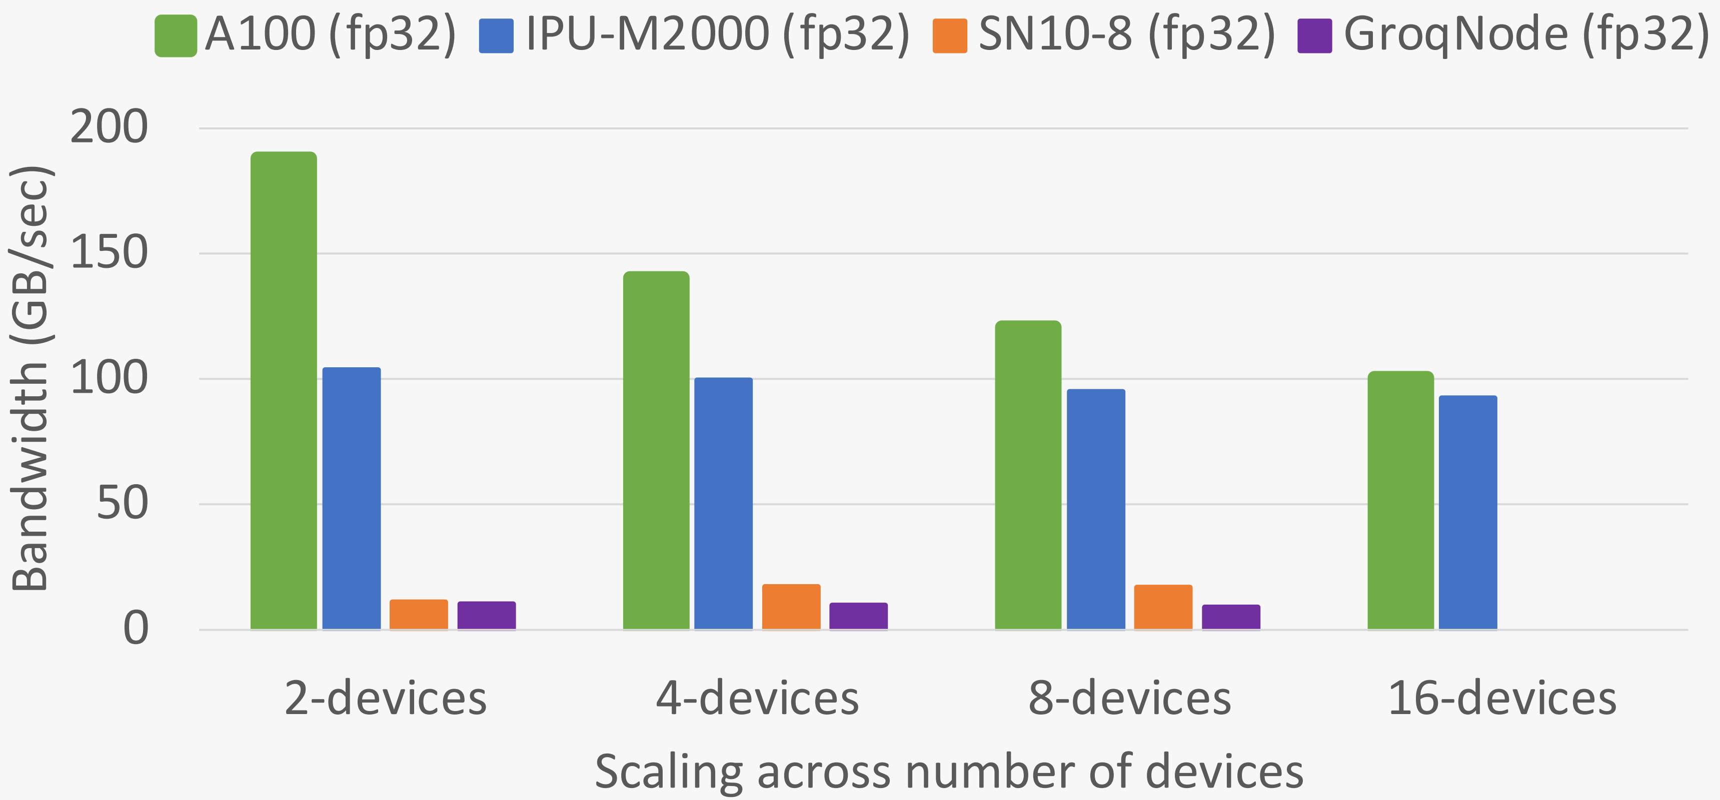

DeepBench and OSU MPI Benchmarks used for the all_reduce communication evaluation and we scale the number of devices to 16. We use only 8 devices for Groq and SambaNova

We observe that Nvidia DGX3 achieves higher All Reduce performance in comparison to other AI systems

AI Testbed Community Engagement

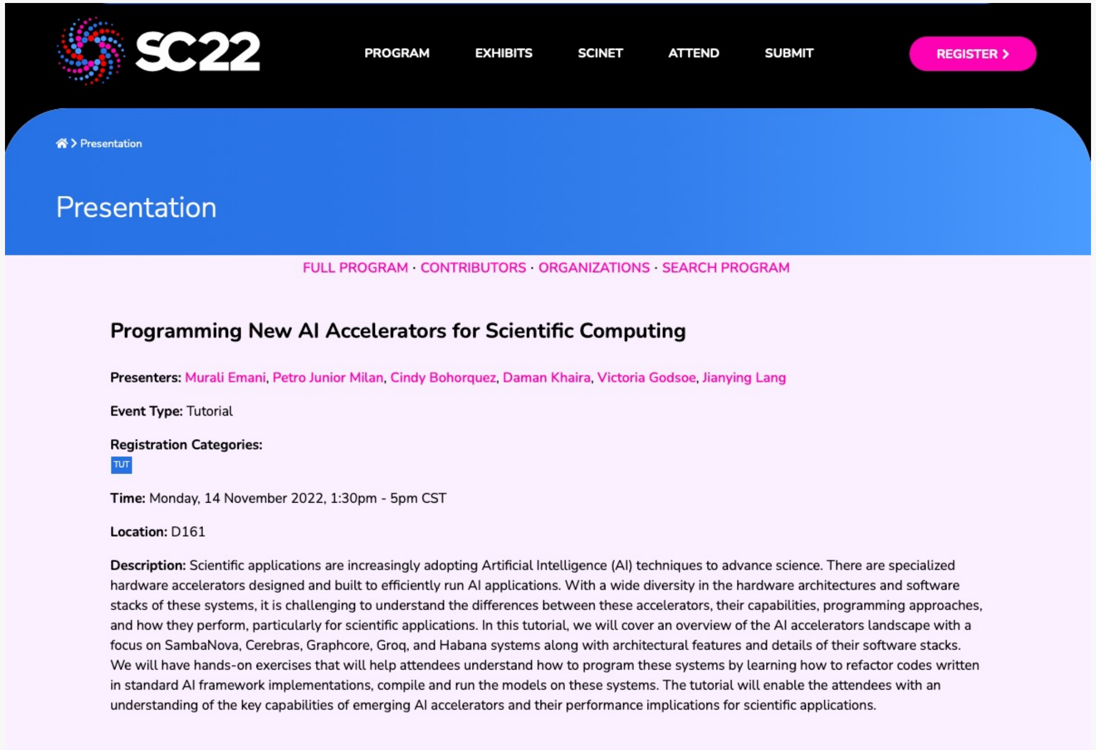

SC22 PROGRAM EXHIBITS SCINET ATTEND SUBMIT REGISTER >

> Presentation

Presentation

FULL PROGRAM · CONTRIBUTORS · ORGANIZATIONS · SEARCH PROGRAM

Programming New AI Accelerators for Scientific Computing

Presenters: Murali Emani, Petro Junior Milan, Cindy Bohorquez, Daman Khaira, Victoria Godsoe, Jianying Lang

Event Type: Tutorial

Registration Categories:
TUT

Time: Monday, 14 November 2022, 1:30pm - 5pm CST

Location: D161

Description: Scientific applications are increasingly adopting Artificial Intelligence (AI) techniques to advance science. There are specialized hardware accelerators designed and built to efficiently run AI applications. With a wide diversity in the hardware architectures and software stacks of these systems, it is challenging to understand the differences between these accelerators, their capabilities, programming approaches, and how they perform, particularly for scientific applications. In this tutorial, we will cover an overview of the AI accelerators landscape with a focus on SambaNova, Cerebras, Graphcore, Groq, and Habana systems along with architectural features and details of their software stacks. We will have hands-on exercises that will help attendees understand how to program these systems by learning how to refactor codes written in standard AI framework implementations, compile and run the models on these systems. The tutorial will enable the attendees with an understanding of the key capabilities of emerging AI accelerators and their performance implications for scientific applications.

- SambaNova AI training workshop – July 19-20, 2022
- ATPESC H/W Architecture Day on August 1, 2022 will cover five AI accelerators
- ALCF AI for Science training series for students in the fall will include the AI testbed
- Cerebras CS-2 training workshop planned for August 2022

SC'22 Tutorial on Programming AI accelerators for Scientific Computing ***in collaboration with Cerebras, Intel Habana, Graphcore, Groq and SambaNova* accepted**

AI Testbed Expeditions – Argonne LDRD Program

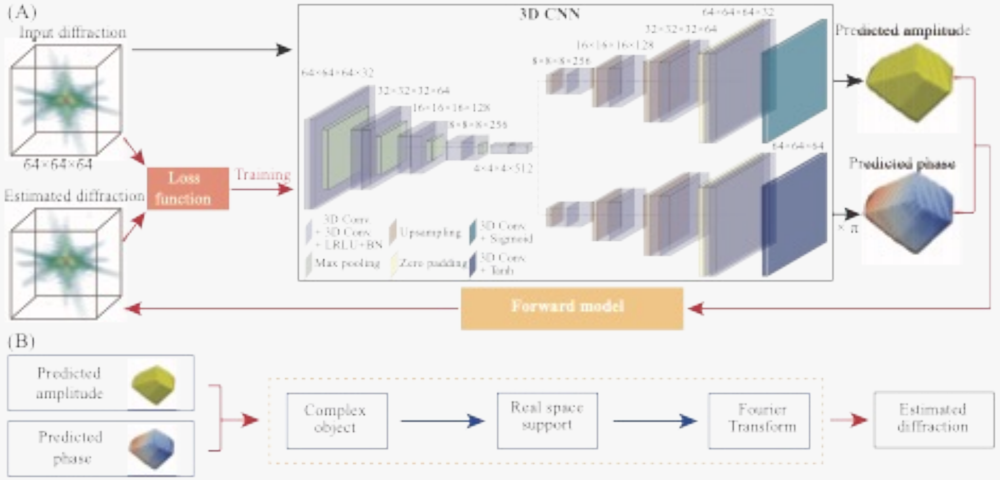

AutoPhaseNN for coherent diffraction imaging
Courtesy: Yudong Yau, Argonne

- Expeditions projects target Argonne-related AI/ML science, autonomous discovery, or computational science problem areas; make use of this new testbed; and, ideally, promote collaboration across domains
- Supported 18 projects in 2021 (~1 month effort) and supporting 11 projects in 2022 (~2 months effort)
- AutophaseNN achieved a 39% improvement in training time on SambaNova over A100

Table of Contents

SambaNova

Groq and GraphCore

Cerebras

LDRD 2021 Report

Getting Started on ALCF AI Testbed:

Apply for a Director’s Discretionary (DD) Allocation Award

Director’s Discretionary (DD) awards support various project objectives from scaling code to preparing for future computing competition to production scientific computing in support of strategic partnerships.

Cerebras CS-2 and SambaNova Datascale are available for allocations

[Allocation Request Form]

[AI Testbed User Guide]

Ongoing Efforts

- System upgrade plans include a two-rack SambaNova Datascale system (from ½ rack), a Graphcore Bow-200 (3^{rd} generation) Pod64 rack, and rack-scale Groq system
- Work with AI vendors to facilitate AI for Science applications, including support for large-language models.
- Evaluate new AI accelerators offerings and incorporate promising solutions as part of the testbed
- Integrate AI testbed systems with the PBSPro scheduler to facilitate job scheduling across the accelerators in the testbed and improve user experience
- Evaluate traditional HPC on AI Accelerators
- Understand how to integrate AI accelerators with ALCF's existing and upcoming supercomputers to accelerate science insights

Observations, Challenges and Insights

Significant speedup achieved for a wide-gamut of scientific ML applications

- Easier to deal with larger resolution data and to scale to multi-chip systems

Room for improvement exists

- Porting efforts and compilation times
- Coverage of DL frameworks and support for performance analysis tools and debuggers

Good progress made in integration of AI accelerators, in production, at a national user facility and significant more work is needed for effective coupling

Training and Outreach is critical to educate users to effectively use AI systems

Close collaboration with vendors is necessary to realize the vision of AI for science

Recent Publications

- **Intelligent Resolution: Integrating Cryo-EM with AI-driven Multi-resolution Simulations to Observe the SARS-CoV-2 Replication-Transcription Machinery in Action***
  Anda Trifan, Defne Gorgun, Zongyi Li, Alexander Brace, Maxim Zvyagin, Heng Ma, Austin Clyde, David Clark, Michael Salim, David Hardy,Tom Burnley, Lei Huang, John McCalpin, Murali Emani, Hyenseung Yoo, Junqi Yin, Aristeidis Tsaris, Vishal Subbiah, Tanveer Raza,Jessica Liu, Noah Trebesch, Geoffrey Wells, Venkatesh Mysore, Thomas Gibbs, James Phillips, S.Chakra Chennubhotla, Ian Foster, Rick Stevens, Anima Anandkumar, Venkatram Vishwanath, John E. Stone, Emad Tajkhorshid, Sarah A. Harris, Arvind Ramanathan, International Journal of High-Performance Computing (IJHPC'22)
  DOI: https://doi.org/10.1101/2021.10.09.463779

- **Stream-AI-MD: Streaming AI-driven Adaptive Molecular Simulations for Heterogeneous Computing Platforms**
  Alexander Brace, Michael Salim, Vishal Subbiah, Heng Ma, Murali Emani, Anda Trifa, Austin R. Clyde, Corey Adams, Thomas Uram, Hyunseung Yoo, Andrew Hock, Jessica Liu, Venkatram Vishwanath, and Arvind Ramanathan. 2021 Proceedings of the Platform for Advanced Scientific Computing Conference (PASC'21). DOI: https://doi.org/10.1145/3468267.3470578

- **Bridging Data Center AI Systems with Edge Computing for Actionable Information Retrieval**
  Zhengchun Liu, Ahsan Ali, Peter Kenesei, Antonino Miceli, Hemant Sharma, Nicholas Schwarz, Dennis Trujillo, Hyunseung Yoo, Ryan Coffee, Naoufal Layad, Jana Thayer, Ryan Herbst, Chunhong Yoon, and Ian Foster, 3rd Annual workshop on Extreme-scale Event-in-the-loop computing (XLOOP), 2021

- **Accelerating Scientific Applications With SambaNova Reconfigurable Dataflow Architecture**
  Murali Emani, Venkatram Vishwanath, Corey Adams, Michael E. Papka, Rick Stevens, Laura Florescu, Sumti Jairath, William Liu, Tejas Nama, Arvind Sujeeth, IEEE Computing in Science & Engineering 2021 DOI: 10.1109/MCSE.2021.3057203.

*** Finalist in the ACM Gordon Bell Special Prize for High Performance Computing-Based COVID-19 Research, 2021**

Thank You

- This research was funded in part and used resources of the Argonne Leadership Computing Facility (ALCF), a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.

- Murali Emani, Michael Papka, William Arnold, Bruce Wilson, Varuni Sastry, Sid Raskar, Corey Adams, Rajeev Thakur, Anthony Avarca, Arvind Ramanathan, Alex Brace, Zhengchun Liu, Hyunseung (Harry) Yoo, Ryan Aydelott, Sid Raskar, Zhen Xie, Kyle Felker, Craig Stacey, Tom Brettin, Rick Stevens, and many others have contributed to this material.

- Our current AI testbed system vendors – Cerebras, Graphcore, Groq, Intel Habana and SambaNova