# ATLAS Distributed Computing Technical Interchange Meeting

## Welcome from BNL, SBU and US ATLAS

Chris Bee (SBU), Alexei Klimentov (BNL), Torre Wenaus (BNL)

January 21, 2025

# Welcome to SBU and BNL !

- Thank you very much for coming! Many have come a long way, we appreciate it
- Keep your badge on while at BNL
- BNL and SBU have an anti-harassment policy, available on the website
- Observe speed limits! Watch for humans !
- Group photo: Today
- REDWOOD splinter meetings will be in CDS Building (Bldg 725, conference room A)

TIM :

- A discussion-oriented technical gathering
- We will keep live notes, at least for some of it, contributions appreciated
  - Mario will present TIM highlights at SW&C week

BNL Maps & Directions

Brookhaven National Laboratory

# Brookhaven Lab Today

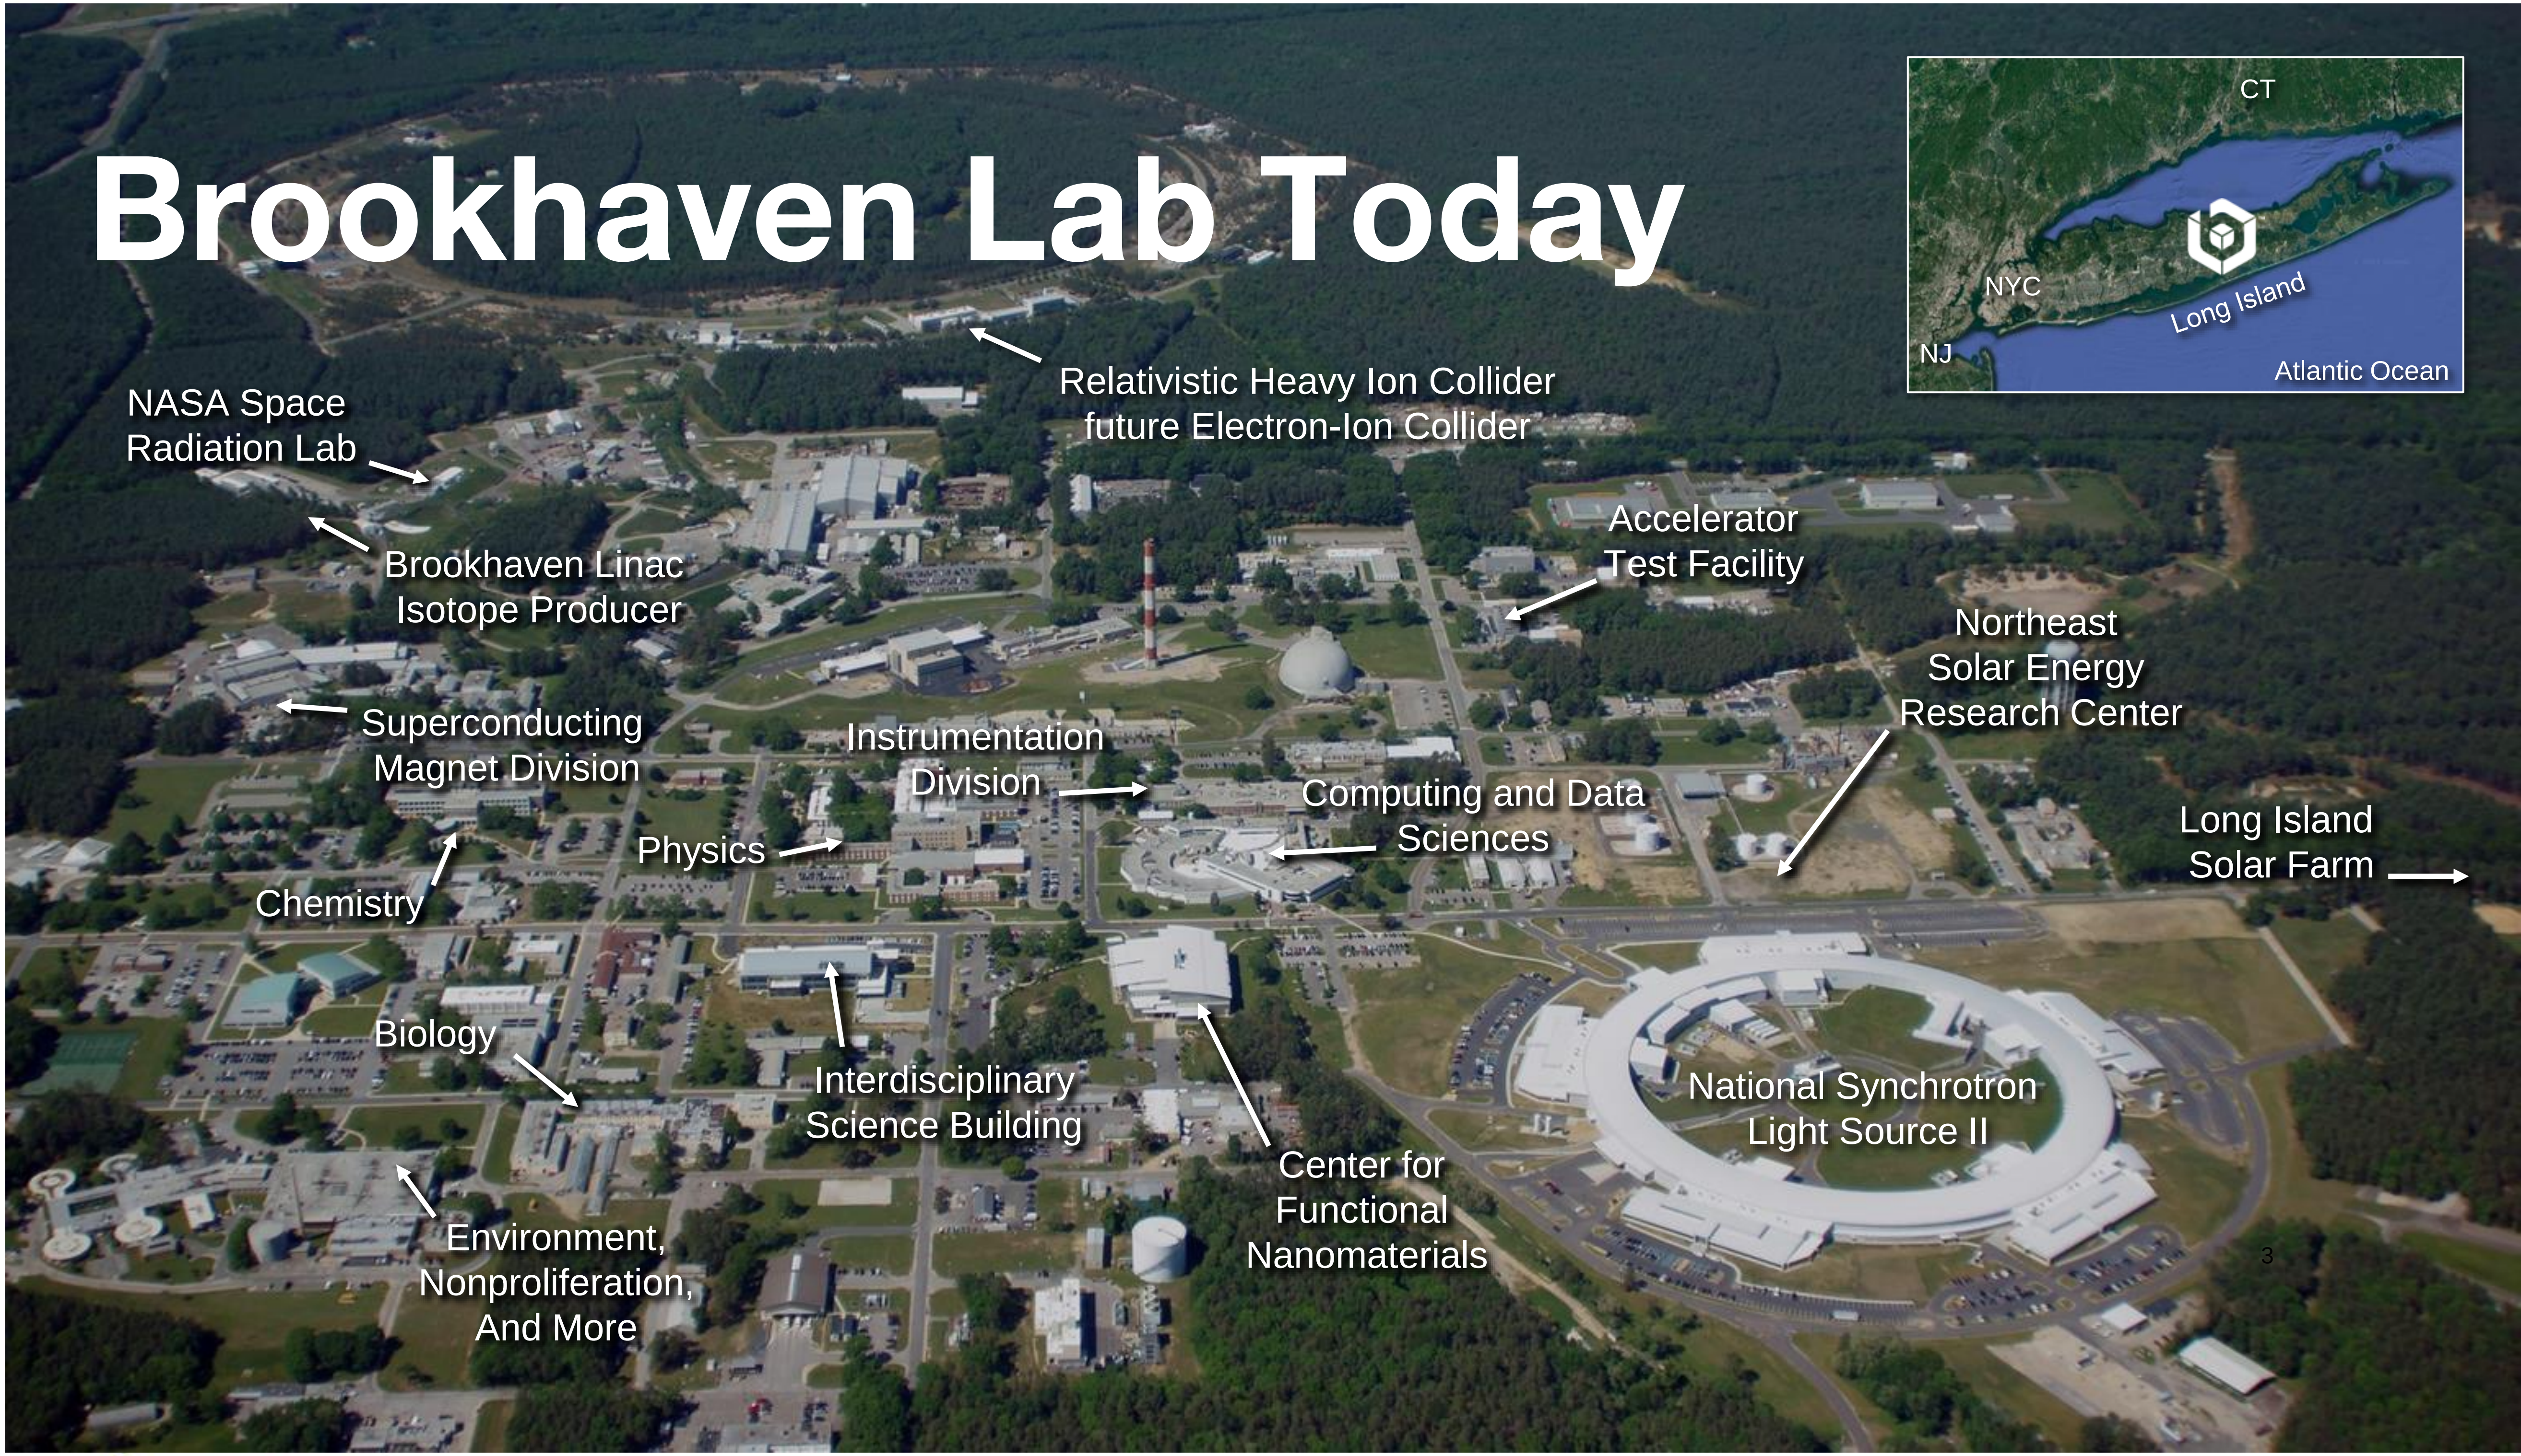

# Brookhaven Supports Data-rich Experimental and Computational Facilities and Programs

Relativistic Heavy Ion Collider (**RHIC**): Supports nearly 2000 scientists worldwide

Electron-Ion Collider (**EIC**): Echelon0 and Echelon1 data storage and processing center, a new paradigm in physics and frontier for data science*

Large Hadron Collider (**LHC**): Largest ATLAS Tier-1 center

**Super KEKB**: Tier-1 and RAW data storage center for high energy physics Belle II experiment

Quantum chromodynamics (**QCD**): computing facilities for Brookhaven Lab, RIKEN, and U.S. QCD communities

National Synchrotron Light Source II (**NSLS-II**): Newest and brightest synchrotron in the world; supports a multitude of scientific research in academia, industry, and national security

Center for Functional Nanomaterials (**CFN**): Combines theory and experiment to probe materials

Accelerator Test Facility (**ATF**): User facility for for advanced accelerator and laser research

Atmospheric Radiation Measurement (**ARM**) program: Partner in multi-site facility, operating its external data center

RHIC → EIC

LHC

QCD

NSLS-II

CFN

Brookhaven National Laboratory

*)*Torre's talk tomorrow about ePIC/EIC computing*

# Computing and Data Sciences Directorate

# Computing and Data Sciences Directorate Organization

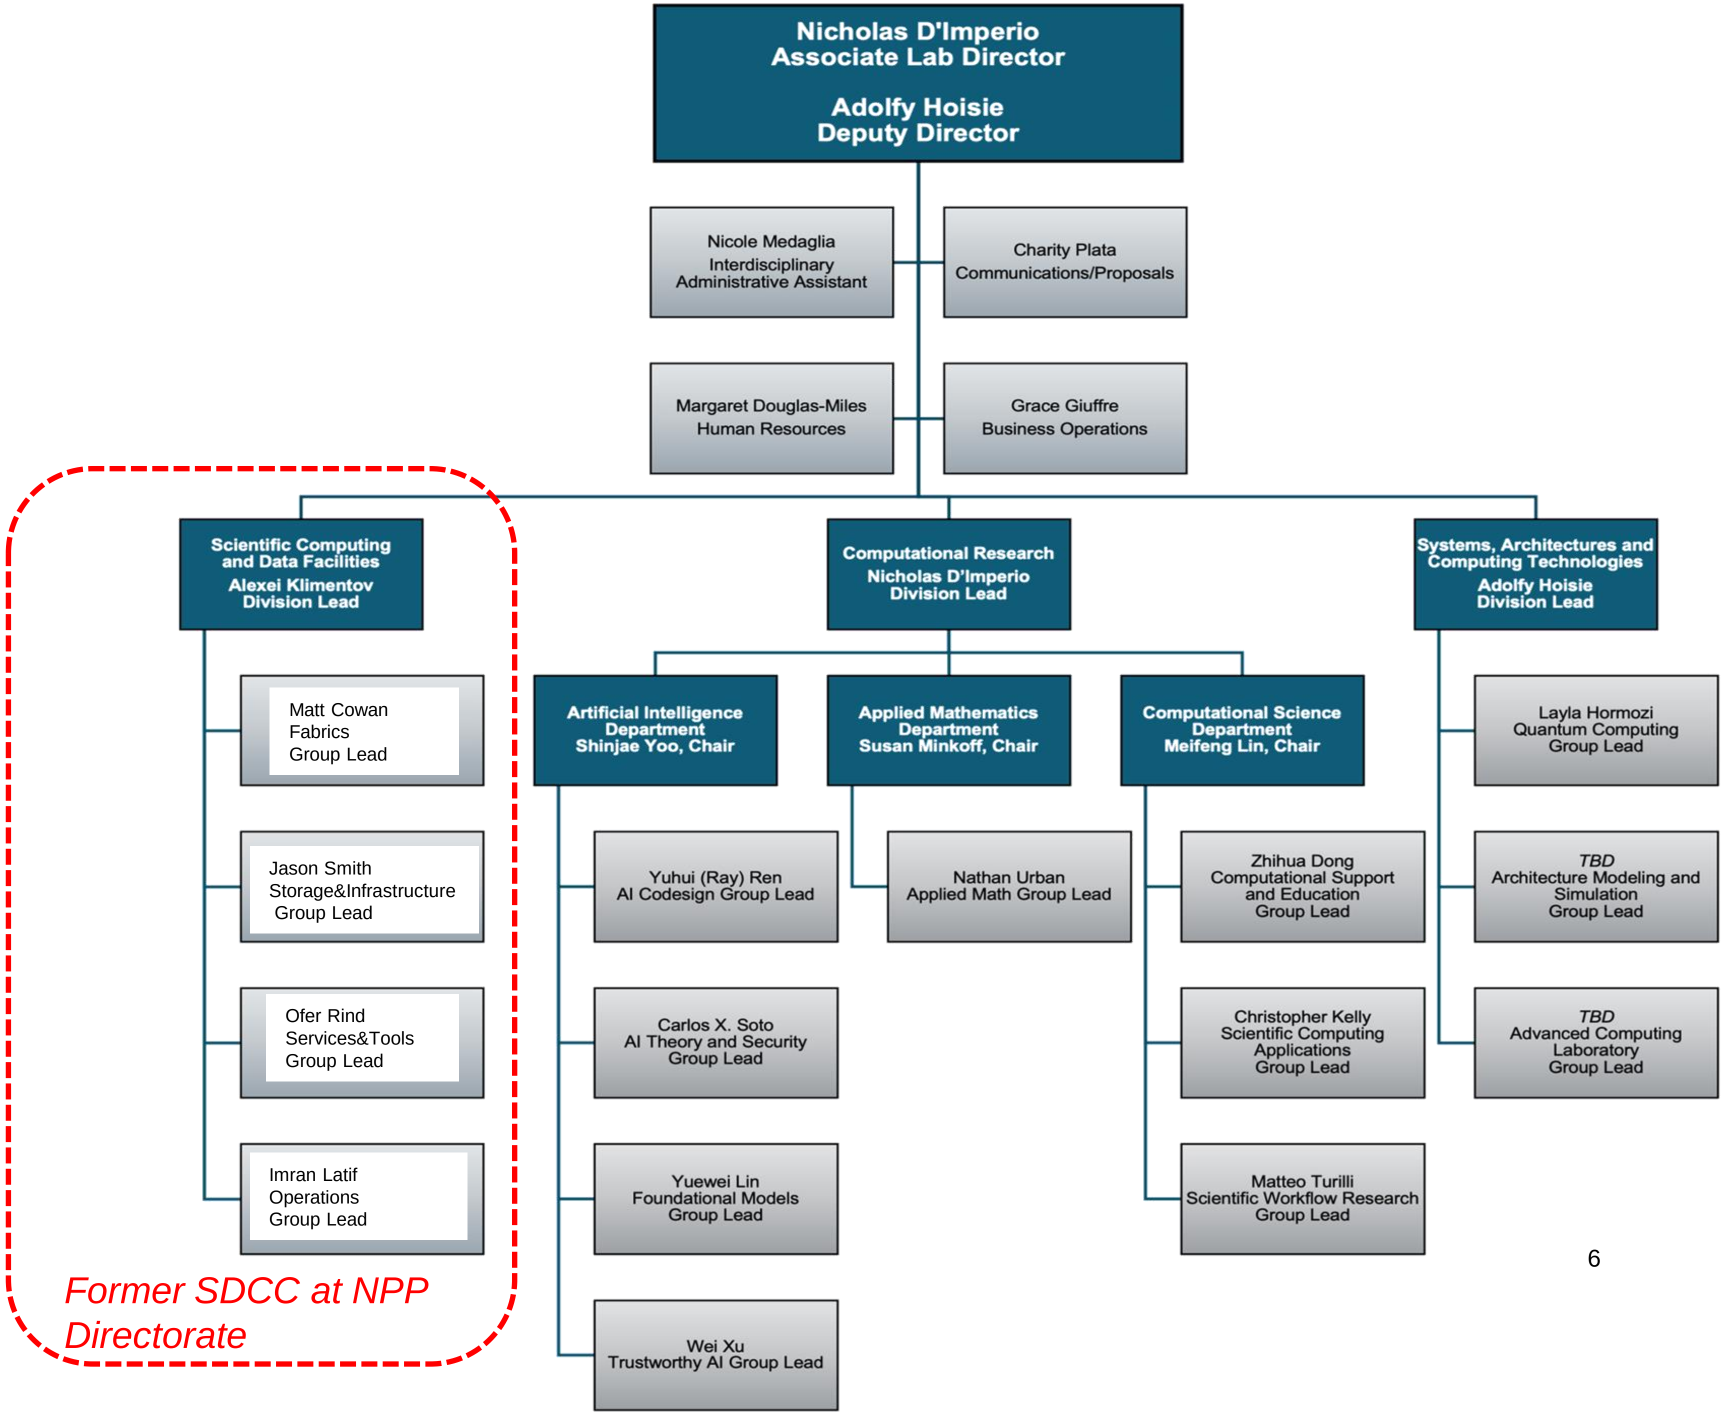




Brookhaven National Laboratory

# CDS at Brookhaven Lab

***Provides Brookhaven Lab with leading-edge computing capabilities that enhance its ability to fulfill its mission today and into the future by*:**

- Conducting mission-informed foundational research in computer science, applied mathematics, and quantum information science
- Creating early hardware and software prototypes in close collaboration with other Lab science directorates and the broader DOE community
- Providing computing and data services to enable research for the Lab's broad scientific user community through Scientific Computing and Data Facilities
- Acting as a coordination point for cross-cutting computer science, applied mathematics, and computational science research at Brookhaven Lab
- Developing computing-oriented strategic initiatives for the Laboratory
- Being mapped to DOE-ASCR's strategic directions, programs, and facilities.

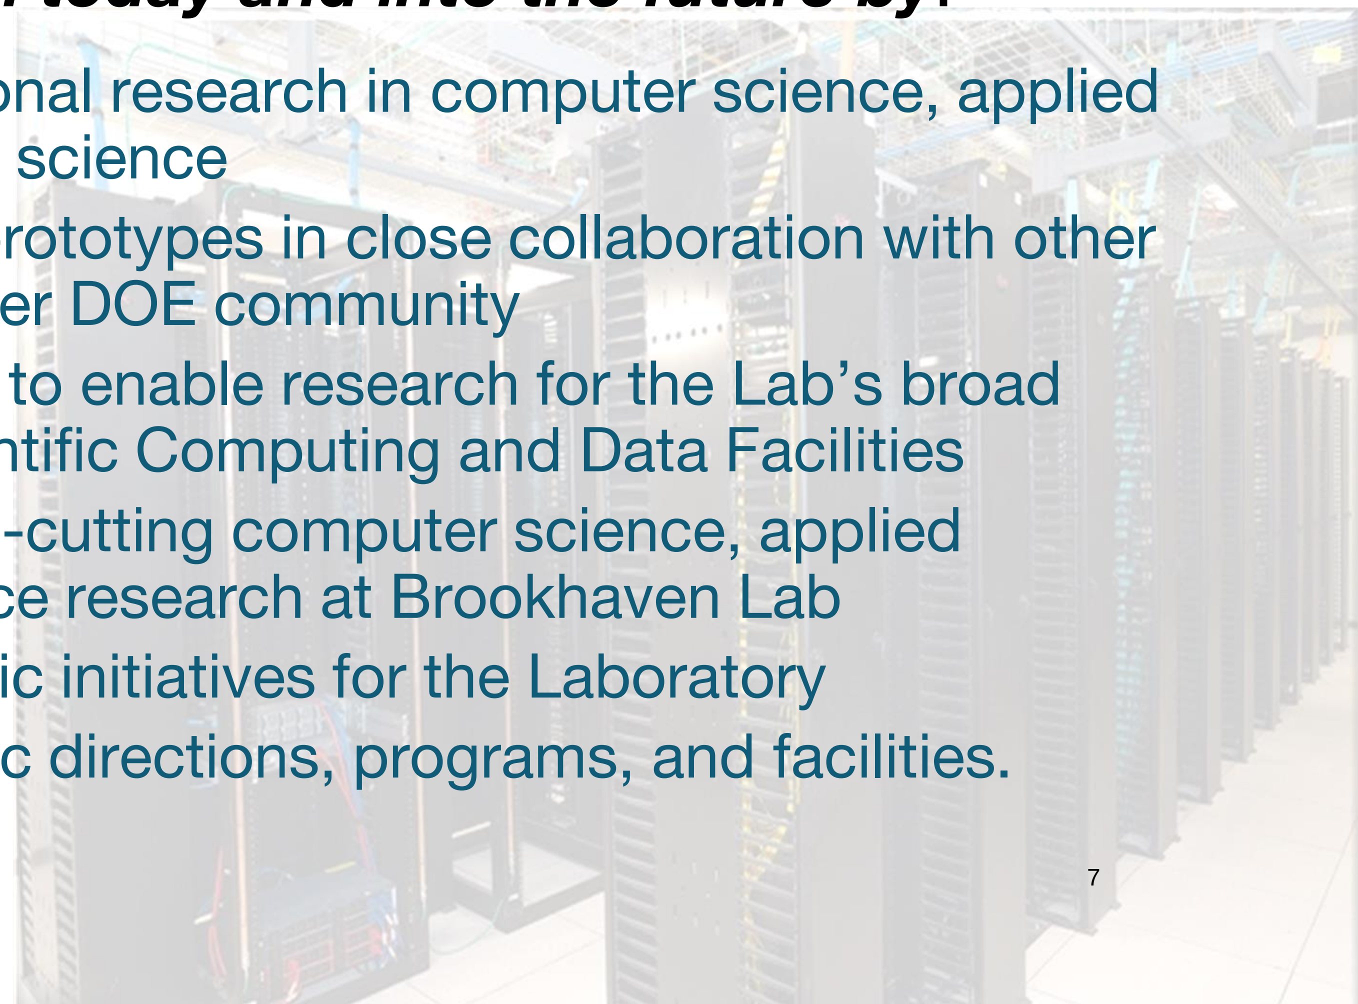

# Scientific Computing and Data Facilities Division

Brookhaven National Laboratory

# Scientific Computing and Data Facilities By the Numbers

SCDF today is a leading computing center for high-throughput computing and scientific data. SCDF supports high energy and nuclear experiments, as well as other Lab-based and international projects, including EIC, LQCD, NSLS-II, CFN, BES, Worldwide LHC Computing Grid (WLCG), and Open Science Grid (OSG).

One of the top-10 scientific data centers in the world

- The largest NP and HEP data archive in the United States
- ~300+ PB of data archived (exabytes by 2030)
- 160+ PB on disks
- The mass storage HPSS system is used in Data Carousel mode when data are actively migrated between disk and tape

~2000 nodes (~175k CPU cores and 350 GPU)

- additional 2200 nodes (~240k CPU cores will be deployed in 2025)

1.5 EB of data analyzed

BNL network externally connected to global network at 1.6 Tbps

Infrastructure: 59,000 sq-ft$^2$ data center began operations in 2021

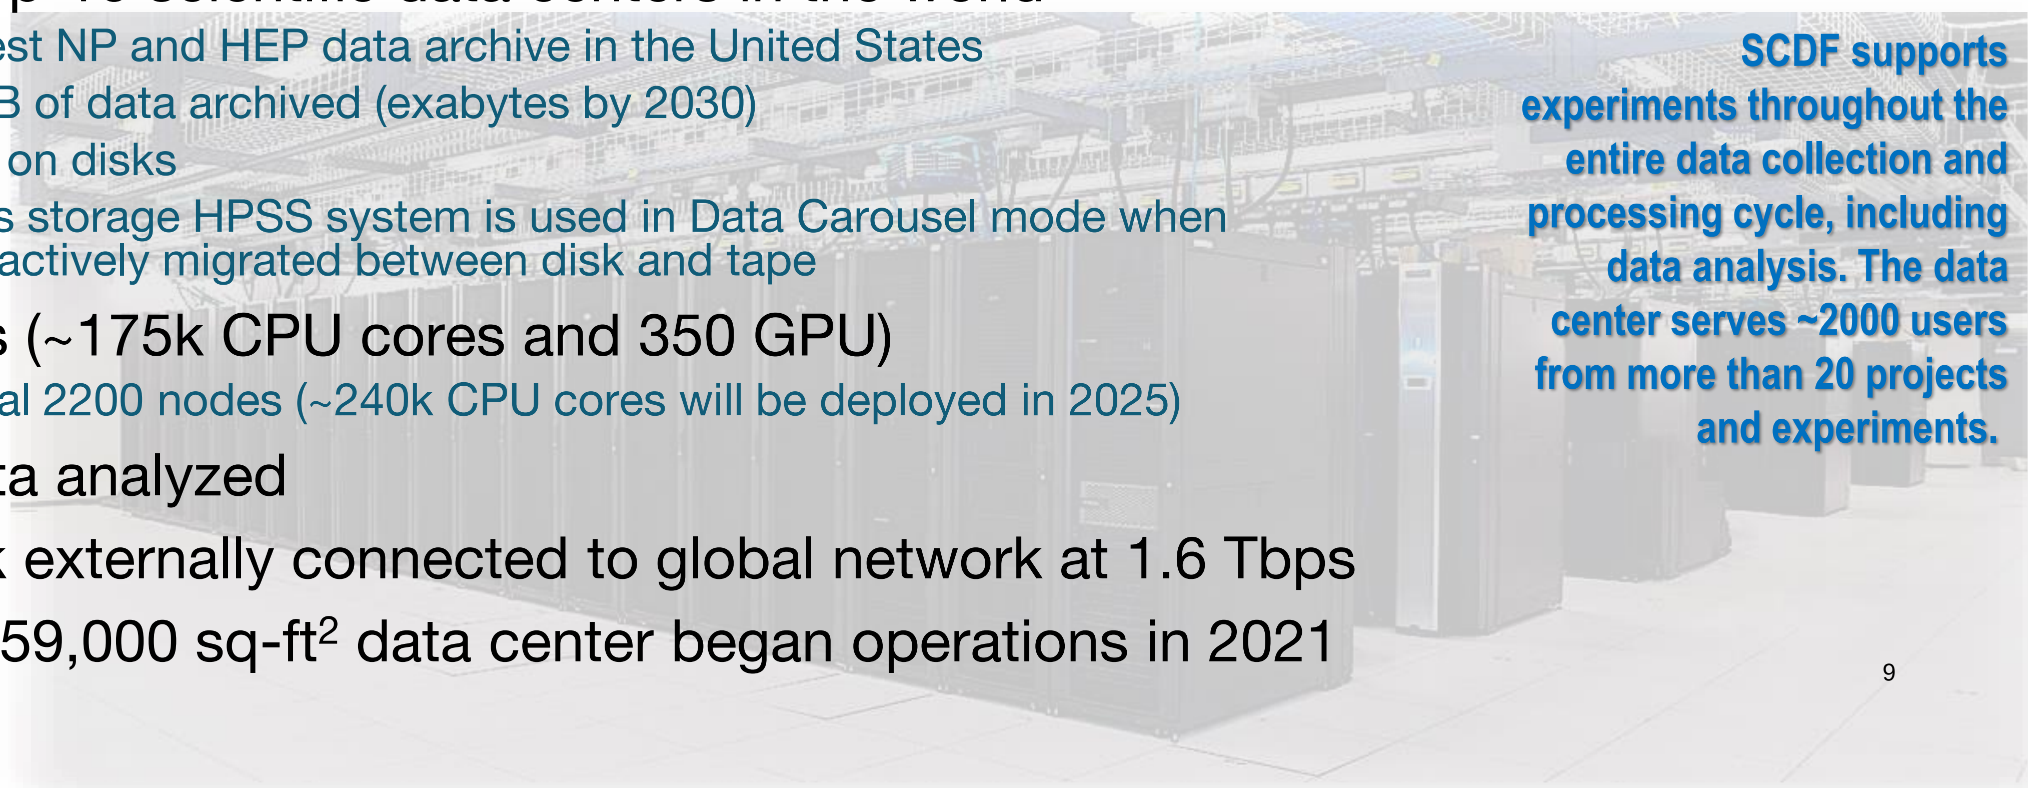

**SCDF supports experiments throughout the entire data collection and processing cycle, including data analysis. The data center serves ~2000 users from more than 20 projects and experiments.**



Brookhaven National Laboratory

# Data Center Key Features


- Modular, scalable, and robust design: 9.6 MW of ultimate IT payload capacity
- Currently, 3.6 MW (UPS/Generator) backed-up power available
- Additional 1.2 MW power block to be available in early FY26
- Cooling with high-efficiency chillers and rear door heat exchangers
- Liquid (direct to chip/immersion) cooling ready for latest GPU-based IT hardware deployment for AI, HPC, and Digital Twin applications
- Energy-efficient data center with 1.3 current PUE, aiming for 1.2 PUE with full IT payload deployment
- Streamlined operations through Data Center Infrastructure Management (DCIM), including node level electric billing, asset management, environmental monitoring, and capacity planning



I.Latif topics for Thu Jan 23rd meeting
"BNL-ATLAS Collaboration Ideas: Advancing Data Center Sustainability"

# sPHENIX: Newest Experiment in the RHIC Program

Completes the RHIC physics program with a three-year run (2023-2025)

Unique dual-streaming at very high rates (20 GB/s in aggregate)

- First stream to archival tape for storage and backup
- Second stream to disk for near-real-time (preliminary) data analysis

**Experience builds know-how to support similar high-rate programs in the future (i.e., ePIC at the EIC and ATLAS at the HL-LHC).**

Requires massive amounts of resources to keep up with experimental needs

- CPU deployment for RHIC in 2025 grows ~79% to 160,736 computing cores
- Tape archival storage estimated needs in the three-year run is ~0.5+ EB
- For reference, the total data currently archived on tape at the SCDF is ~300+ PB
- Total disk storage capacity (for data analysis) to be deployed in 2025 is ~87 PB

# Projects

# Two New Research Centers

## Computing for National Security
**Adolfy Hoisie, Director***

- Coordinate efforts across different departments and directorates with a range of agencies

## Digital Twins Institute
**Meifeng Lin, Director**
**Vanessa Lopez-Marrero, Chief Scientist**

- Coordinate digital twin research and development (R&D)
- Develop new funding opportunities

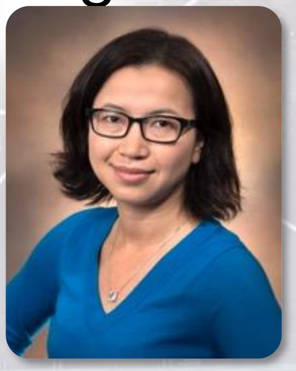

# CDS and NPP Projects*

DOE ASCR, BNL LDRD, BNL PD, BER, SciDAC and Quantum Computing selected projects

- Hierarchical, AI-Enabled Modeling of Future Supercomputers
- Novel Quantum Algorithms from Classical Transforms
- Quantum Algorithms Across Topological and Quantum Circuit Models
- Optimal Experimental Design (OED) for Quantum Technologies
- REDWOOD: Resilient Federated Workflows in a Heterogeneous Computing Environment
  BNL : NPPS, CDS; SLAC; ORNL; UMass; CMU; UPitt
- EIC Computing at BNL
  CDS, NPPS, CAD, Nuclear Physics
- DUNE Computing at BNL
- FCC (including FCC Analysis Center prototype)
- Data and Analysis Preservation for RHIC
- NP Collaborative Tools
- Codesign in Action for Experimental Science Computing: Architectures, Systems, and Testbeds

*) the list is incomplete (and many projects are BNL wide)
Torre's talk about AI/ML @NPP

# REDWOOD: Resilient Federated Workflows in a Heterogeneous Computing Environment

**Goal**: Optimal data placement and workload scheduling enhancing the resilience, throughput, and resource utilization.

**Approach**

- Multi-objective, multi-constraint optimization algorithm to assign jobs and allocate datasets
- Enhanced monitoring to provide deeper understanding of the workflows and distributed systems
- Remote streaming capability with data reduction to support near-real-time workflows
- System modeling approaches for large-scale distributed and heterogeneous computation grids

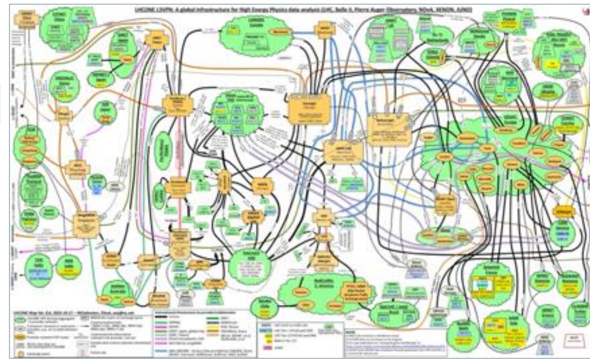

World Nuclear and Particle Physics Research Network

**Project Structure**

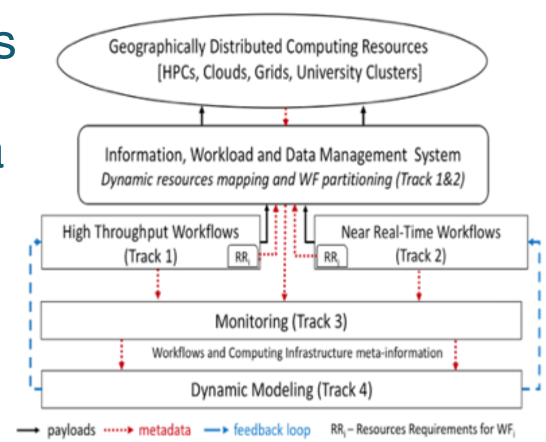


**Research Accomplishments**

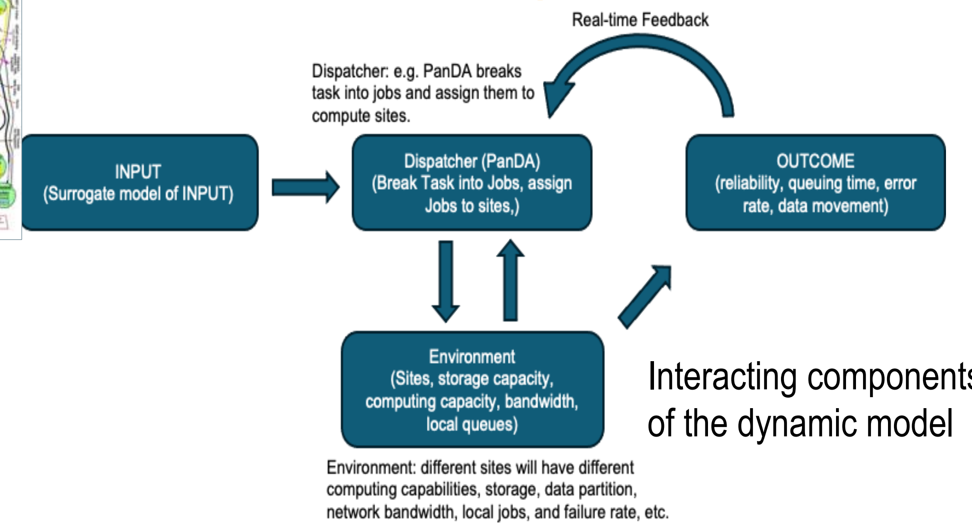


Interacting components of the dynamic model

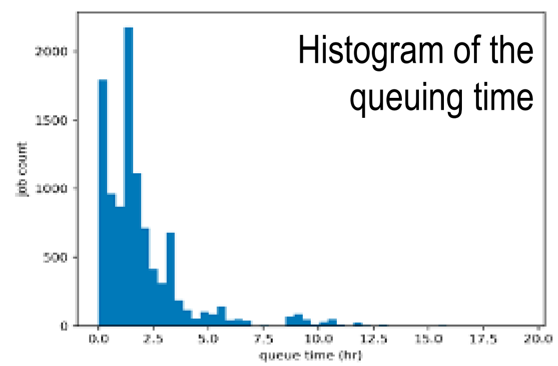


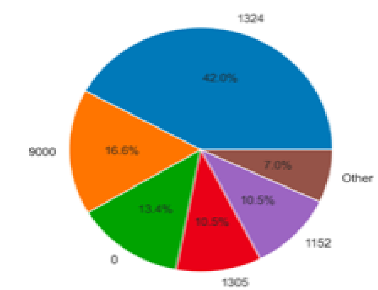


Error type breakdown

Oak Ridge National Laboratory · SLAC · University of Massachusetts Amherst · Carnegie Mellon University · University of Pittsburgh

Brookhaven National Laboratory

# REDWOOD: Developing an Innovative Computing Ecosystem to Impact Science

- Three National Labs, Brookhaven, Oak Ridge and SLAC, and three universities, Carnegie Mellon, University of Massachusetts Amherst, and University of Pittsburgh
- Lead PI A. Klimentov (BNL) with co-PIs A. Hoisie, T. Maeno and S. Yoo (BNL); S. Klasky (ORNL); and W. Yang (SLAC); as well as a BNL team of computing scientists, IT engineers, and physicists
- Ecosystem of research platforms connected to address scientific challenges involving complex workflows and exabyte data volume in heterogeneous computing environments.
- Data placement and complex workflows created by brokering and partitioning algorithms and associated runtime to help science applications better exploit architectural features found in DOE's computing infrastructure.
- Incorporate timely new algorithms to provide near-term high impact on science (with domain scientists).
- Software and algorithms will be demonstrated at scale for several scientific domains, e.g., particle physics, astronomy, and nuclear fusion.
- Make software and algorithms easily available to the research community for broader, long-term impact.



*Splinter meeting on Thu Jan 23rd organized by Paul Nilsson and Frederic Suter*

# Codesign in Action for Experimental Science Computing: Architectures, Systems, and Testbeds

## Advanced Computing Lab

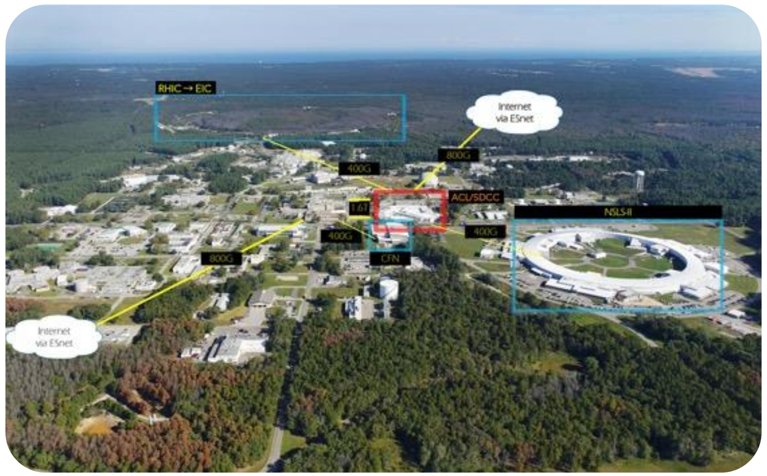

Unique collaborative testbed facility with access to live, actual data from diverse experiments, such as CFN (microscopy), NSLS-II, and RHIC, for codesign of architectures and experimental workflows.

## Experimental Science Workflows

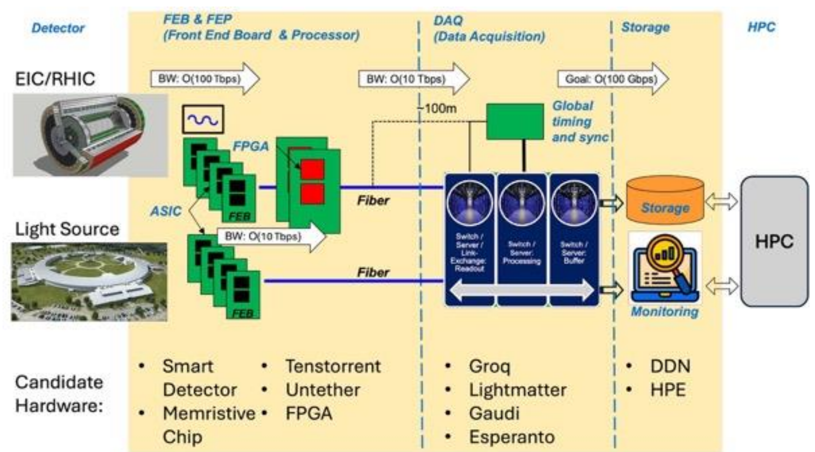


Extreme data challenges, heterogeneity, large spectrum of spatial and temporal computing scales from the edge to the extreme – real-time to long computational campaigns.

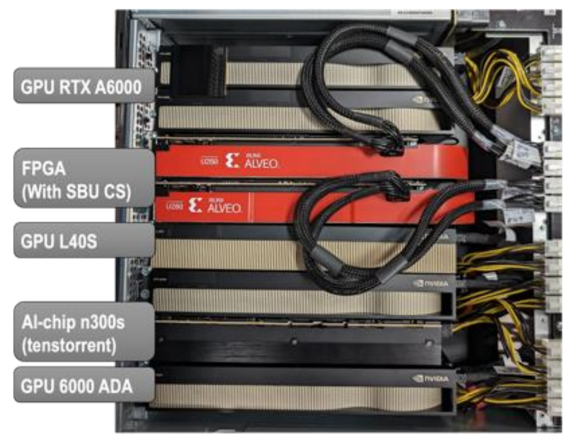


## AI-based Modeling and Simulation

Leading the charge with SimNet and PerfVec
Accurate simulation faster by orders of magnitude compared with Discrete-Event Simulation

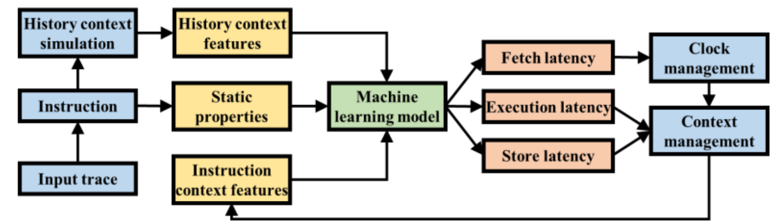


*SimNet: AI-based architecture simulation https://github.com/lingda-li/simnet*

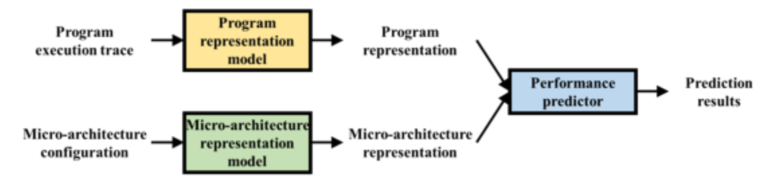


*PerfVec: AI-based Architecture modeling https://github.com/PerfVec/PerfVec*

Li, L. S. Pandey, T. Flynn, H. Liu, N. Wheeler, and A. Hoisie. 2022. SimNet: Accurate and High-Performance Computer Architecture Simulation using Deep Learning. POMACS 6(2):Article 25. DOI: 10.1145/3530891.

Pandey, S., L. Li, T. Flynn, A. Hoisie, and H. Liu. 2022. Scalable Deep Learning-Based Microarchitecture Simulation on GPUs. SC22, pp. 1-15. DOI: 10.1109/SC41404.2022.00084.



Brookhaven National Laboratory

Staff: Lingda Li; Thomas Flynn; Yuhui Ren; Alexei Klimentov; Adolfy Hoisie

# Supplementary Slides

# Emmy Noether Fellowship

Fellow **Sanket Jantre** has worked on two activities:

## Active subspace reduction for Bayesian deep learning

- Reduce high-dimensional neural network weight space to <100 dimensions
- Can be applied to differentiable hybrid models
- *Outlook*: transitioned to Data Reduction

## Neural system identification of coarse-grained PDE surrogates

- Learn effective PDE (partial differential equation) governin equations that describe coarse-grained dynamics of higher-fidelity simulations
- *Outlook*: Developing synthetic data augmentation methods for greatly enhanced training sample efficiency. Evaluating for possible transition into other projects (e.g., MMICC)

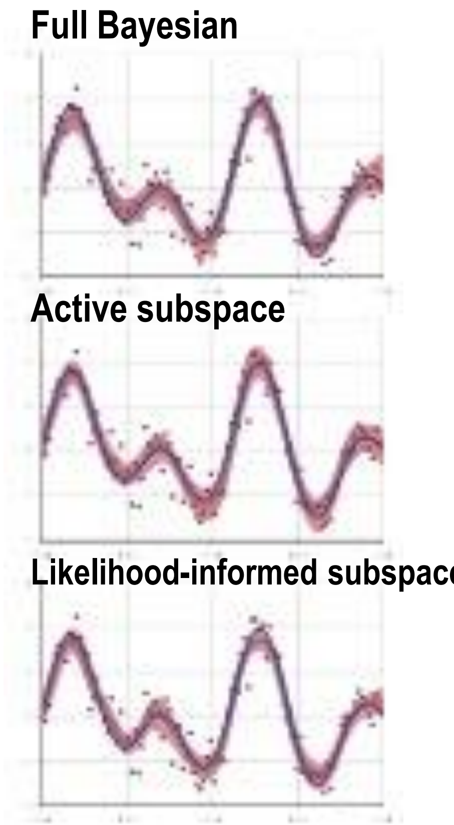


Reduction of a 6-layer deep neural net's weight space to a 30D parameter space shows almost no loss of quality in uncertainty quantification.

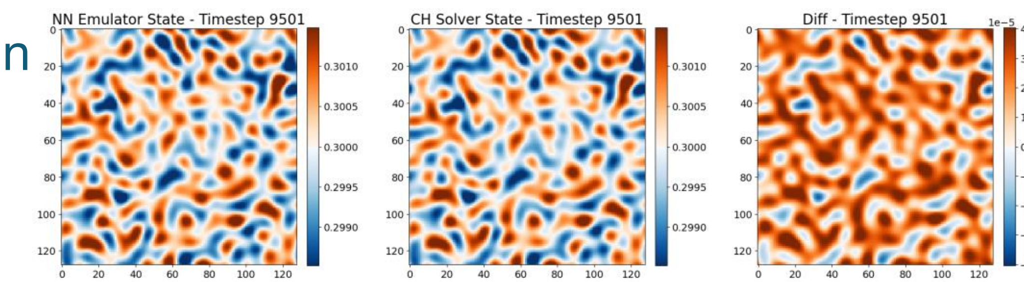


Learning a neural approximation to a 2D PDE's governing equations shows high accuracy when the data-driven neural PDE solver is time integrated.



Brookhaven National Laboratory

# SciDAC ImPACTS

## (Online) Tensor Factorization (TF)

- TF techniques are used to discover dominant spatio-temporal modes in the system and to quantify dominant AMOC modes.
- Beneficial to detect change point that is exascale-ready with TF and compare modes of different modeling results.

## Explainability Study for Latent Representations

- Developed several dimension reduction methods, including TF, correlations, embedding, clustering, and ablations, to understand and compare latent representations.
- Demonstrated that learned latents of a climate reconstruction model incorporate the contextual information in their representations, leading to improved model performance over traditional methods.

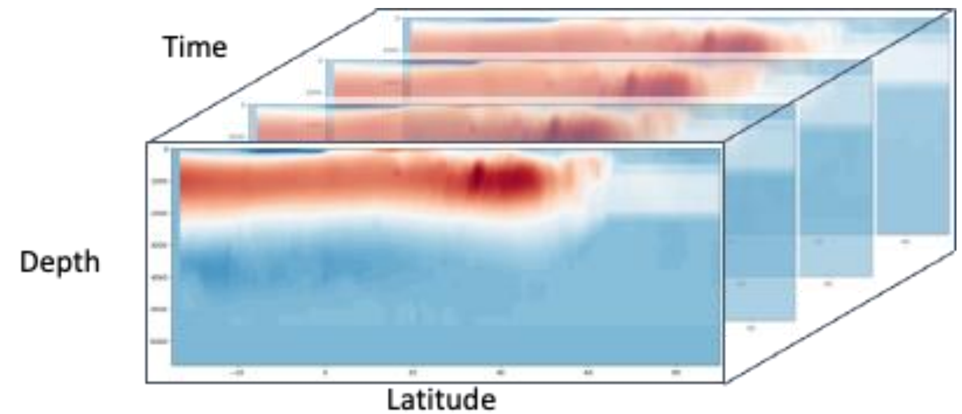


A Tucker Decomposition result of a CESM2 simulation

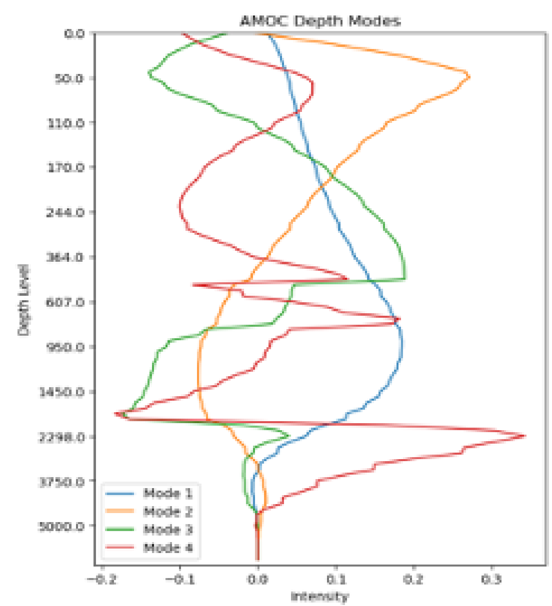


Tucker can extract additional information e.g., depth mode strength

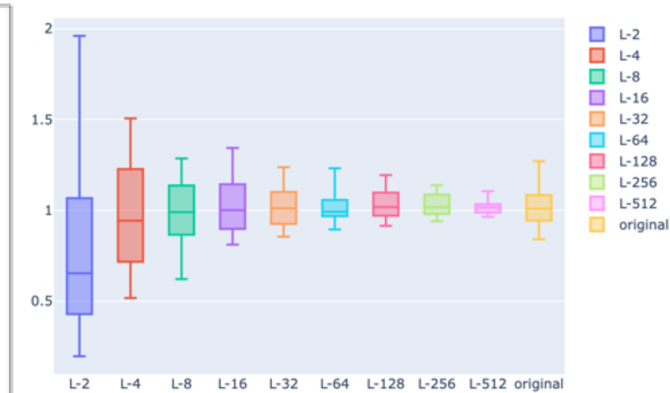


Explaining latent representations

W. Xu, D. F. DeSantis, X. Luo, A. Parmar, K. Tan, B. T. Nadiga, Y. Ren, and S. Yoo, "Studying the Impact of Latent Representations in Implicit Neural Networks for Scientific Continuous Field Reconstruction," AAAI co-held XAI4Sci workshop.

Staff: Wei Xu; Xihaier Luo; Yihui Ren; Shinjae Yoo

Brookhaven National Laboratory

# BER: LUCID and RADBio

Improve modeling of effects of low-dose radiation on the human body.

Working in various areas: biological pathway latent space dynamics identification; prompt engineering for constructing pathways from bio literature; uncertainty quantification for large language model (LLM) knowledge extraction; optimal experimental design.

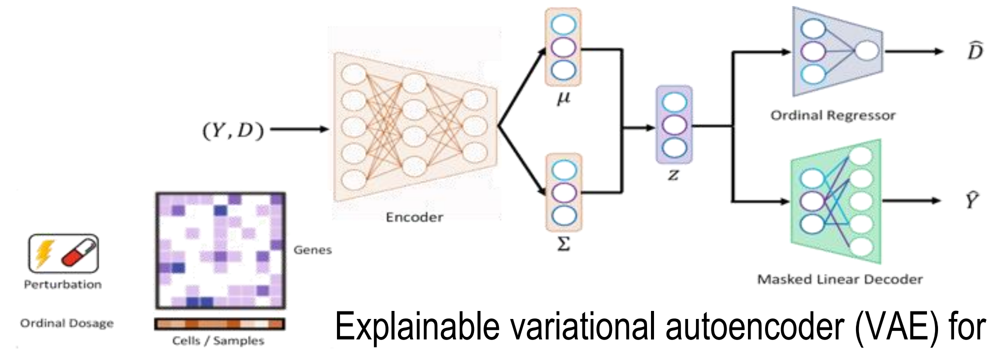


Explainable variational autoencoder (VAE) for ordinally-perturbed transcriptomics data.

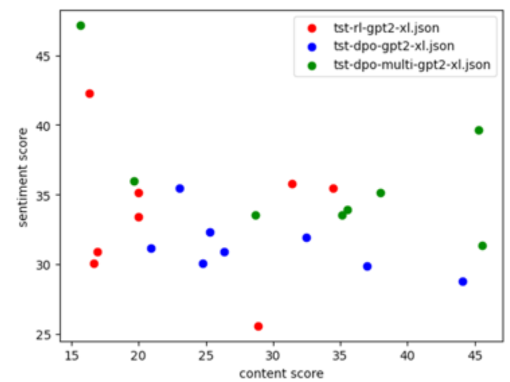


Content and style functions generate an optimal Pareto frontier of LLM prompts for biological knowledge elicitation.

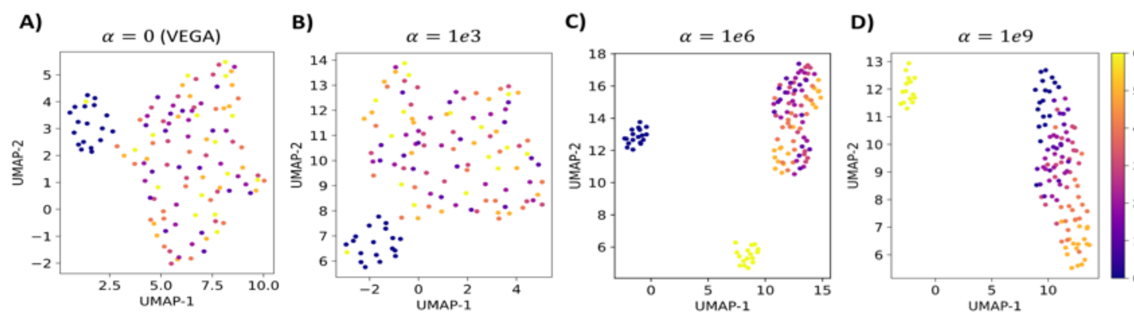




VAE identifies changes within latent space after gradual exposure to radiation.

Staff: Gilchan Park; Guang Zhao; Sanket Jantre; Xihaier Luo; Vanessa Lopez-Marrero; Xiaoning Qian; Matteo Turilli; Mikhail Titov; Nathan; Urban, Byung-Jun Yoon

Brookhaven National Laboratory

# Hierarchical, AI-Enabled Modeling of Future Supercomputers

**Goal**: Develop a modular and hierarchical modeling framework to explore and optimize system-level impacts of beyond-CMOS technologies

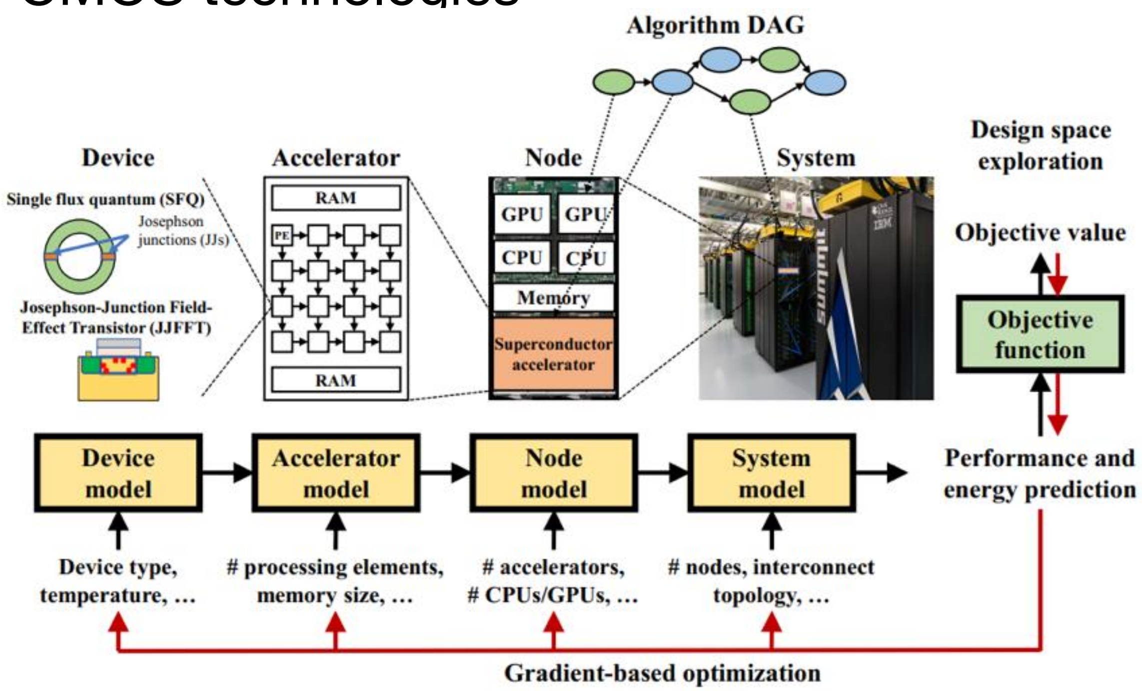


CNN = Convolutional neural network

## Achievements

- Modeled the accelerator performance and power of CNNs using Timeloop
- Developed loop fusion strategies to alleviate memory bottleneck
- Provided a system model for GEMM (general matrix multiply)

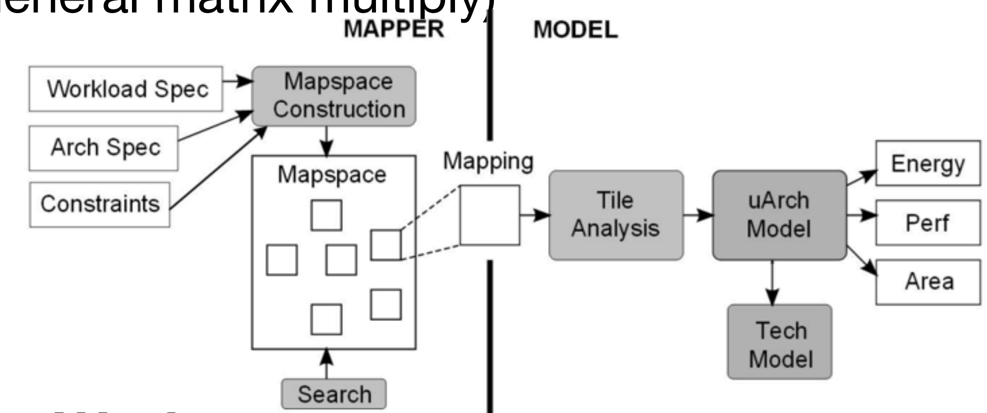


## Future Work

- Expand and integrate with system models
- Model broader scientific applications
- System-level design space exploration

Staff: Lingda Li; Adolfy Hoisie

Brookhaven National Laboratory

# SciDAC RAPIDS

**Scientific Achievement:** A new, implicit neural representations (INRs) method to accurately reconstruct continuous physical fields from few sensor points. The model achieves state of the art in continuous field reconstruction and outperforms other models on two benchmark datasets.

**Significance and Impact**

- Addresses a challenging problem, where the locations and number of sensors keep changing.
- Predicts a continuous physical field for an arbitrary resolution.
- Model surpasses existing methods for scientific data.
- Technology has been applied to SciDAC ImPACTS.

**Technical Approach**

- Introducing a context-aware indexing mechanism that compared to standard time index (t)-based INR models, incorporates additional semantic information.
- Factorizing target signals into a set of multiplicative basis functions and applying element-wise shift and scale transformations to the latent codes.

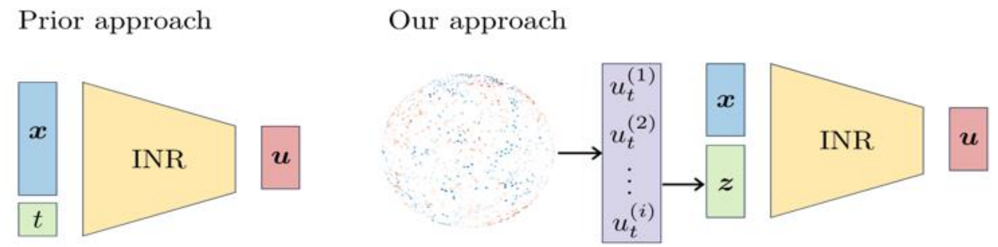


Overview of our method, supporting sparse and irregular data.

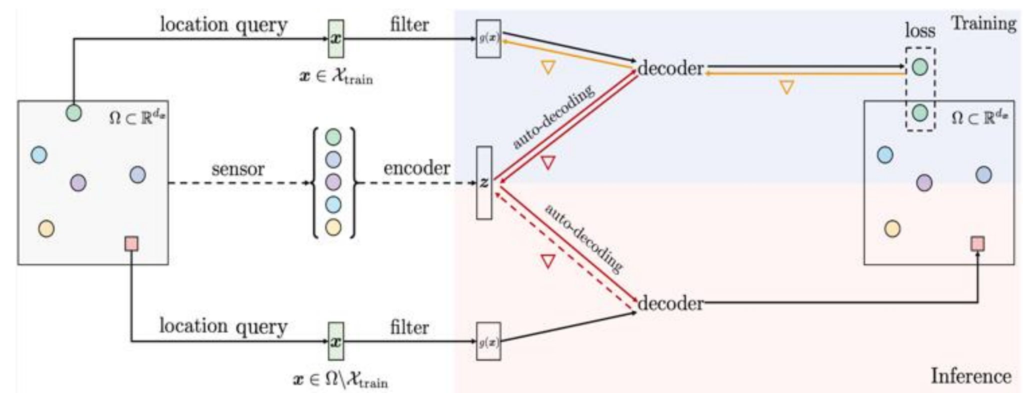


INR training and inference framework to create continuous field from sparse sensing data.

| Model | Simulation-based Data | | | | Satellite-based Data | | | |
|---|---|---|---|---|---|---|---|---|
| | Task 1 | Task 2 | Task 3 | Task 4 | Task 1 | Task 2 | Task 3 | Task 4 |
| | Sampling ratio $s = 5\%$ | | | | Sampling ratio $s = 0.1\%$ | | | |
| ResMLP | 1.951e-2 | 1.672e-2 | 1.901e-2 | 1.468e-2 | 1.717e-3 | 1.601e-3 | 1.179e-3 | 1.282e-3 |
| SIREN | 2.483e-2 | 2.457e-2 | 2.730e-1 | 2.455e-2 | 3.129e-1 | 4.398e-2 | 1.304e-2 | 9.338e-2 |
| FFN+P | 2.974e-2 | 1.121e-2 | 1.495e-2 | 8.927e-3 | 2.917e-3 | 2.392e-3 | 7.912e-4 | 7.565e-4 |
| FFN+G | 2.943e-2 | 1.948e-2 | 1.980e-2 | 1.426e-2 | 4.904e-3 | 7.969e-3 | 1.005e-3 | 1.044e-3 |
| MMGN | **4.244e-3** | **4.731e-3** | **3.148e-3** | **3.927e-3** | **1.073e-3** | **1.131e-3** | **6.309e-4** | **6.298e-4** |
| *Promotion* | 78.24% | 57.79% | 78.94% | 56.01% | 37.51% | 29.35% | 20.26% | 16.74% |

Our model (MMGN) outperforms all others on two benchmark datasets.



Luo, X., Xu, W., Nadiga, B., Ren, Y., & Yoo, S. (2023, October). Continuous Field Reconstruction from Sparse Observations with Implicit Neural Networks. In The Twelfth International Conference on Learning Representations.

Brookhaven National Laboratory

Staff: Shinjae Yoo; Yuhui Ren; Wei Xu; Xihaier Luo; Ghao Zhang

# Novel Quantum Algorithms from Classical Transforms

### Achievements

- Discovered a quantum Hermite transform that is exponentially faster than the classical Hermite transform
- Implementation in progress
- Algorithm permits fast-forwarding of certain Hamiltonian simulations that previously was exponential time

### Future Work

- Identification of sufficient conditions for classical numerical transforms supporting exponential quantum speedups based on recursive implementability

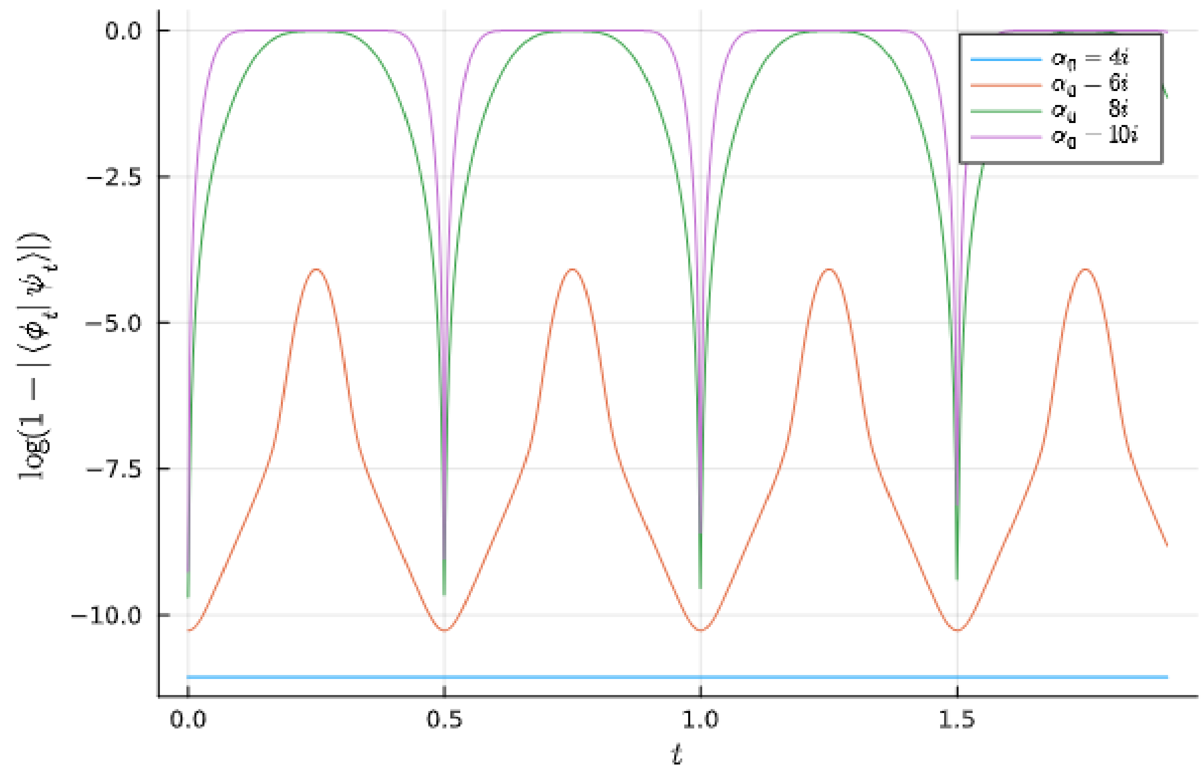


Staff: Ning Bao; Kwangmin Yu

Brookhaven National Laboratory

Microsoft DARTMOUTH University of Maryland

# Quantum Algorithms Across Topological and Quantum Circuit Models

## Achievements

- Created a method for initialization of logical qubits for the Fibonacci non-Abelian topological code. For a 2D lattice the number of steps grows as the square root of the number of plaquettes
- Devised methods for efficient braiding of pairs of logical anyons essential for implementing multi-qubit gates, reducing the cost to ~2.5 x braiding of single anyons
- Extended the application of quantum algorithms for calculating homology to knot theory

## Future Work

- Resource estimation for topological algorithms
- Combining logical computation with error correction steps
- Conceiving a novel quantum method to calculate the Khovanov homology and integrating it with topological fault tolerance

**Staff: Layla Hormozi; Seth Lloyd**

Brookhaven National Laboratory

# Optimal Experimental Design (OED) for Quantum Technologies

**Goals**: Make measurable predictions using physical models (Geant4/G4CMP) and optimize experiments to maximize observable differences

## Current Application

- Use simulations to test and optimize error mitigation strategies for quasiparticle poisoning (QPP)
- Optimize microstructure designs to improve quantum circuitry operation

## Achievements

- Reduced error rates: predicted 36% improved performance in error mitigation strategy.
- Better efficiency: Allows for less material to be removed from the substrate. Chips will be less fragile.
- More flexibility: Simulations now can include other chip materials for a generalized materials optimization approach.

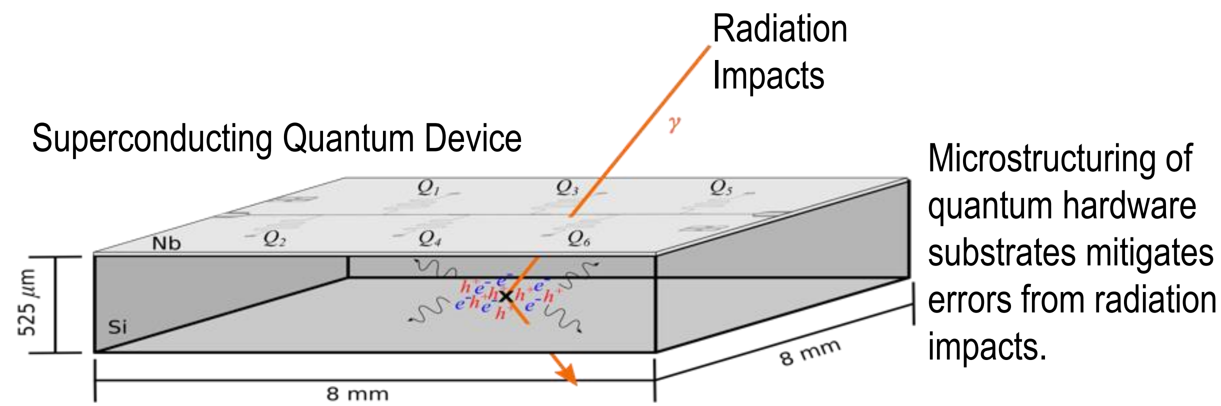


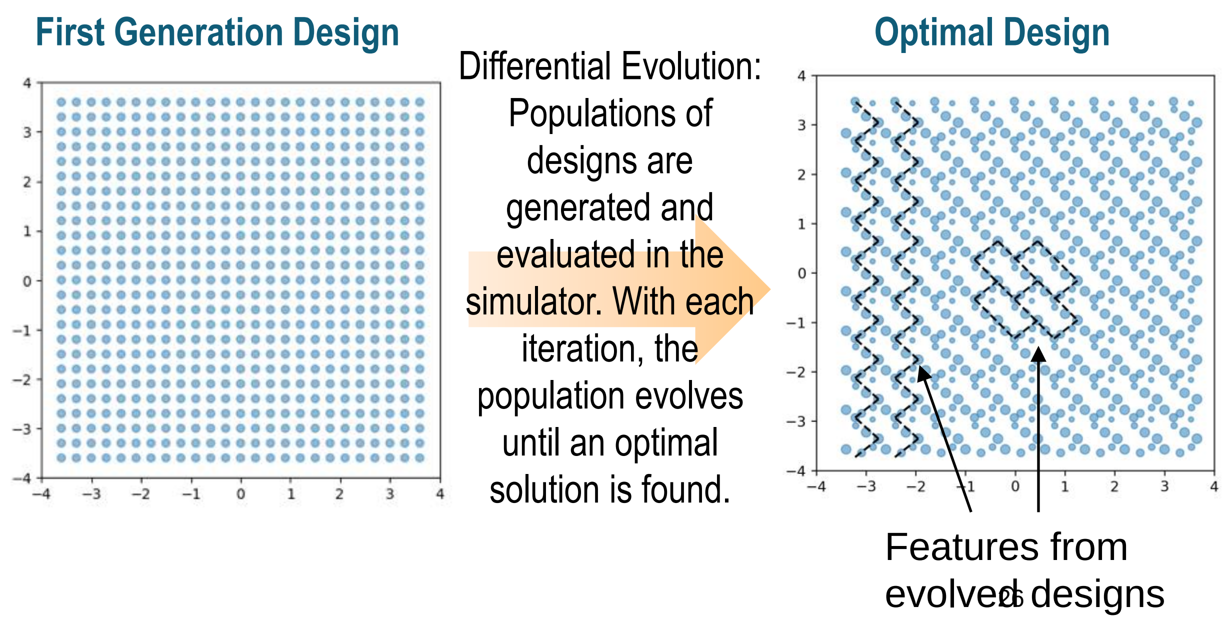




Staff: Paul Baity; Kristofer Reyes; Byung-Jun Yoon; Gilchan Park; Nathan Urban; Peter Love; Adolfy Hoisie

Brookhaven National Laboratory

Syracuse University

WISCONSIN UNIVERSITY OF WISCONSIN-MADISON

E. Yelton et al. *Phys. Rev. B* **110**, 024519 (2024).

# Lightning Talks: Subjects and Speakers

## Tuesday, October 15

- Advancing Integrative Structural Biology through HPC and AI, **Kriti Chopra**
- Multimodal Foundation Models, **Xi Yu**
- ML-driven Prediction of Resource, **Tasnuva Chowdhury**
- Using AI/ML for Data Placement Optimization in a Multi-tiered Storage System within a Data Center, **Qiulan Huang**

## Wednesday, October 16

- Modeling & Simulation of Performance and Resilience in Distributed Systems, **Yihui (Ray) Ren**
- An Active Learning-based Streaming Workflow for Reduced Data Training of Structure Finding Models in Neutron Diffractometry, **Tianle Wang**
- Neural Compression for sPHENIX, **Yi Huang**
- Learning Active Subspaces for Effective and Scalable Uncertainty Quantification in Deep Neural Networks, **Sanket Jantre**
- Decision Making for Autonomous Experiments, **Kristofer (Kris) Reyes**
- Future Supercomputer Modeling, **Lingda Li**