March 07, 2025

Overview of Computing at Argonne

Michael E. Papka
Division Director, Argonne Leadership Computing Facility
Deputy Associate Laboratory Director, Computing, Environment and Life Sciences, Argonne National Laboratory

Warren S. McCulloch Professor of Computer Science, University of Illinois Chicago

Argonne innovations in computing

From AVIDAC to Aurora … and everything in between

Argonne has always been a leader in computing innovation.

Complementing theoretical and experimental research, **computing** is a third pillar of science— revolutionizing how scientists learn, experiment, and theorize.

Argonne Computing Research

- Argonne Leadership Computing Facility (ALCF)
  - Provides world-class supercomputing resources.
  - Supports scientific breakthroughs via advanced HPC capabilities.
- Mathematics and Computer Science Division (MCS)
  - Develops mathematical algorithms and software technologies.
  - Drives innovation in computational sciences and engineering.
- Data Science and Learning Division (DSL)
  - Focuses on artificial intelligence and machine learning research.
  - Enables transformative data-driven scientific discoveries.
- Computational Science Division (CPS)
  - Advances computational modeling, simulation, and analysis.
  - Solves complex scientific and engineering challenges.

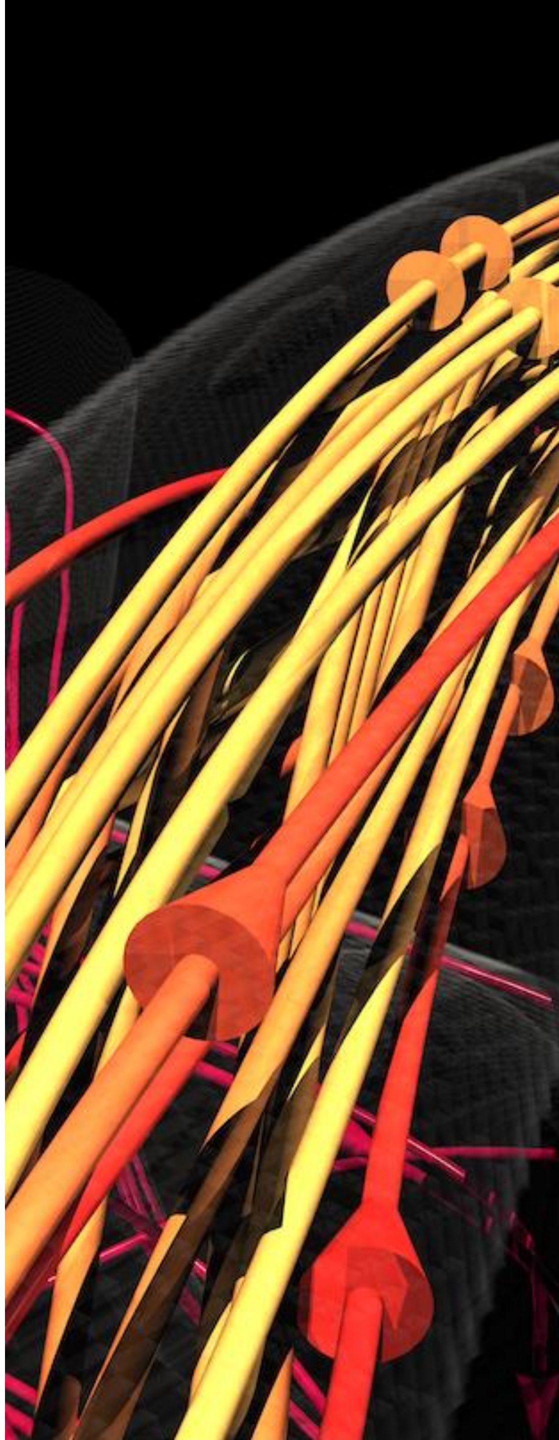

Laboratory Computing Resource Center (LCRC)

- Lab-wide resource for mid-range computing and data; **LCF on-ramp**
- Institutional resources: 1500 nodes, 12 petabytes

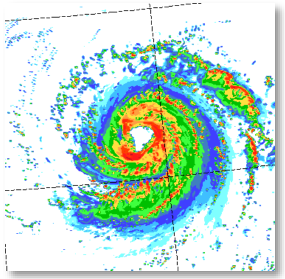

Climate
Predicting hurricane tracks to mitigate risks, hindcasting with climate model data to gauge impact of global change.

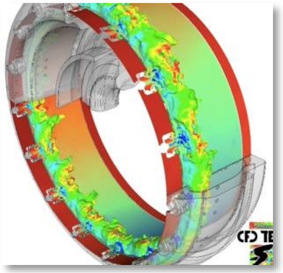

Power generation
Modeling two-phase flow and combustion for the design of more efficient aircraft engines.

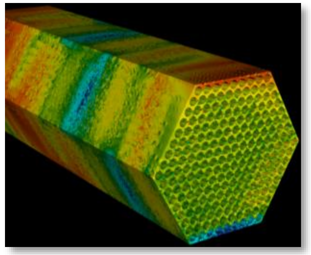

Nuclear Energy
High-fidelity fluid flow and heat transfer simulation of next-generation reactor designs, aiming to reduce the need for costly experimental facilities.

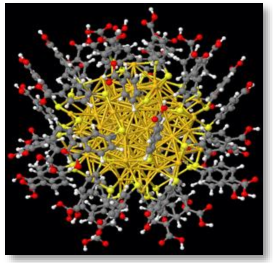

Materials Science
Molecular simulation of fracture dynamics in structural materials in next-generation nuclear reactors.

	Improv	Bebop	Swing
Nodes	825	672	6
Processors/node	2 AMD EPYC 7713	2 Intel Broadwell	8 NVIDIA A100 *GPU*s
Cores/node	128	36	
Memory /node	256GB (12 Nodes have 1TB)	128GB	320GB (1 node with 640 GB)
Total cores	105,600	24,192	
Peak	3.38PF	0.7PF	1 PF

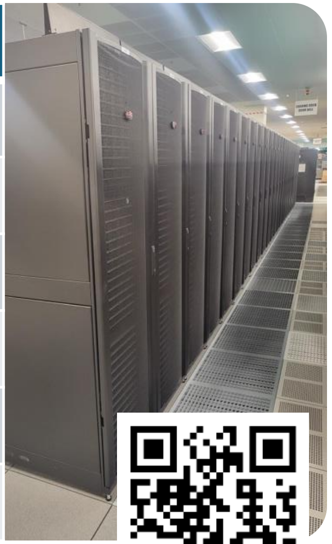

ALCF Systems Evolution

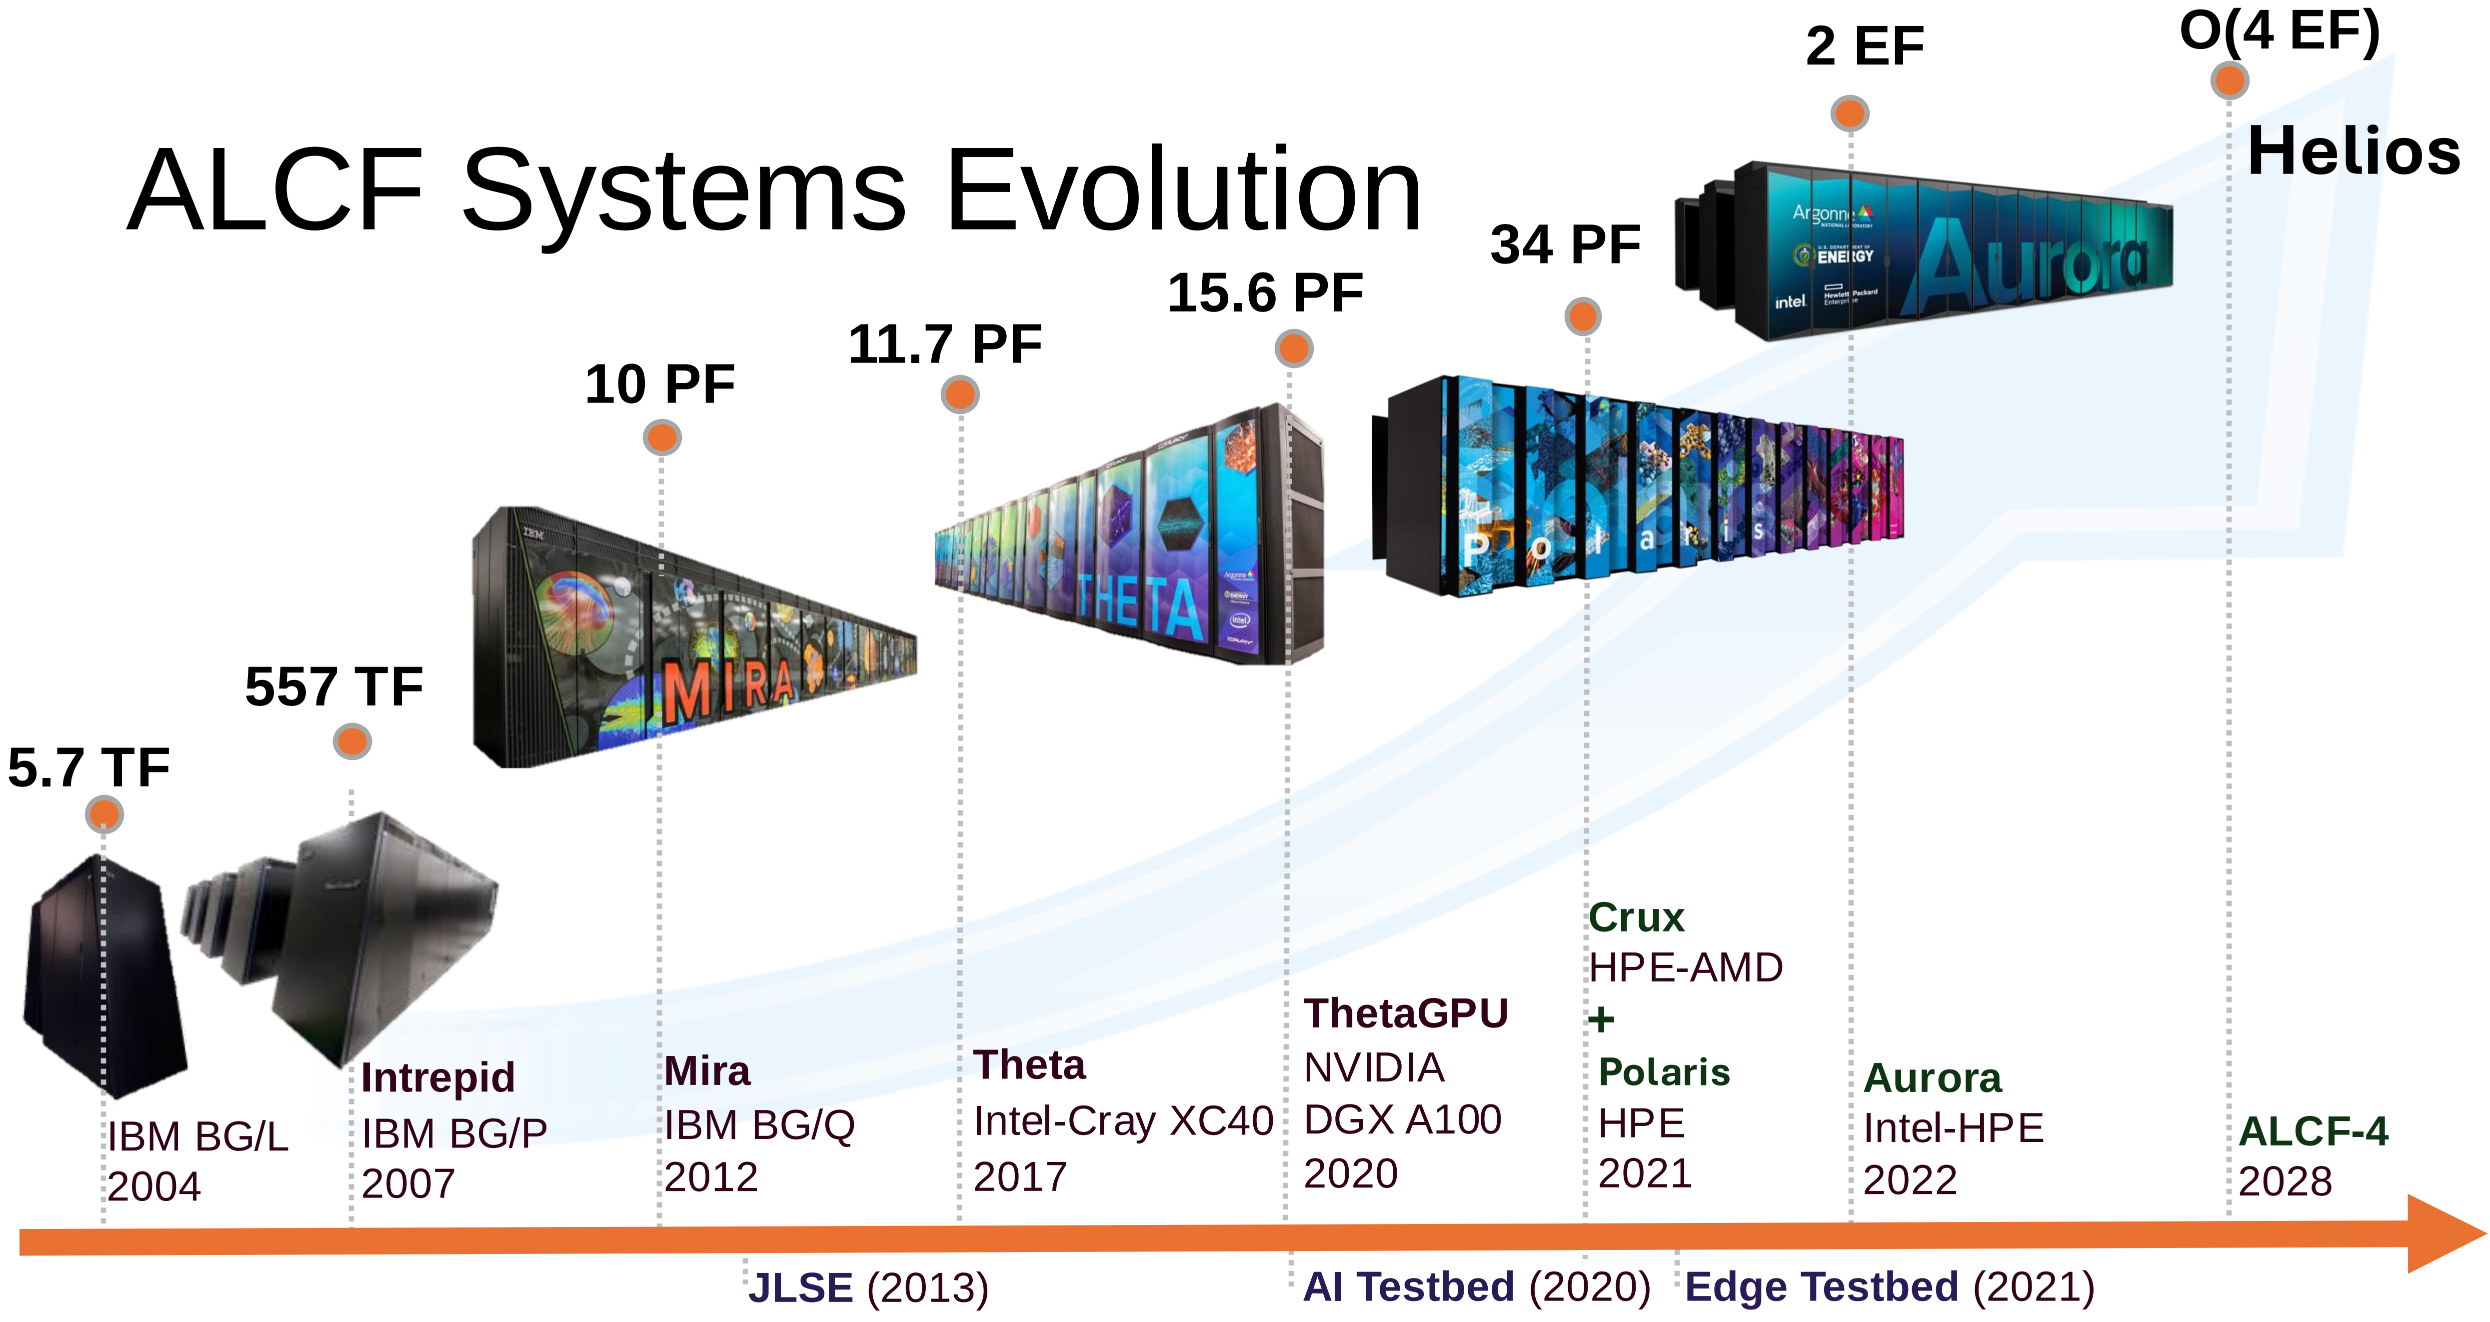

Aurora Specifications

Compute

21,248
CPUs

63,744
GPUs

10,624
Nodes

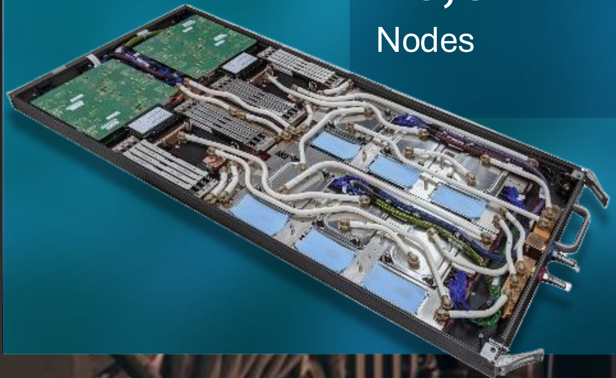

Fabric

Peak Injection Bandwidth
2.12 PB/s

Peak Bisection Bandwidth
0.69 PB/s

Dragonfly Topology

Memory

10.9PB DDR Capacity	1.36PB HBM CPU Capacity	8.16PB HBM GPU Capacity
5.95PB/s Peak DDR BW	30.5PB/s Peak HBM BW CPU	208.9PB/s Peak HBM BW GPU

Storage

230PB DAOS Capacity	31TB/s DAOS Bandwidth	1024 DAOS Node #

Computing Resources

Polaris
- HPE Apollo Gen10+
- AMD processors/NVIDIA GPUs
- 44 petaflops (double precision)
- NVIDIA GPU A100; HBM stack
- AMD EPYC Processor Milan
- 560 nodes

Sophia
- NVIDIA DGX A100
- 3.9 petaflops
- AMD EPYC 7742
- 24 nodes
- 24 TB of DDR4; 7.7 TB (HBM) of memory

JLSE Experimental Testbeds
- 150 nodes
- Intel/AMD/IBM/Marvell/GPGPU
- EDR/100GbE/OPA
- Lustre/GPFS/DAOS

Grand and Eagle
(Storage)
Each system has:
- HPE ClusterStor E1000
- 100 petabytes of usable capacity
- 8,480 disk drives
- Lustre filesystem
    - 160 Object Storage Targets
    - 40 Metadata Targets
- HDR InfiniBand network
- 650 GB/s rate on data transfers

ALCF AI Testbed

Cutting-edge AI accelerators for science

GroqRack (Available for Allocation Requests)

GroqRack Inference

System Size: 72 Accelerators (9 nodes x 8 Accelerators per node)

Compute Units per Accelerator: 5120 vector ALUs

Performance of a single accelerator (TFlops): >188 (FP16) >750 (INT8)

Software Stack Support: GroqWare SDK, ONNX

Interconnect: RealScale TM

Cerebras CS-2 (Available for Allocation Requests)

Cerebras CS-2 Wafer-Scale Cluster WSE-2

System Size: 2 Nodes (each with a Wafer scale engine) including Memory-X and Swarm-X

Compute Units per Accelerator: 850,000 Cores

Performance of a single accelerator (TFlops): >5780 (FP16)

Software Stack Support: Cerebras SDK, Tensorflow, Pytorch

Interconnect: Ethernet-based

SambaNova Dataflow (Available for Allocation Requests)

SambaNova DataScale SN30

System Size: 64 Accelerators (8 nodes and 8 accelerators per node)

Compute Units per Accelerator: 1280 Programmable compute units

Performance of a single accelerator (TFlops): >660 (BF16)

Software Stack Support: SambaFlow, Pytorch

Interconnect: Ethernet-based

Graphcore Bow Pod64 (Available for Allocation Requests)

Graphcore Intelligent Processing Unit (IPU)

System Size: 64 Accelerators (4 nodes x 16 Accelerators per node)

Compute Units per Accelerator: 1472 independent processing units

Performance of a single accelerator (TFlops): >250 (FP16)

Software Stack Support: PopArt, Tensorflow, Pytorch, ONNX

Interconnect: IPU Link

Habana Gaudi-1

Habana Gaudi Tensor Processing Cores

System Size: 16 Accelerators (2 nodes x 8 Accelerators per node)

Compute Units per Accelerator: 8 TPC + GEMM engine

Performance of a single accelerator (TFlops): >150 (FP16)

Software Stack Support: Synapse AI, TensorFlow and PyTorch

Interconnect: Ethernet-based

Community Data Sharing with Eagle

- A global filesystem deployed to bring larger and more capable production-level file sharing to facility users
- A space for broader distribution of reassembled data acquired from various experiments
  - Data originating at the ALCF
  - Greater scientific community
- Science community can access uploaded data, and ALCF users are able to directly access the data for analysis
- Designed to foster experimentation
  - Analysts are able to write new algorithms to attempt analyses that have never been performed

- **HPE ClusterStor E1000**
- **100 petabytes of usable capacity**
- **8,480 disk drives**
- **Lustre filesystem**
- **160 Object Storage Targets**
- **40 Metadata Targets**
- **HDR InfiniBand network**
- **650 GB/s rate on data transfers**

ALCF Allocation Programs

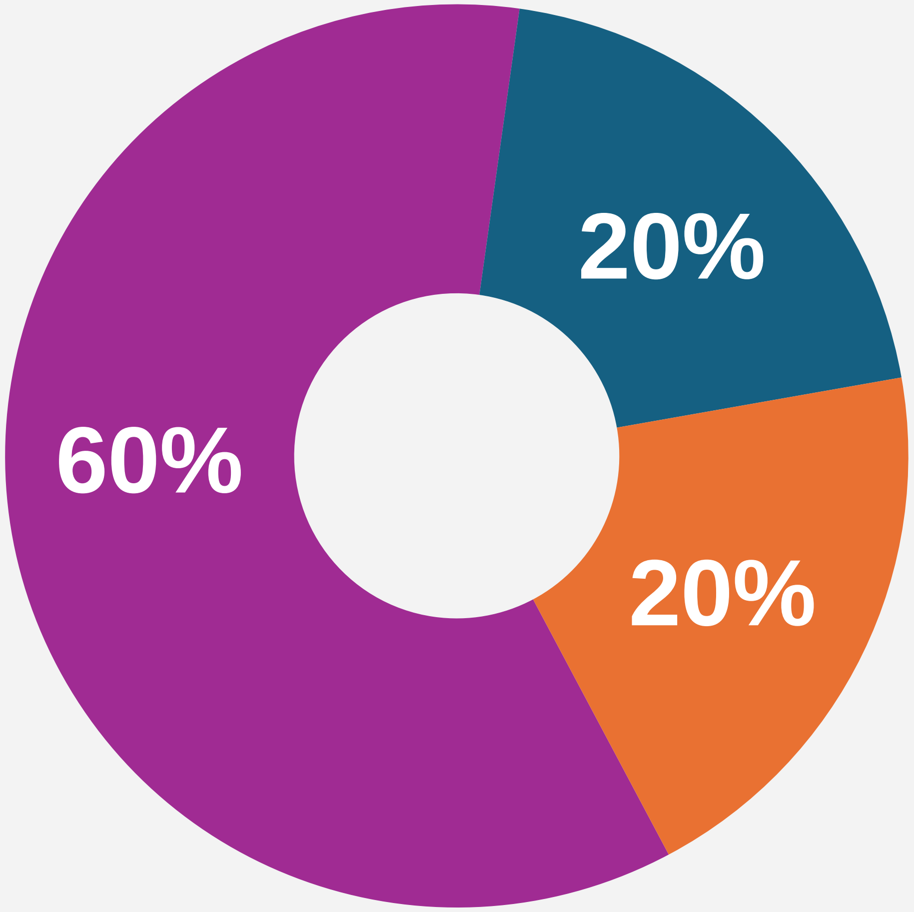

INCITE: Innovative and Novel Computational Impact on Theory and Experiment

- Yearly call with computational readiness and peer reviews
- Open to all domains and user communities

ALCC: ASCR Leadership Computing Challenge

- Yearly call with peer reviews
- Focused on DOE priority

DD: Director's Discretionary Program

- Rapid allocations for project prep and immediate needs
  - Early Science Program (ESP)
  - Exascale Computing Project (ECP)
  - ALCF Data Science Program (ADSP)
  - Proprietary Projects

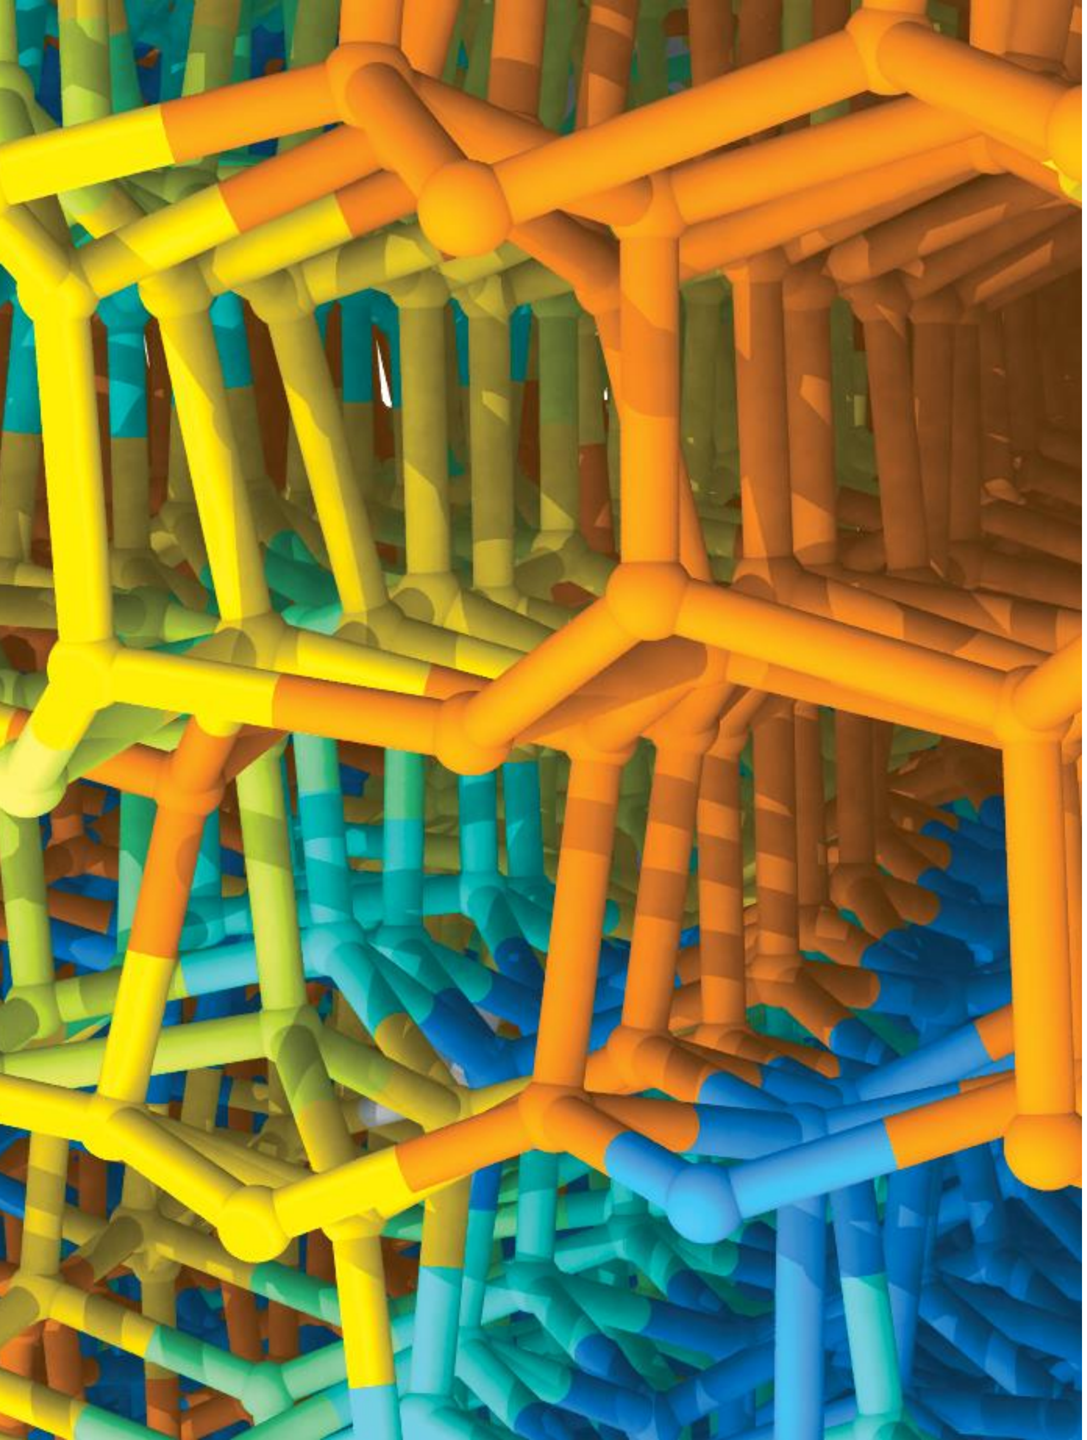

Director’s Discretionary

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.

- **Award Cycle:** Ongoing (available year-round)
- **Award size:** Up to several million compute-hours
- **Duration:** 3–6 months; renewable
- **Total percent of ALCF** resources allocated for DD projects: 20%

Getting Started

ALCF Lighthouse Initiative

Goal

- Broaden Argonne Leadership Computing Facility (ALCF) user base by cultivating enduring academic partnerships
- Extend high-performance computing (HPC) to a wider audience and engage the next generation of computing professionals

Partnered with UIC ACER

Key Activities

- Acquire knowledge of ALCF's HPC resources and computational science expertise
- Promote ALCF capabilities to faculty, staff, and students
- Identify potential university projects suited to ALCF resources
- Support researchers in pilot usage of ALCF systems
- Provide quarterly updates to ALCF, share best practices, and coordinate growth

https://acer.uic.edu/services/argonne-leadership-computing-facility-lighthouse-initiative/

Intensive two-week training on the key skills, approaches, and tools to design, implement, and execute Computational Science and Engineering (CSE) applications on current and next-generation supercomputers.

ATPESC 2024

JULY 28 – AUGUST 9

Deadline Extended:

03/16/2025

PROGRAM CURRICULUM

Renowned computer scientists and high-performance computing (HPC) experts from U.S. National Laboratories, universities, and industry serve as lecturers and effectively guide hands-on training sessions.

ATPESC participants will be granted access to **U.S. Department of Energy (DOE) Office of Science User Facilities,** which are home to some of the world's most powerful supercomputers.

The core curriculum includes:

- Computer architectures and predicted evolution
- Numerical algorithms and mathematical software
- Software productivity and sustainability
- Data analysis, visualization, I/O, and methodologies and tools for big data applications
- Performance measurement and debugging tools
- Machine learning and data science

COST

There are no fees to participate. Domestic airfare, meals, and lodging are provided.

ELIGIBILITY

Doctoral students, postdocs, and computational scientists are encouraged to submit applications. Visit the website for eligibility details.

APPLICATION

The program provides **advanced training to 70 participants.**

Qualified applicants must have:

- Substantial experience in MPI, OpenMP, and/or Data Science Frameworks,
- Used at least one HPC system for a complex application, and
- Plans to conduct CSE research on large-scale computers.

The call for applications for ATPESC 2024 opens January 2, 2024. For updates via email, please subscribe on our website.

SUBSCRIBE

For more information, visit:
extremecomputingtraining.anl.gov

CONTACT

Email: support@extremecomputingtraining.anl.gov

Argonne NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

Thank You