U.S. DEPARTMENT OF ENERGY | Office of Science

Advanced Scientific Computing Research (ASCR)

Hal Finkel

Presentation MODSIM 2023

ASCR – over 70 years of Advancing Computational Science

Beginnings: During the Manhattan Project, John Von Neumann advocated for the creation of a Mathematics program to support the continued development of applications of digital computing

ASCR has a rich history of investment in computational science and applied mathematics research, and revolutionary computational and network infrastructure.

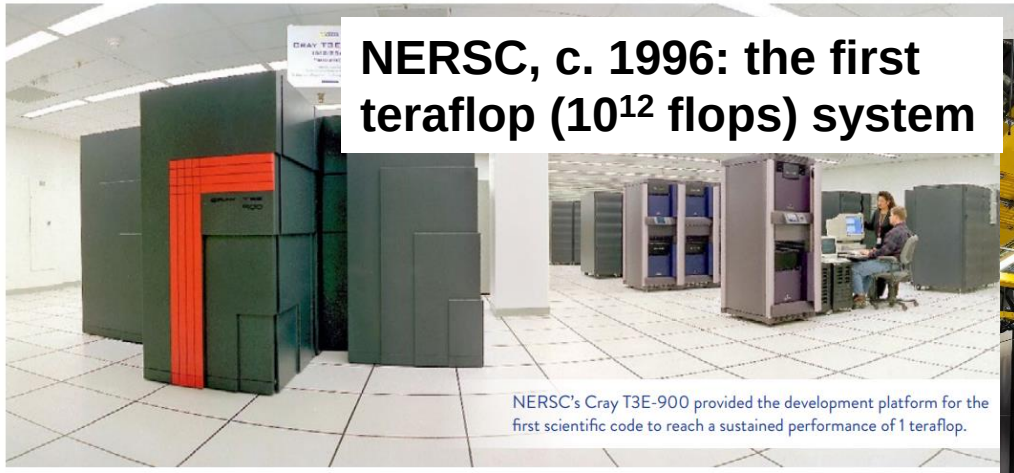

NERSC, c. 1996: the first teraflop (10^{12} flops) system

NERSC's Cray T3E-900 provided the development platform for the first scientific code to reach a sustained performance of 1 teraflop.

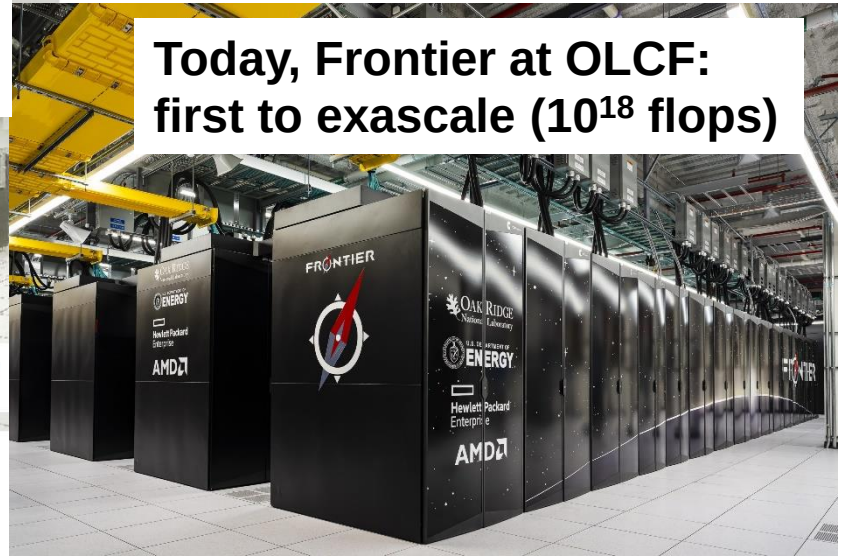

Today, Frontier at OLCF: first to exascale (10^{18} flops)

WHY COMPUTATIONAL SCIENCE?

- Computational science added a third pillar to researcher's toolkit along side theory and experiments
- Valuable when experiments are too expensive, dangerous, time-consuming or impossible
- Facilitates idea-to-discovery that leads from equations to algorithms
- Virtually every discipline in science and engineering has benefited from DOE's sustained investments in computational science

Frontier: An Efficient and Performant Exascale Supercomputer

TOP500

#1*

1.2 exaflops of performance on the TOP500 List.

*May 2023

GREEN500

#2*

62 gigaflops/watt power efficiency on a single cabinet.

*November 2022

HPL-MxP

#1*

10 exaflops on the HPL-MxP (formerly HPL-AI) benchmark.

*May 2023

- 74 HPE Cray EX cabinets
- 9,408 AMD EPYC CPUs, 37,632 AMD GPUs
- 700 petabytes of storage capacity, peak write speeds of 11 terabytes per second using Cray Clusterstor Storage System
- 90 miles of HPE Slingshot networking cables
- www.olcf.ornl.gov

Sources: May 30, 2022, and November 14, 2022, Top500 releases

DOE OLCF System Optimization with Digital Twins

What? A virtual supercomputer that can be used to simulate and optimize the scheduling, cooling, energy consumption, asset life, and even user behavior.

How? Uses a combination of complex simulations, AI/ML, and telemetry data to make autonomous decisions in real-time that realize value.

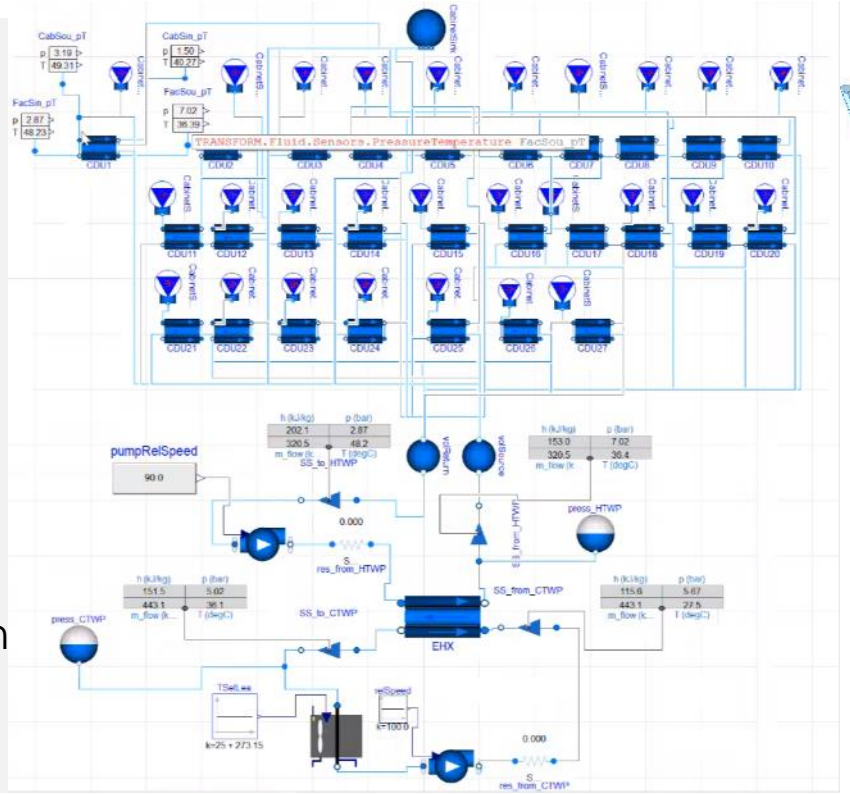

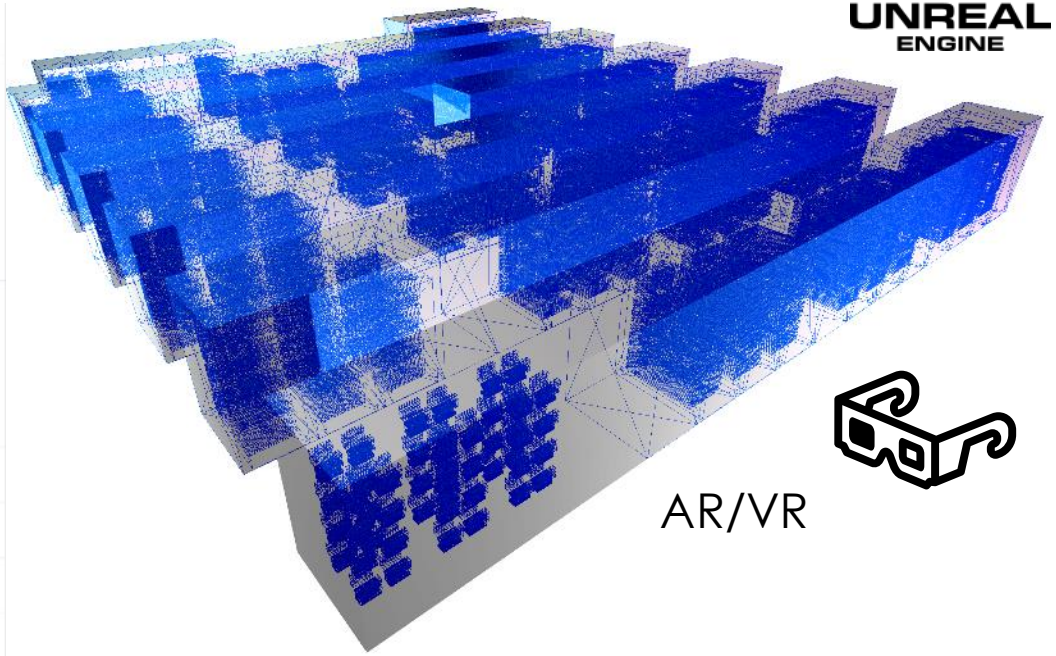

Slurm Sim

Simulated Job Scheduling

gem5

Simulated energy consumption

Modelica Language

Simulated Thermo-Fluids Cooling

Example energy efficiency use cases:

- Virtual prototyping – study design of future systems for energy efficiency and sustainability.
- Automated setpoint control for optimized cooling efficiency.
- Study optimized scheduling of jobs based on reduced energy consumption.

Energy-Efficient Computing - Public/Private Partnerships

- **Frontier, DOE's first exascale supercomputer, is #2 in the world on energy efficiency (on the Green500) in addition to being #1 in the world on performance (on the Top500).**
- **Frontier's performance and efficiency were enabled by the public/private R&D partnerships in the Exascale Computing Program (ECP)'s PathForward Program.**

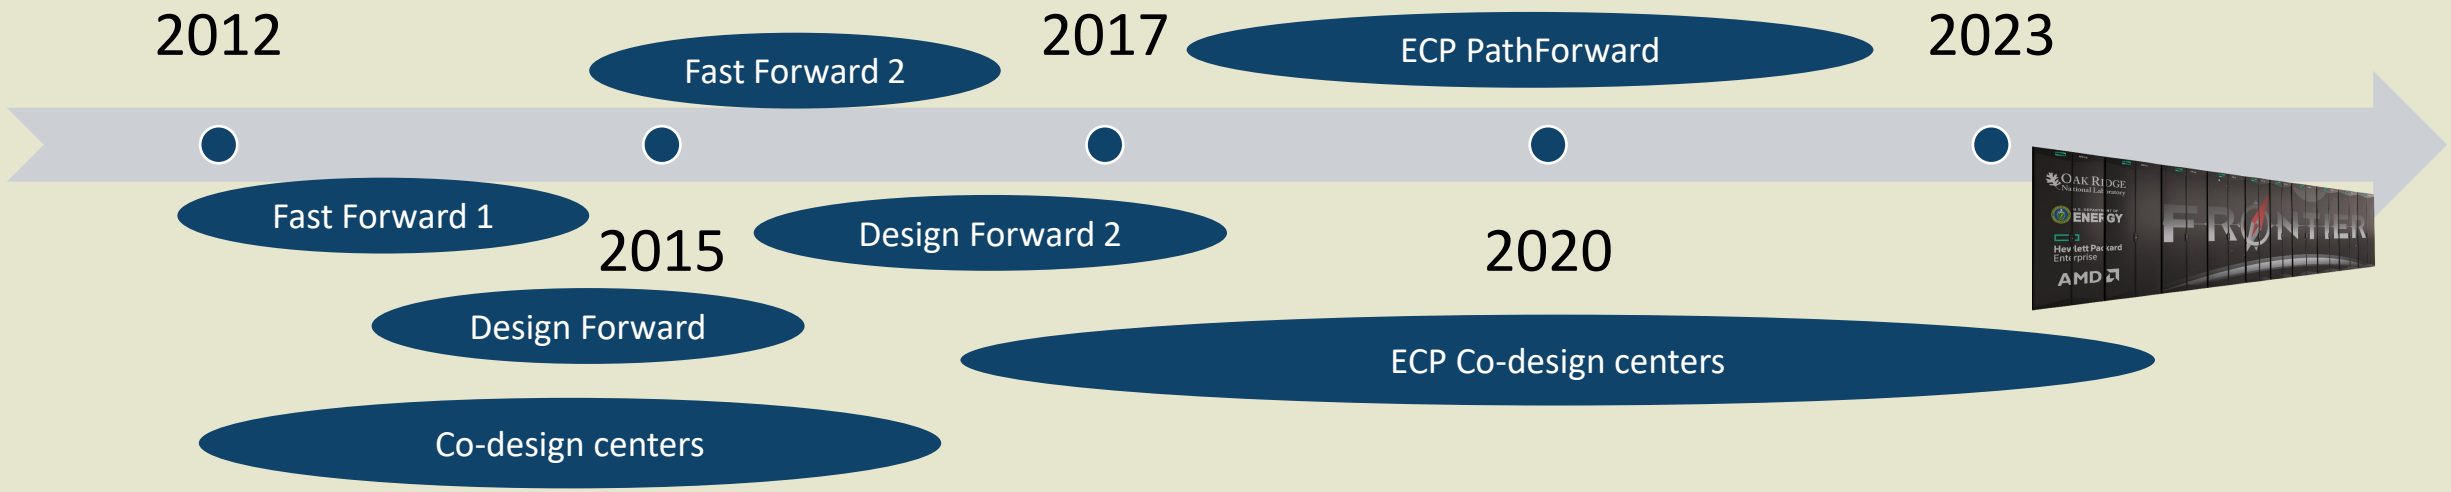

- **The ECP PathForward program built on a long history of DOE *Forward programs and co-design centers.**
  - **Partnered with AMD, Cray, HPE, IBM, Intel, NVIDIA on technologies related to compute, memory, networking, and storage.**
- **ASCR supports the Artificial Intelligence (AI) for Science (AMAIS) project at PNNL, an ongoing partnership with Micron. With ORNL and PNNL, ASCR has supported the DOE/Micron collaboration since 2017.**

Energy-Efficient Computing: A Critical National Challenge

- **SRC Decadal Plan (*) Computing Grand Goals:**
  - Discover computing paradigms/architectures with a radically new computing trajectory
  - Demonstrating > 1,000,000x improvement in energy efficiency
- **PCAST September 2022 Report on Revitalizing the US Semiconductor Ecosystem:**
  - Advanced Computing into the Zettascale Era (which requires breakthroughs in energy efficiency)
  - Energy-Efficient Computing via Domain-Specific Accelerators
- **CHIPS and Science Act 2022 Appropriated:**
  - National Semiconductor Technology Center (NSTC)
  - National Advanced Packaging Manufacturing Program (NAPMP)
- **NSTC and NAPMP will provide key capabilities to DOE's energy-efficient computing R&D, *and also*, will depend on DOE's user facilities and research pipeline.**

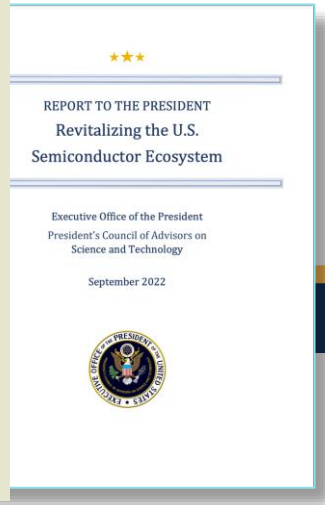

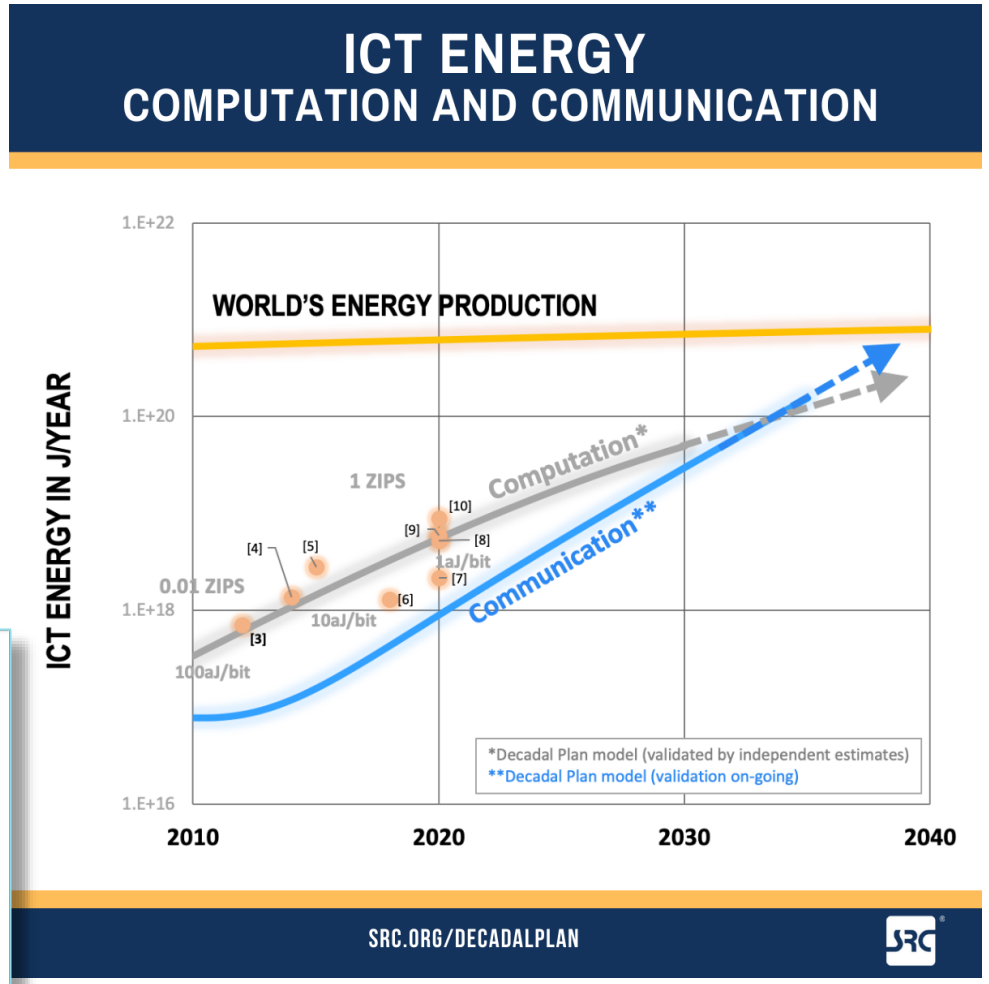

(*) DOE was a sponsor of the SRC Decadal Plan workshop series.

Energy-Efficient Computing – National Strategy Alignment

- **2020 National Science and Technology Council (NSTC) Advanced Computing Ecosystem Strategic Plan Objectives:**
  - **Ensure hardware leadership in a post-Moore/von Neumann**
  - **Advance software and software-hardware research**
  - **Address challenges and opportunities related to growing data volumes**
  - **Enhance AI capabilities**.
  - **Expand availability of and access to testbeds and prototyping**
  - **Address the need for technologies for hardware supply chain security**

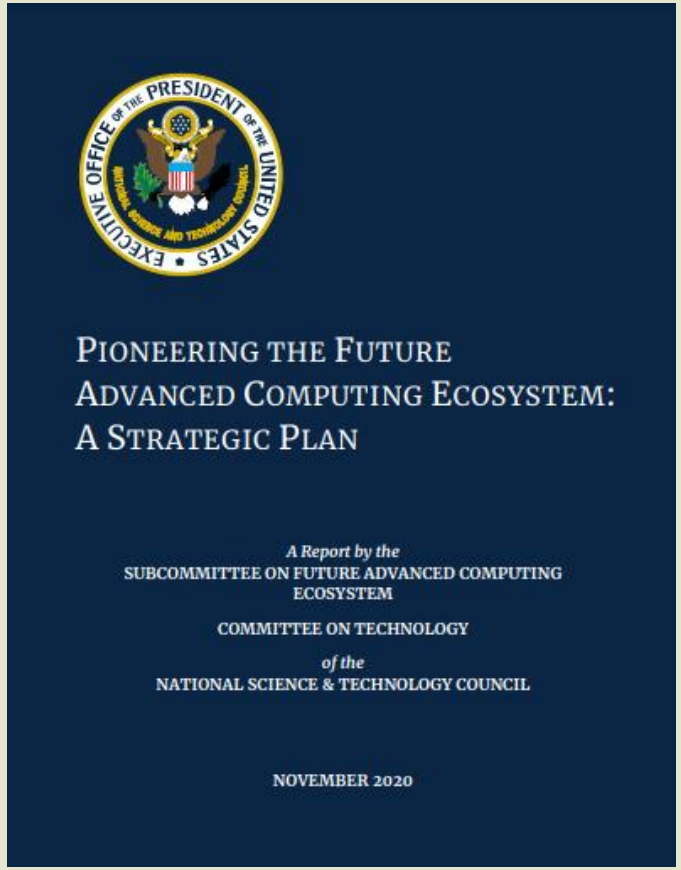

- **DOE's investments in energy-efficient computing help address all objectives.**

- **DOE is actively participating the NSTC Subcommittee on Microelectronics Leadership (SML) and the National Strategy on Microelectronics Research** (the draft strategy was released in September 2022 following by a public Request for Information and the strategy is now being finalized).

Critical Technology Trends Motivating ASCR Today

Data, Privacy, and Scientific Integrity

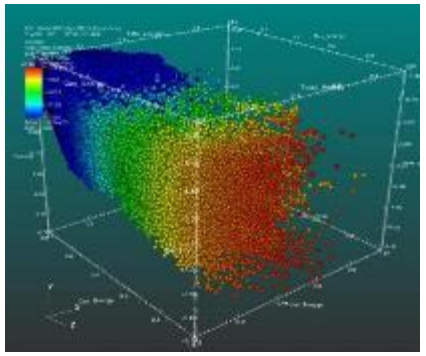

Artificial Intelligence and Deep Learning

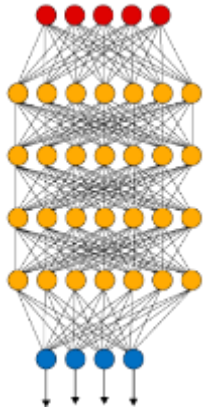

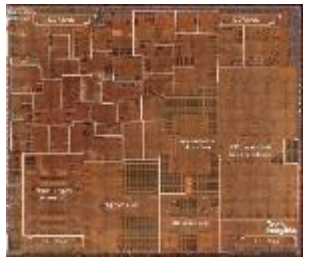

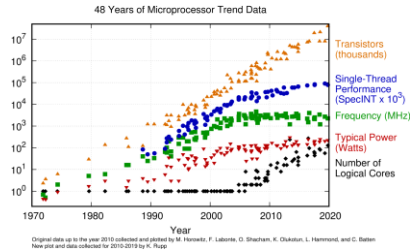

Heterogeneous, Distributed, Special-Purpose, Energy-Efficient Computing

Exploding Software Complexity

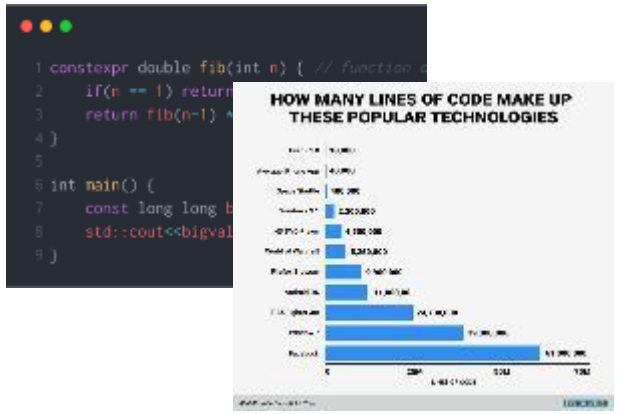

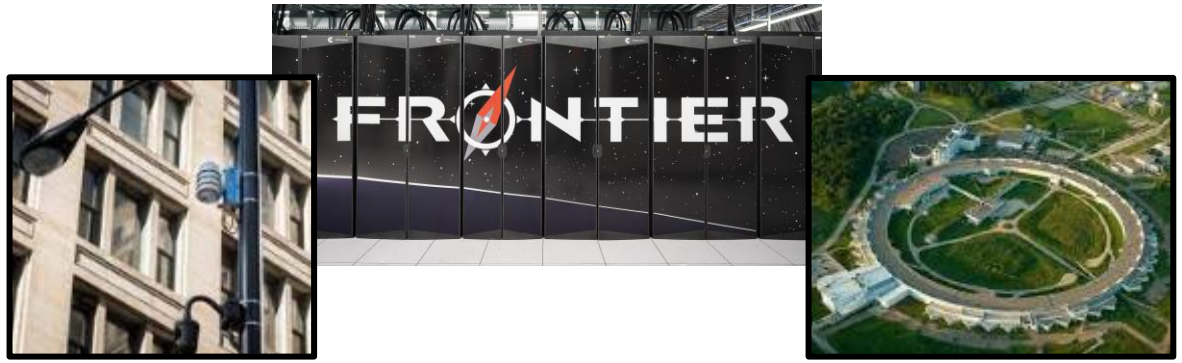

Scientific Computing and Networking: from HPC to the Edge

Interconnectivity and integration of instrumentation, data and computing are essential requirements for national R&D objectives

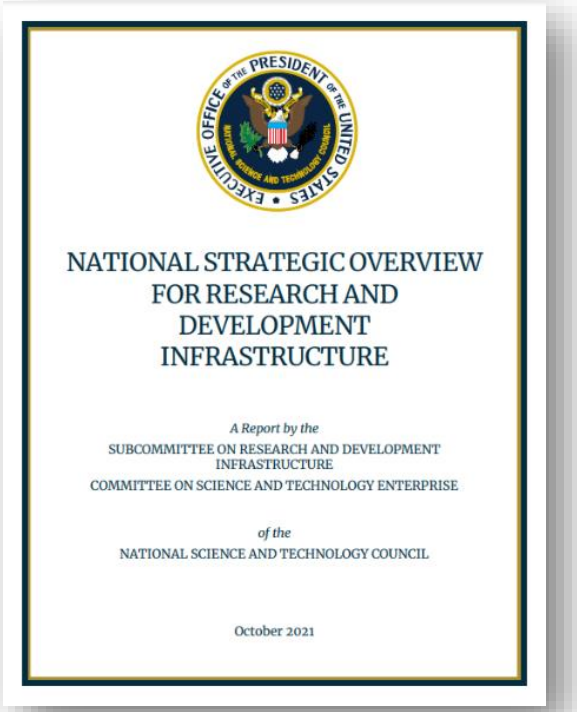
NATIONAL STRATEGIC OVERVIEW FOR RESEARCH AND DEVELOPMENT INFRASTRUCTURE

A Report by the
SUBCOMMITTEE ON RESEARCH AND DEVELOPMENT INFRASTRUCTURE
COMMITTEE ON SCIENCE AND TECHNOLOGY ENTERPRISE

of the
NATIONAL SCIENCE AND TECHNOLOGY COUNCIL

October 2021

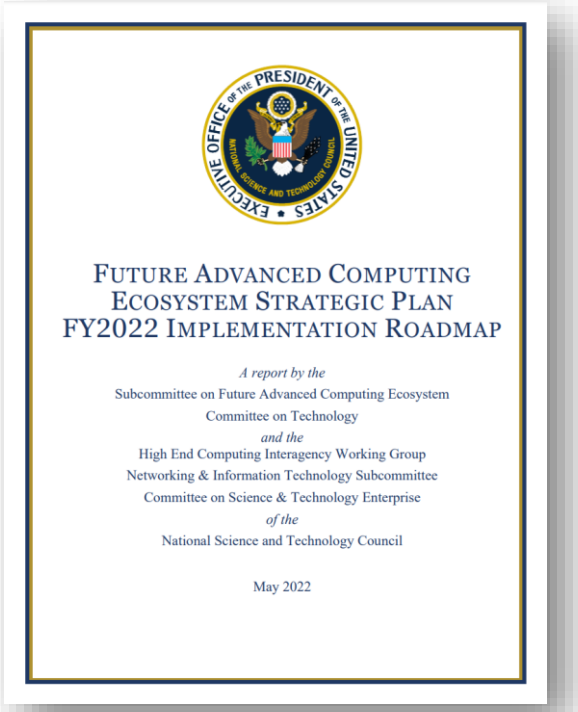
FUTURE ADVANCED COMPUTING ECOSYSTEM STRATEGIC PLAN FY2022 IMPLEMENTATION ROADMAP

A report by the
Subcommittee on Future Advanced Computing Ecosystem
Committee on Technology
and the
High End Computing Interagency Working Group
Networking & Information Technology Subcommittee
Committee on Science & Technology Enterprise
of the
National Science and Technology Council

May 2022

"R&D continues to shift from smaller to bigger science, driven in large part by advances in computing and other research cyberinfrastructure, which interlink[s] research data, analytics, … and experimental instrumentation."

~ 2021 National Strategic Overview of R&D Infrastructure

International Context

China Science and Technology Cloud
European Open Science Cloud
IRIS UKRI SFTC initiative

U.S. DEPARTMENT OF ENERGY | Office of Science

DOE's Integrated Research Infrastructure (IRI) is a new effort to provide researchers with seamless interoperability of DOE's unique data, user facilities, & computing resources.

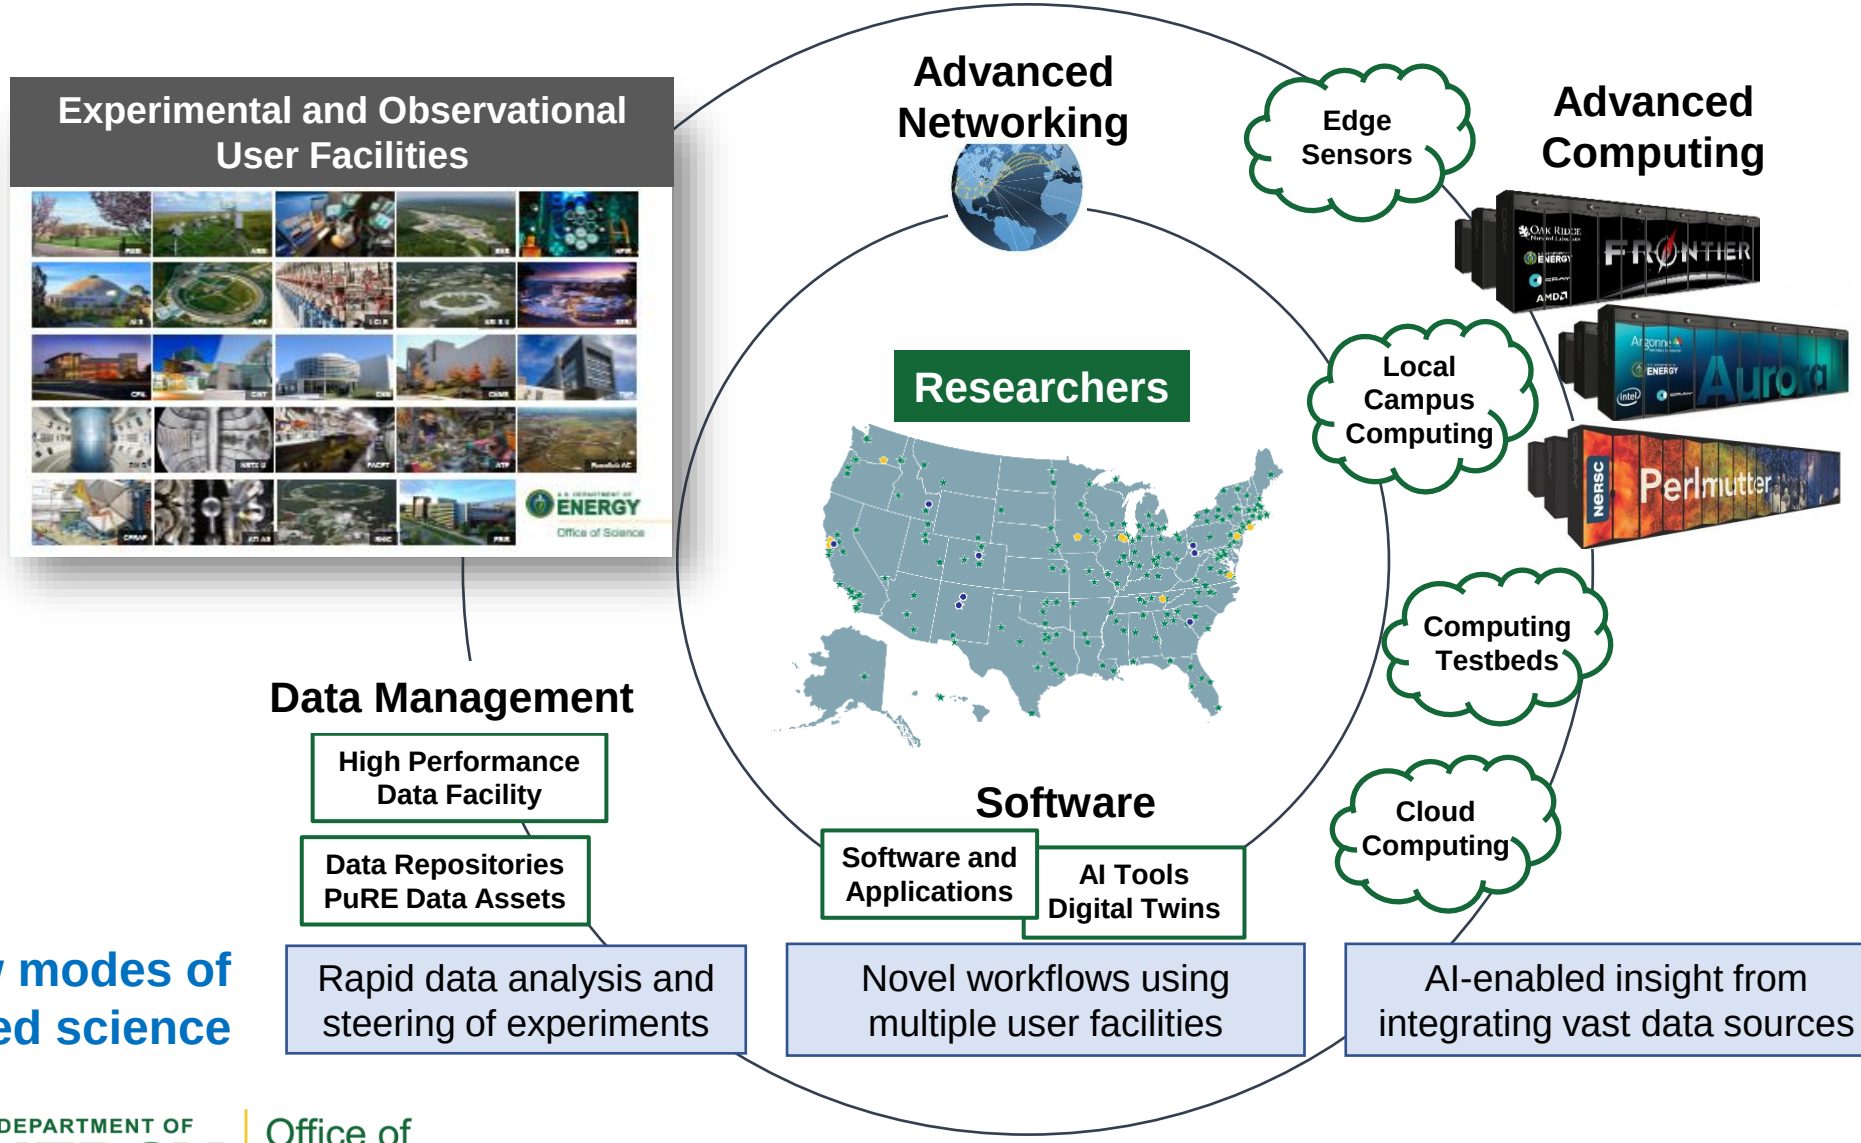

New modes of integrated science

U.S. DEPARTMENT OF ENERGY | Office of Science

Distributed Resilient Systems: DE-FOA-0002902 (FY23)

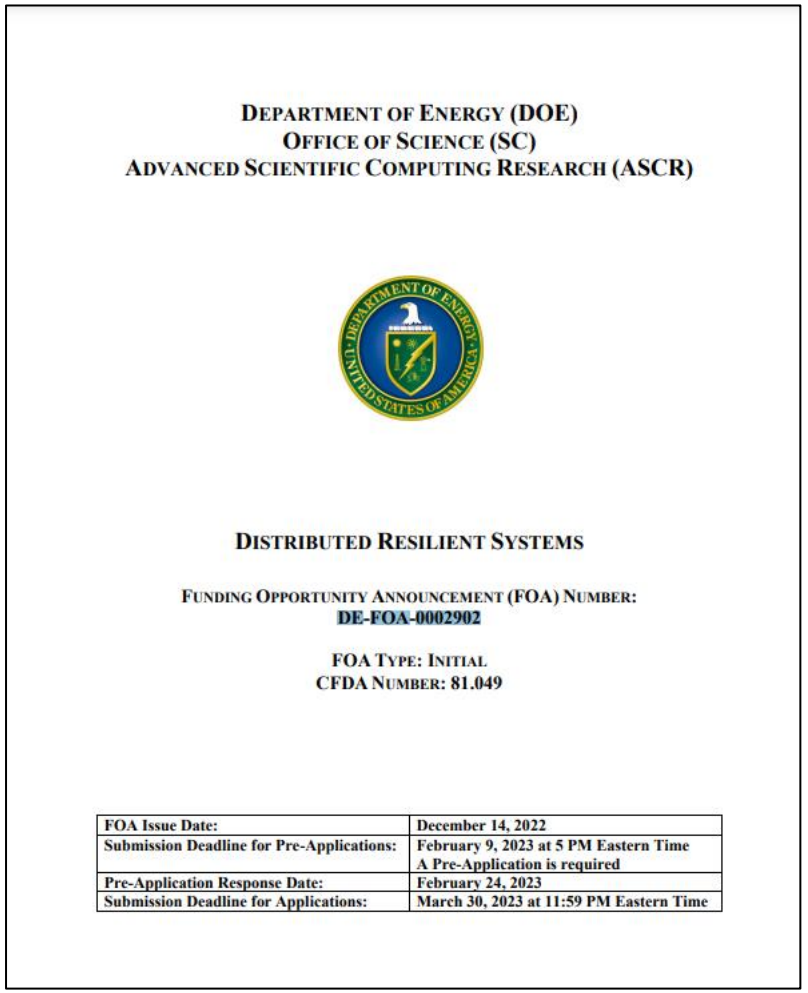

DEPARTMENT OF ENERGY (DOE)
OFFICE OF SCIENCE (SC)
ADVANCED SCIENTIFIC COMPUTING RESEARCH (ASCR)

DISTRIBUTED RESILIENT SYSTEMS

FUNDING OPPORTUNITY ANNOUNCEMENT (FOA) NUMBER:
DE-FOA-0002902

FOA TYPE: INITIAL
CFDA NUMBER: 81.049

FOA Issue Date:	December 14, 2022
Submission Deadline for Pre-Applications:	February 9, 2023 at 5 PM Eastern Time A Pre-Application is required
Pre-Application Response Date:	February 24, 2023
Submission Deadline for Applications:	March 30, 2023 at 11:59 PM Eastern Time

Research areas:

1. **Scalable system modeling**: The computational modeling of large, distributed systems is an important tool in understanding their behavior, both in normal operation and under anomalous conditions.
    - Required: describe the unknown system properties on which the model will provide insight, outline specific metrics and targets for those metrics reasonably necessary to provide the desired insight, and explain why it is reasonable to believe that the proposed approach, if successful, will reach or exceed those targets.
2. **Adaptive management and partitioning of resources**: Scheduling complex workflows on a single system, especially when the resources required for the workflow vary with time and/or the system provides heterogeneous resources, is a challenging endeavor. The scheduling and managing of resources across many systems for complex workflows provides even greater challenges.

Distributed Resilient Systems: DE-FOA-0002902 (FY23)

Diaspora: Resilience-enabling services for science from HPC to edge	Texas Tech University; Argonne National Laboratory (ANL); SLAC National Accelerator Laboratory; Oak Ridge National Laboratory (ORNL)
Exploring the Power of Distributed Intelligence for Resilient Scientific Workflows	Oak Ridge National Laboratory (ORNL); University of Southern California; Lawrence Berkeley National Laboratory (LBNL); The University of North Carolina at Chapel Hill; Argonne National Laboratory (ANL);
Resilient Federated Workflows in a Heterogeneous Computing Environment	Oak Ridge National Laboratory (ORNL); Brookhaven National Laboratory (BNL)
Scalable and Resilient Modeling for Federated-Learning-Based Complex Workflows	Argonne National Laboratory (ANL); University of Iowa; University of California
Tachyon: Intelligent Multi-Scale Modeling of Distributed Resilient Infrastructure and Workflows for Data Intensive HEP Analyses	Rensselaer Polytechnic Institute; Illinois Institute of Technology; University of California, Davis; Fermi National Accelerator Laboratory (FNAL); Argonne National Laboratory (ANL)

https://science.osti.gov/-/media/funding/pdf/Awards-Lists/2902-ASCR-Distributed-Resilient-Systems-Awards-List.pdf

Energy-Efficient Computing – Technology Diversity is Key

2020 NSTC Advanced Computing Ecosystem Strategic Plan and industry roadmaps all stress the importance of investing in a wide variety of promising techniques post Exascale.

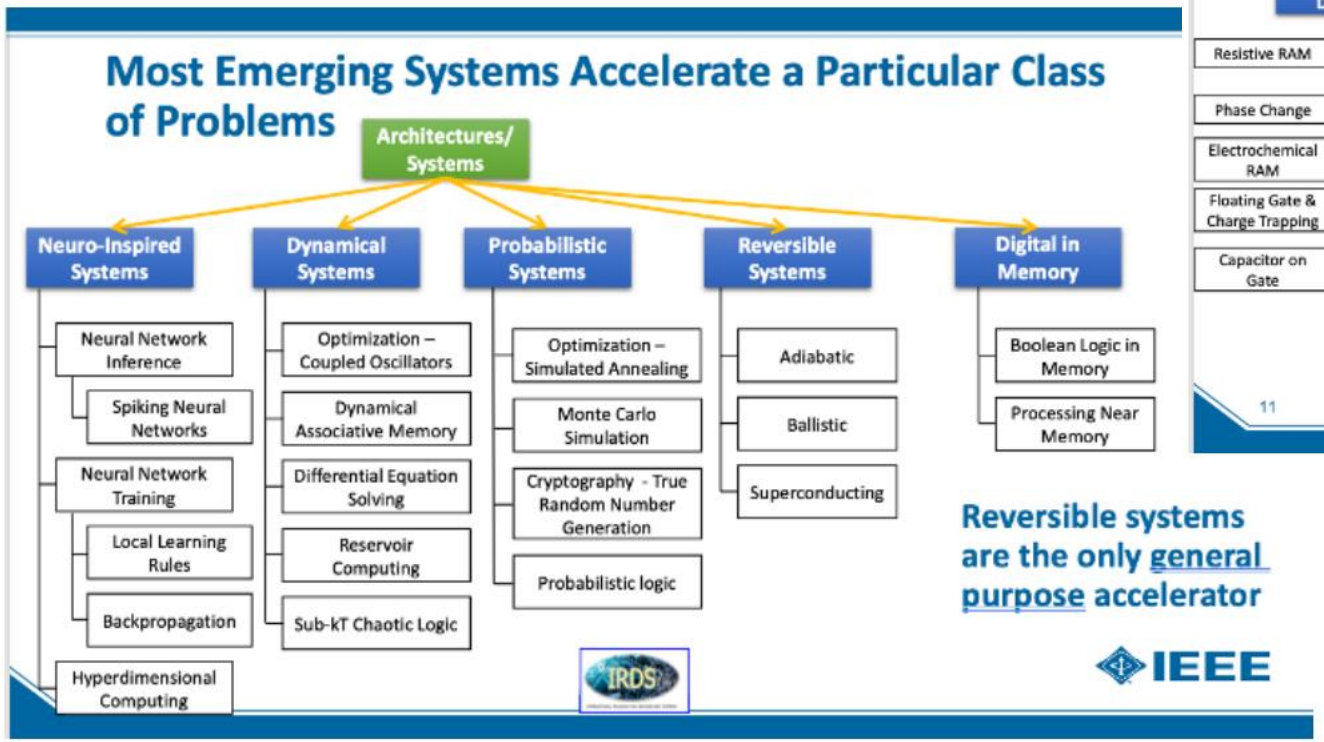

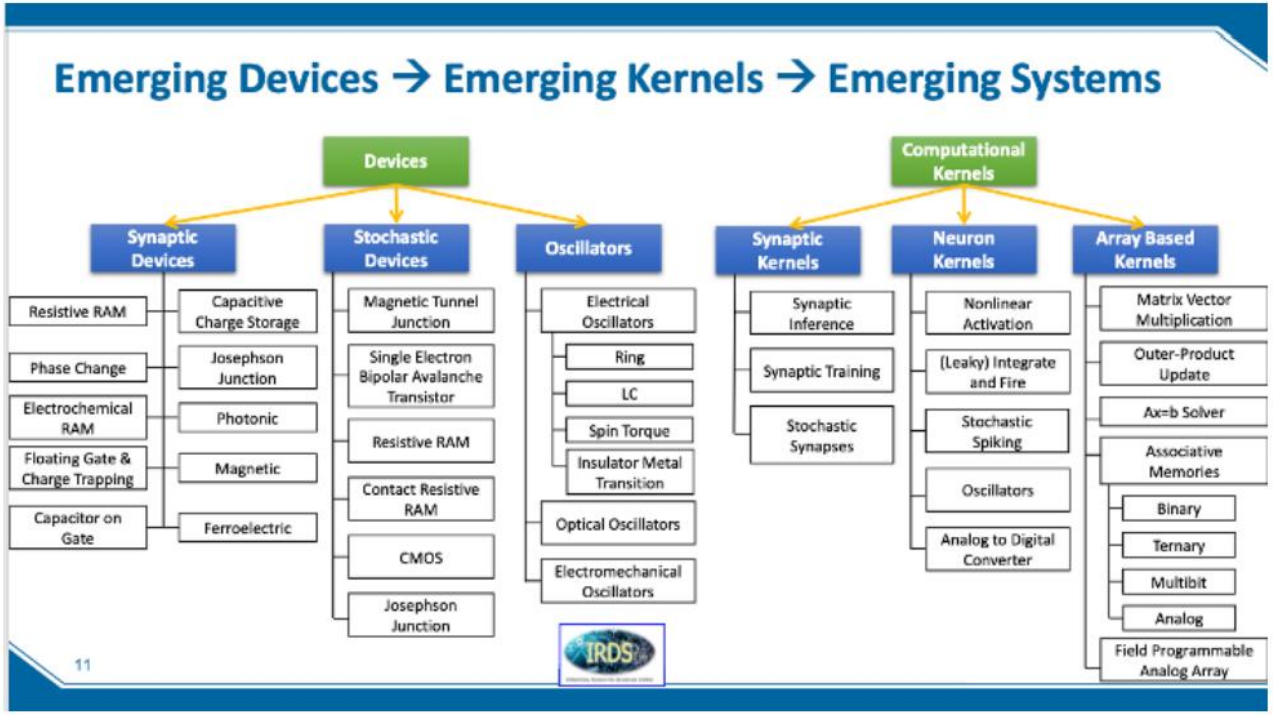

(THE INTERNATIONAL ROADMAP FOR DEVICES AND SYSTEMS: 2022; IEEE)

A holistic approach is needed, from device technologies and architectures, including both computing and memory/storage, to advanced packaging and cooling and data-center-level optimization.

Energy-Efficient Computing – Preparing for the Future

- **Basic Research Needs Workshop in Microelectronic, Extreme Heterogeneity, and Co-Design, identified priority research directions applicable to compute, memory, storage, and power electronics.**
- **2023 ASCR EXPRESS Funding Announcement has topics focused on modeling future computing systems, covering both digital computing and novel computing paradigms.**
- **2024 (Request), Microelectronics Research Centers (authorized in the CHIPS and Science Act).**

DOE Microelectronics Research Centers

- As authorized by the CHIPS and Science Act (Section 10731, Micro Act) the DOE FY 2024 President's Budget request includes $60M/year to establish **Microelectronics Science Research Centers.**
- Centers would perform mission-driven research to address foundational challenges in the design, development, characterization, prototyping, demonstration, and fabrication of microelectronics.
- **Centers would focus on fundamental science and early-stage research, complementing the investments already made through the CHIPS Act, such as:**
  - ❖ DOC National Semiconductor Technology Center (NSTC) and National Advanced Packaging Manufacturing Program (NAPMP)
  - ❖ DOD Microelectronics Commons
- Centers would leverage infrastructure and expertise at the DOE National Labs.

ASCR Workshop on Reimagining Codesign: Technology Targets

https://www.orau.gov/ASCR-CoDesign/

Codesign Technology Targets:

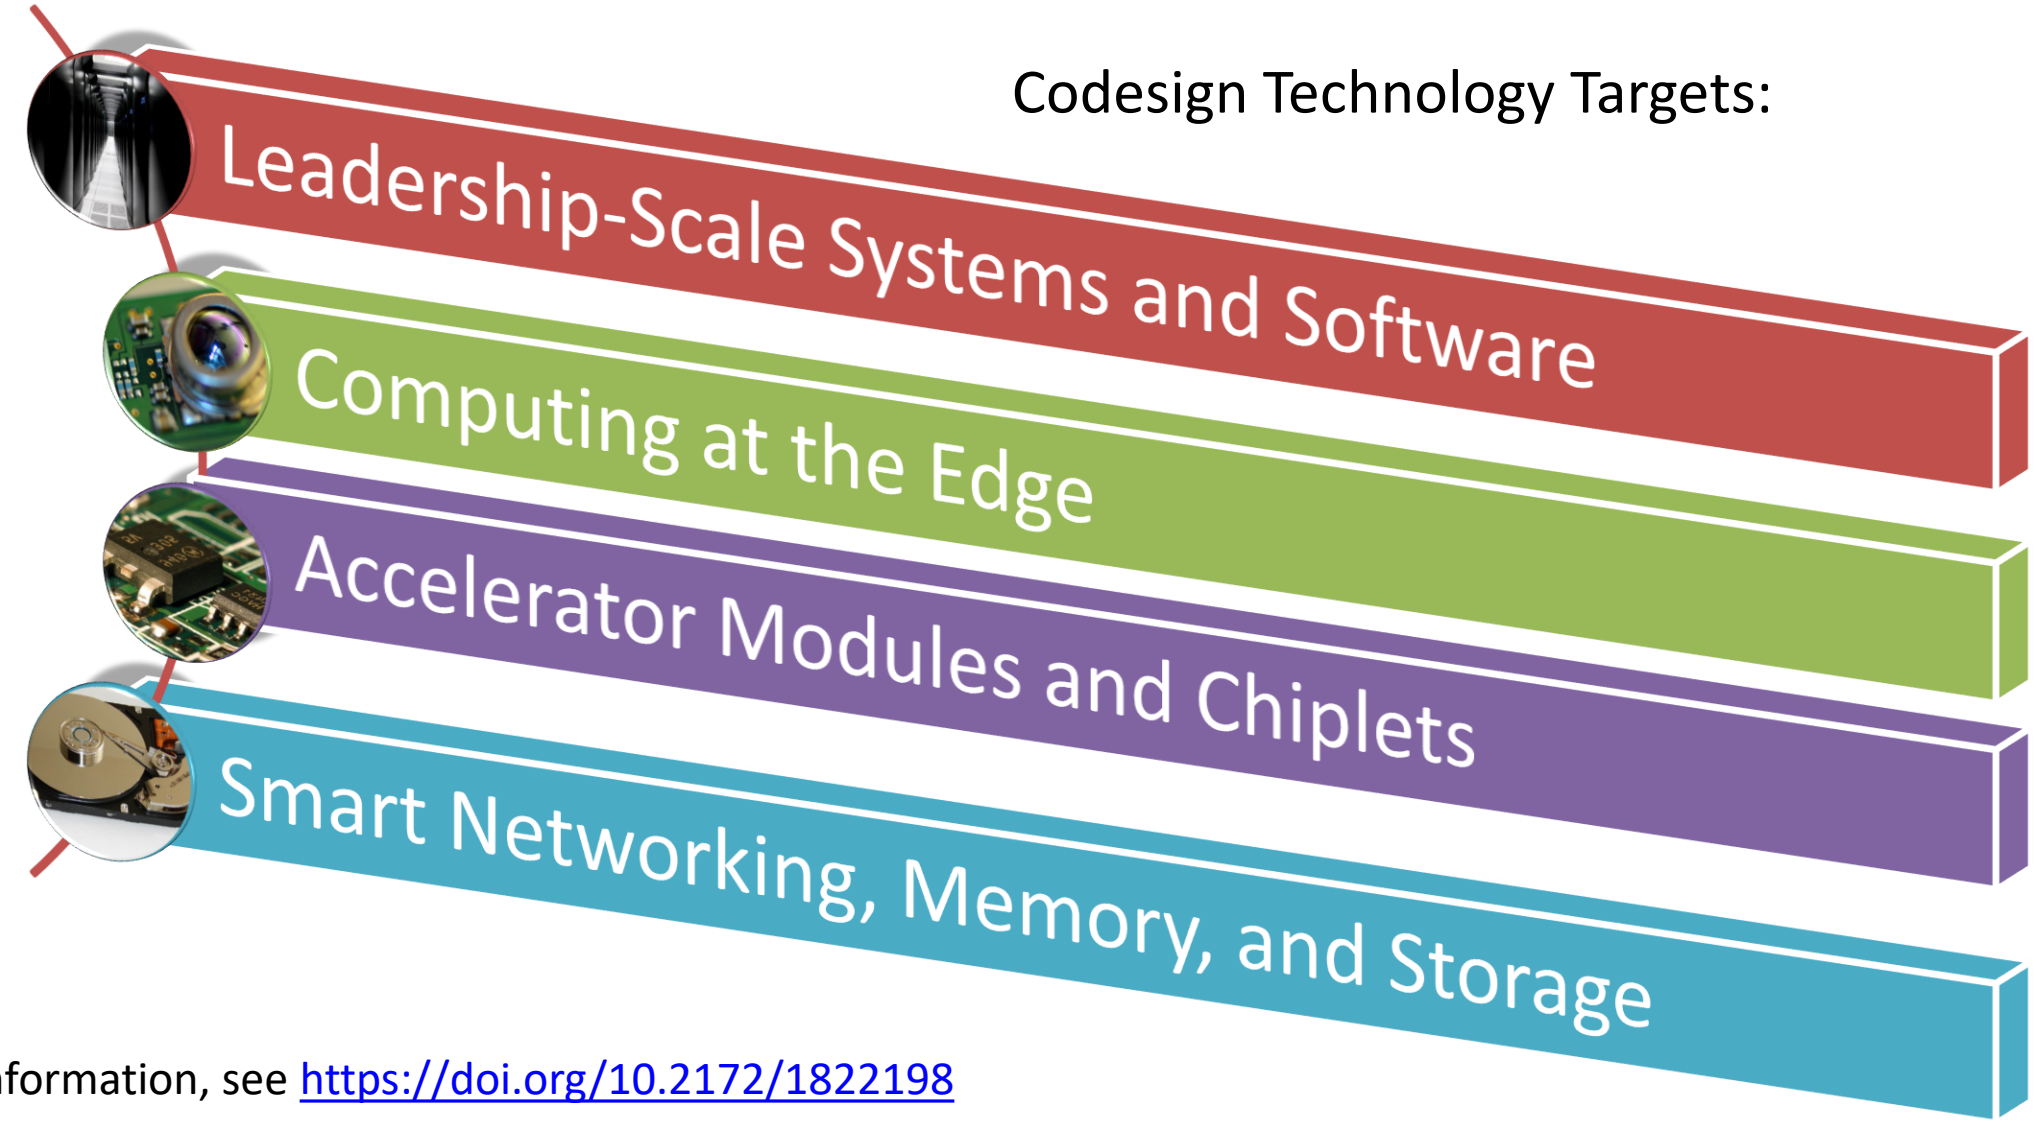

For more information, see https://doi.org/10.2172/1822198

ASCR Workshop on Reimagining Codesign: Enabling Technology Factors

https://www.orau.gov/ASCR-CoDesign/

Enabling Key Technology Factors:

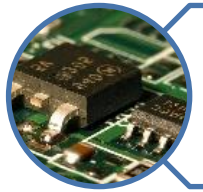

Advanced, modular packaging technologies providing for the high-performance composition of components optimized for different computational motifs, potentially from different organizations

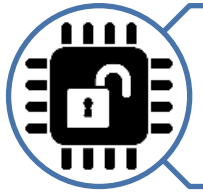

Open-source hardware designs allowing open, low-risk collaboration among academics, laboratories, and industry

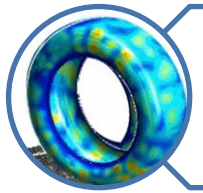

AI-driven technologies, paired with advanced system modeling, creating intelligent, data-driven workflows for hardware design and software development

Critical metrics for energy efficiency, security, and other system properties have joined performance, power usage, and reliability as first-class design constraints

The foundations laid in pursuit of exascale computing have generated applications capable of using first-generation heterogeneous GPU computing resources through **portable programming models and adaptive system software**

For more information, see https://doi.org/10.2172/1822198

ASCR Workshop on Reimagining Codesign: Past vs. Future

https://www.orau.gov/ASCR-CoDesign/

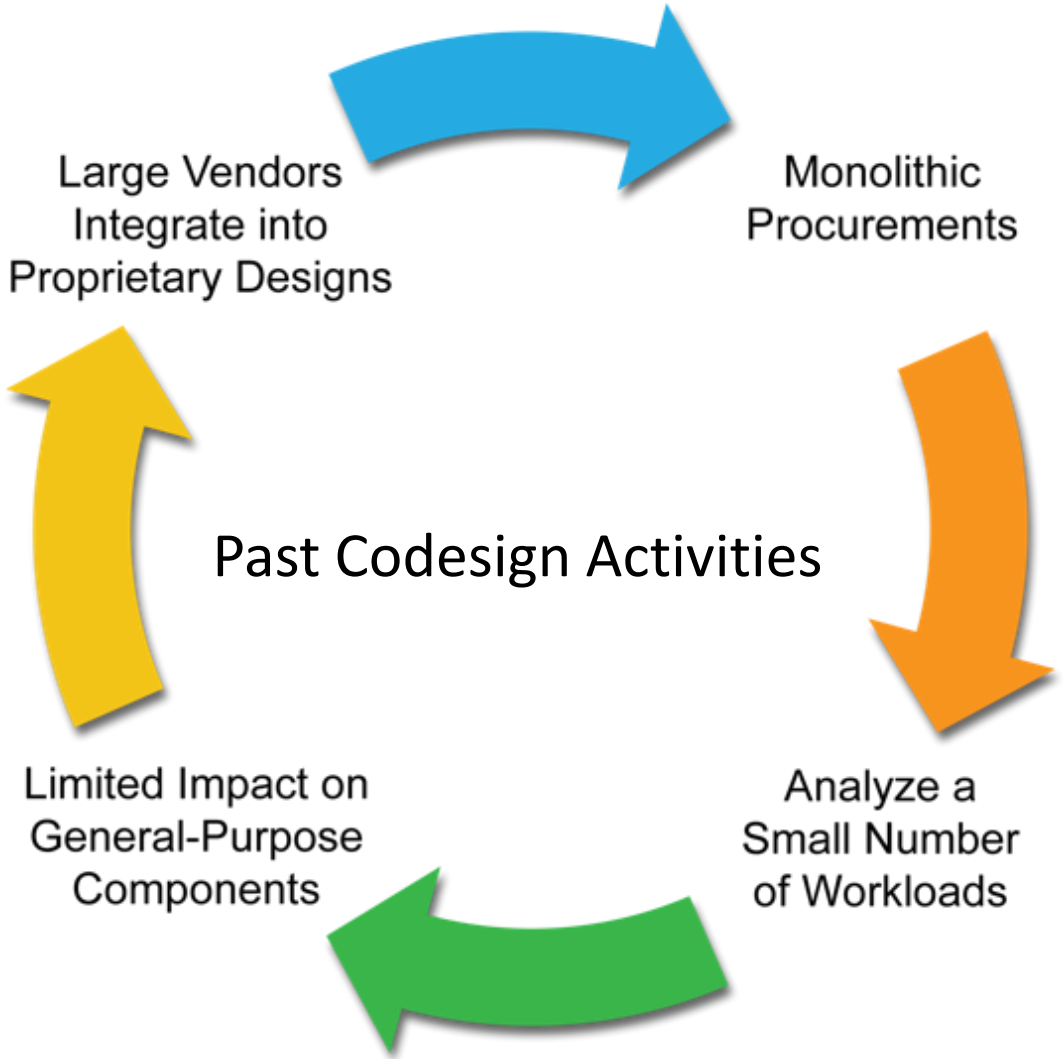

Reimagined Codesign:

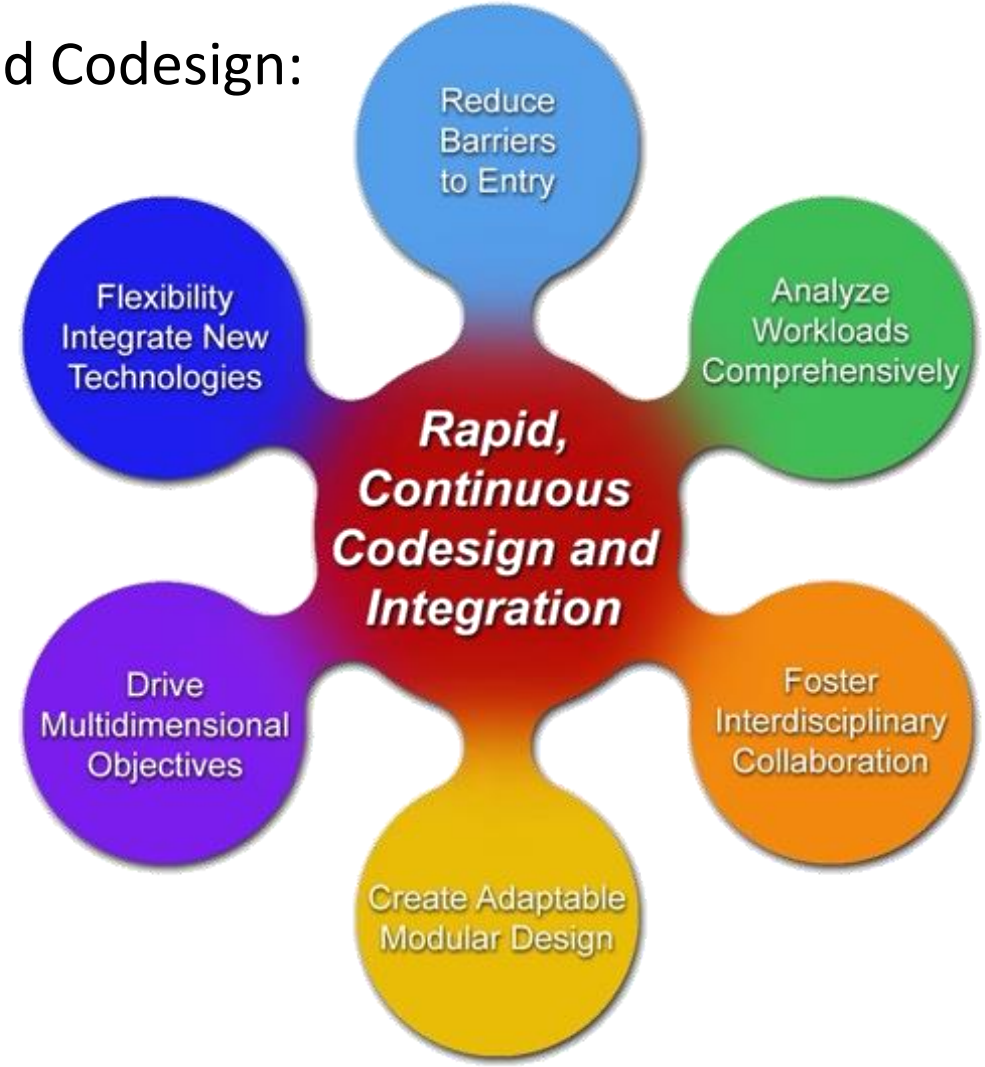

For more information, see https://doi.org/10.2172/1822198

EXPRESS: DE-FOA-0002950 (FY23)

https://science.osti.gov/-/media/funding/pdf/Awards-Lists/2950-ASCR-EXPRESS-Awards-List.pdf

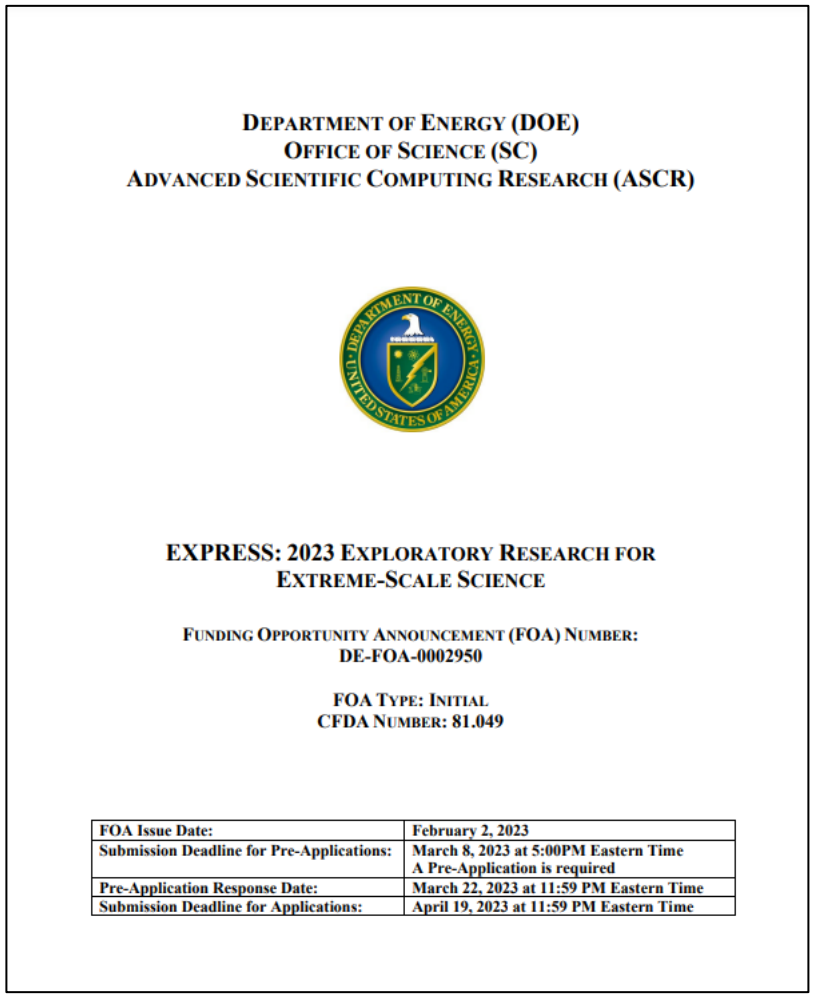

DEPARTMENT OF ENERGY (DOE)
OFFICE OF SCIENCE (SC)
ADVANCED SCIENTIFIC COMPUTING RESEARCH (ASCR)

EXPRESS: 2023 EXPLORATORY RESEARCH FOR EXTREME-SCALE SCIENCE

FUNDING OPPORTUNITY ANNOUNCEMENT (FOA) NUMBER:
DE-FOA-0002950

FOA TYPE: INITIAL
CFDA NUMBER: 81.049

FOA Issue Date:	February 2, 2023
Submission Deadline for Pre-Applications:	March 8, 2023 at 5:00PM Eastern Time A Pre-Application is required
Pre-Application Response Date:	March 22, 2023 at 11:59 PM Eastern Time
Submission Deadline for Applications:	April 19, 2023 at 11:59 PM Eastern Time

Research areas:

1. **Modeling Future Supercomputing Systems:**
   - To model supercomputing systems providing at least a serial computational clock rate of 20 GHz and that might plausibly be capable of providing at least 100 double-precision exaFLOPS (FLoating point Operations per Second) while consuming no more than 20 megawatts (including the power for cooling)
   - The modeling must account for the relatively high latencies associated with communication and memory access that are implied by the higher clock rates.
2. **Programming Techniques for Computational Physical Systems**:
   - To advance our understanding of how complex physical systems can be created whose evolution in time is described by a system of differential equations where those differential equations are produced by compilation of an expressive programming language.
   - A large number of potential uses cases are potentially relevant, from molecular self-assembly to data-storage technologies (including molecular/DNA storage).

U.S. DEPARTMENT OF ENERGY | Office of Science