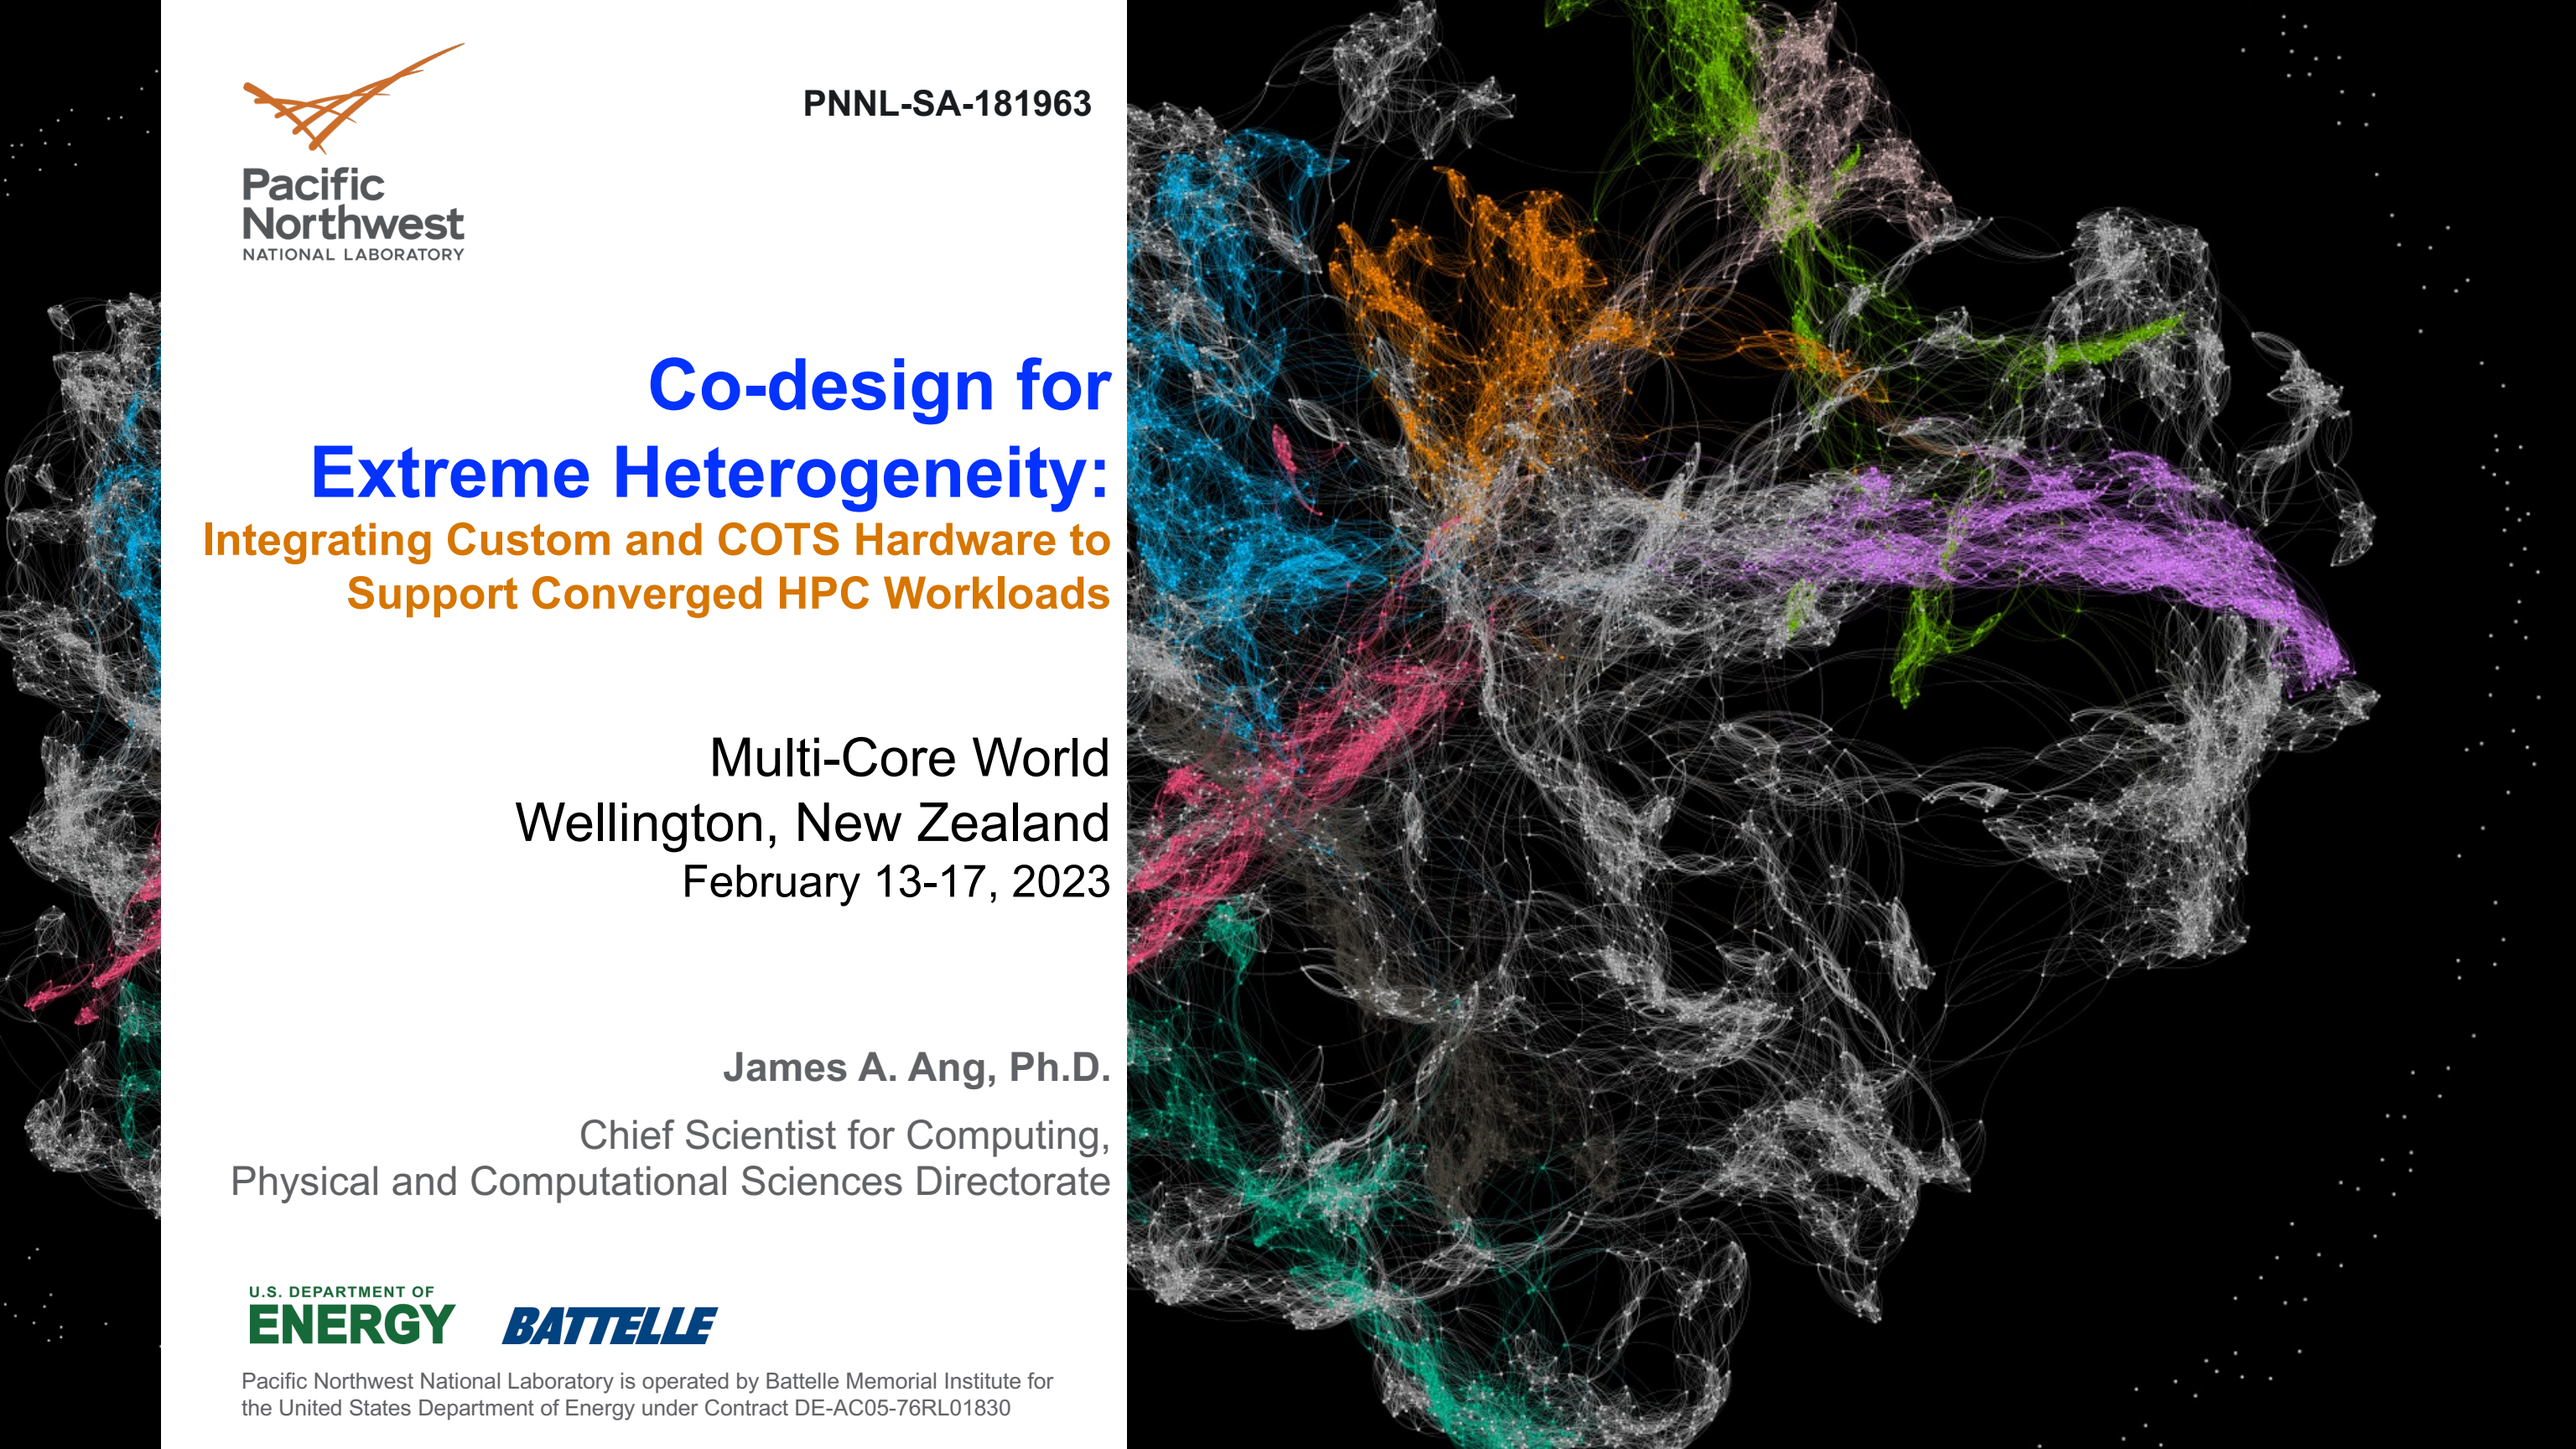

Pacific Northwest
NATIONAL LABORATORY

PNNL-SA-181963

# Co-design for Extreme Heterogeneity:

## Integrating Custom and COTS Hardware to Support Converged HPC Workloads

Multi-Core World
Wellington, New Zealand
February 13-17, 2023

**James A. Ang, Ph.D.**

Chief Scientist for Computing,
Physical and Computational Sciences Directorate

U.S. DEPARTMENT OF ENERGY BATTELLE

Pacific Northwest National Laboratory is operated by Battelle Memorial Institute for the United States Department of Energy under Contract DE-AC05-76RL01830

# New Horizons for HPC

- The Path to Exascale
- The Next Grand Challenge for Computing
- The CHIPS and Science Act of 2022
- PNNL's Data-Model Convergence Initiative
  - Co-design of custom reconfigurable accelerators for heterogeneous computing
  - Integrate with COTS CPU and GPU chiplets
- Continuum computing and metrics

# The Path to Exascale

- Over Two Decades of MPPs
- Moore's Law limits
- End of Dennard Scaling


Top500 #1: Frontier, USA
Performance: 1,102 PFLOPS (Rmax)

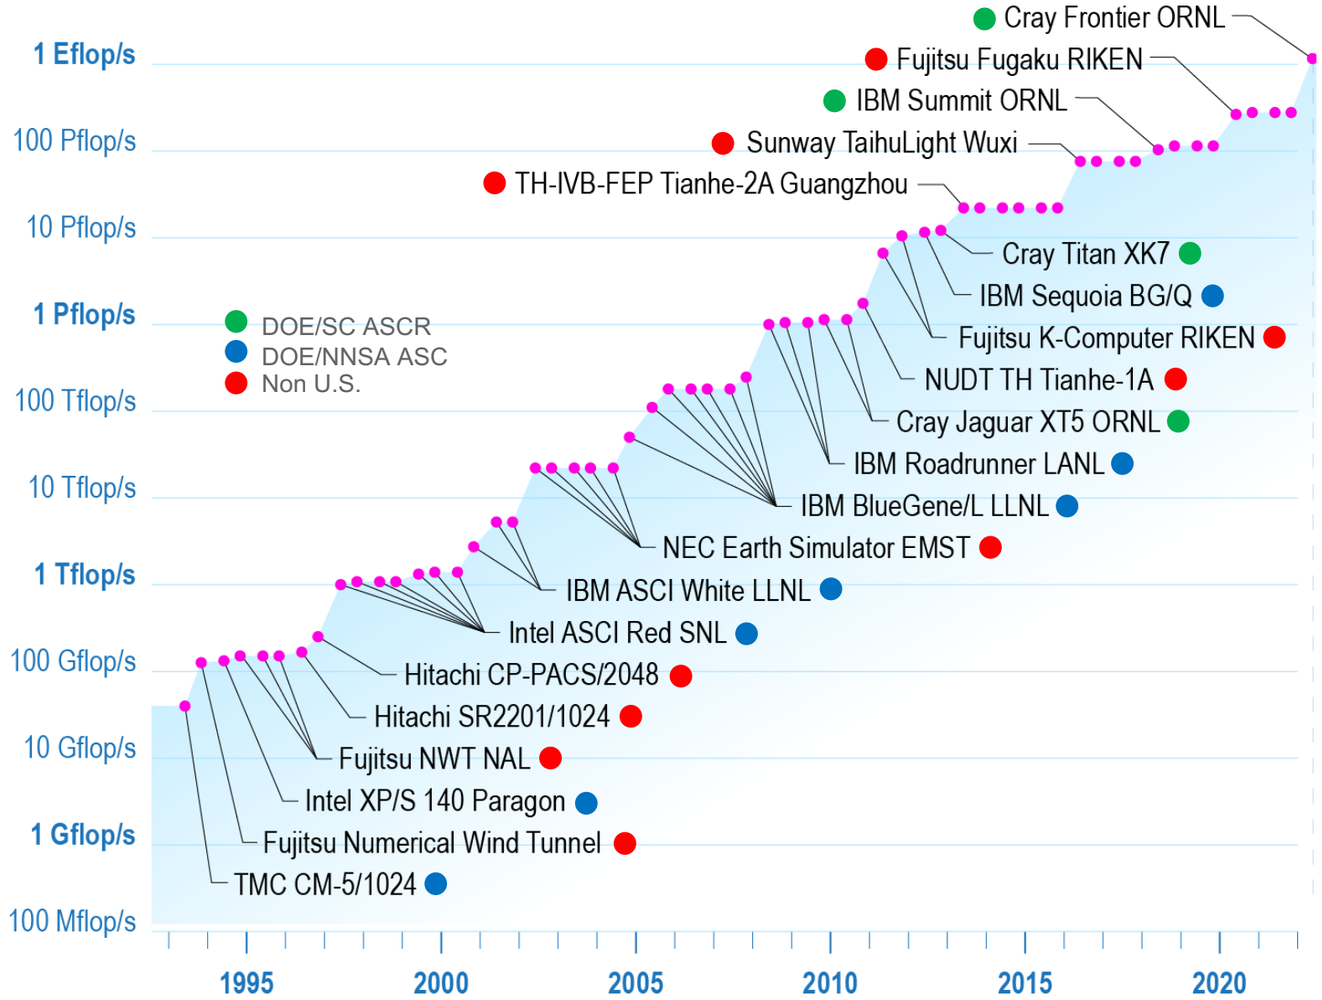


Exascale has been driven by simulation

# Milestone Supercomputers

Milestone supercomputers have required significant increases in power

- TeraFLOP: June 1997, ASCI Red at SNL with 0.85 MW power
- PetaFLOP: May 2008, Roadrunner at LANL with 2.35 MW power
- ExaFLOP: June 2022, Frontier at ORNL with 21.1 MW power
- Future: ZettaFLOPS (?)

# Zettascale: Next Grand Challenge for HPC (?)

- PCAST* Recommendation 9: *NSTC should identify a set of nationwide grand challenges that are enabled through collaboration across the NSTC industrial membership and NSTC-funded research*
- Grand Challenges should span three complimentary areas that benefit from large-scale nationwide collaboration:
  - 9(a) Advanced computing into the zettascale era
  - Significantly reducing design complexity
  - Proliferating semiconductors in life sciences applications

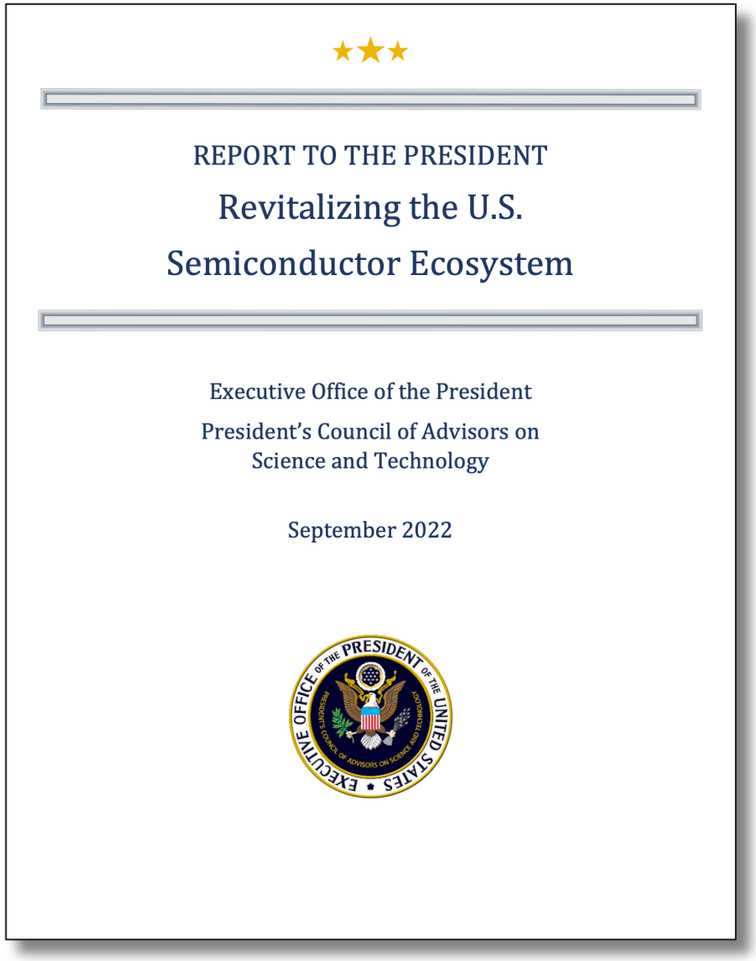


*PCAST Report: Revitalizing the U.S. Semiconductor Ecosystem
https://www.whitehouse.gov/wp-content/uploads/2022/09/PCAST_Semiconductors-Report_Sep2022.pdf

# Next Grand Challenge for HPC:
## Energy Efficient Computing

- PCAST Recommendation 8: The Secretary of Commerce should ensure that the NSTC founding charter allocates a significant portion of the annual funding, on the order of 30-50 percent, to directly fund a national research agenda
- This research agenda should be broad in nature and address the following areas:
  - materials, process, and manufacturing technologies;
  - packaging and interconnect technologies;
  - 8(c) energy-efficient computing and domain-specific accelerators;
  - design automation tools and methods;
  - semiconductor and system security; and
  - semiconductors and life sciences.

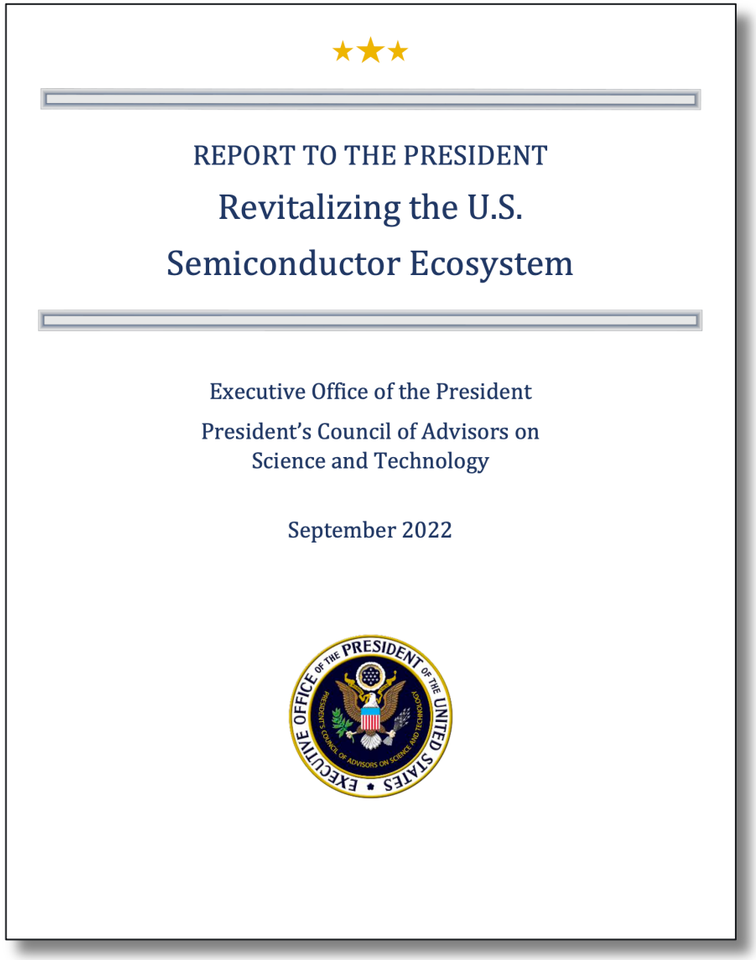

# Next Grand Challenge for Advanced Computing:
## 1,000,000x improvement in energy efficiency

- SRC Decadal Plan Computing Grand Goal
  - Discover computing paradigms/architectures with a radically new computing trajectory
  - Demonstrating > 1,000,000x improvement in energy efficiency

https://www.src.org/about/decadal-plan/

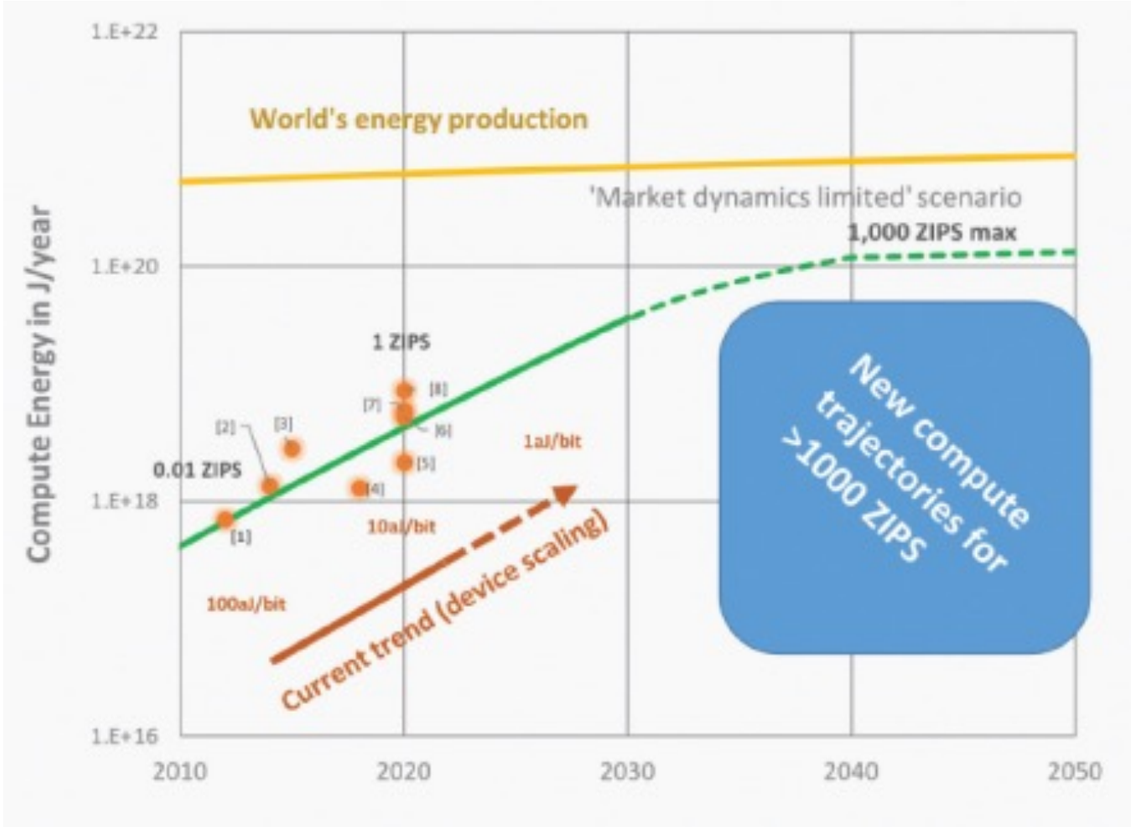


*Total energy of computing.*
*The solid yellow line indicates current energy generation trajectory*

# NIST CHIPS Act R&D Vision


https://www.nist.gov/chips

See *Webinars* and
*Industrial Advisory Committee*

# NIST CHIPS Act R&D Vision

## Update on CHIPS Research and Development

Eric Lin
Interim Director
CHIPS Research and Development Office

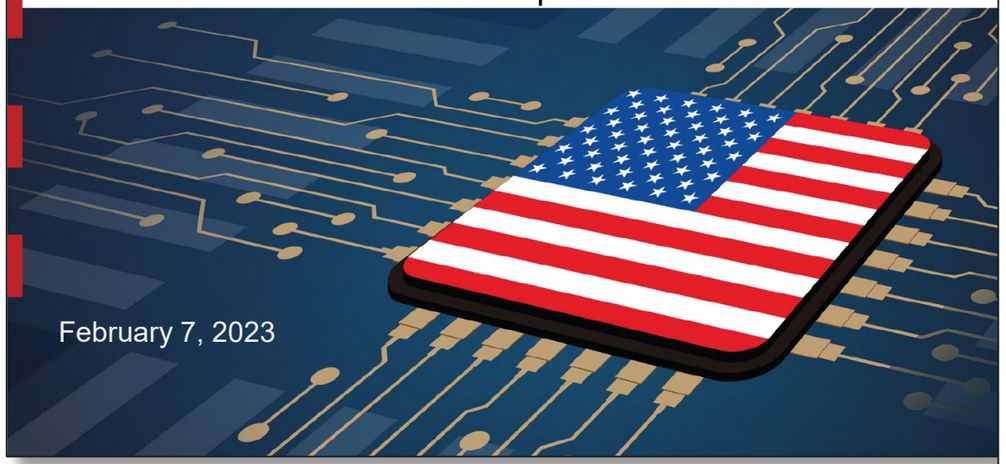

February 7, 2023

## CHIPS R&D Vision

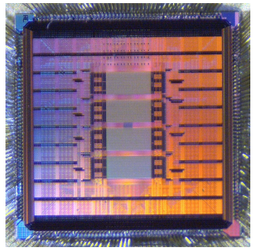

**U.S. Technology Leadership**

The U.S. invents, develops, and deploys the foundational semiconductor technology of the future.

**Accelerate Ideas to Market**

A thriving ecosystem that is focused on getting the best ideas to commercial scale as quickly and cost effectively as possible.

**Talent**

A new generation of skilled workers, inventors, designers, researchers, technicians, and others able to build and sustain semiconductor manufacturing in the U.S.



https://www.nist.gov/chips

See *Webinars* and *Industrial Advisory Committee*

Pacific Northwest
NATIONAL LABORATORY

# NIST R&D Elements:
## National Semiconductor Technology Center

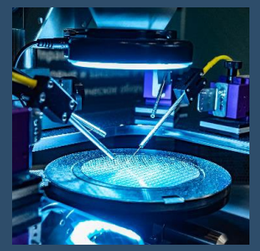

### National Semiconductor Technology Center

CHIPS for AMERICA

**Vision**: Will serve as the **focal point** for research and engineering throughout the semiconductor ecosystem, advancing and enabling disruptive innovation to provide U.S. leadership in the industries of the future.

**Structure**: A public-private consortium as an independent entity with a governing board informed and advised by industry, academia, government, and key stakeholders.

Pacific Northwest NATIONAL LABORATORY

# NIST R&D Elements:
## National Semiconductor Technology Center

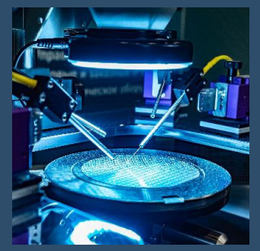

### National Semiconductor Technology Center

CHIPS for AMERICA

**Vision**: Will serve as the **focal point** for research and engineering throughout the semiconductor ecosystem, advancing and enabling d innovation to provide U.S. leadership of the future.

**Structure**: A public-private consortium independent entity with a governing b and advised by industry, academia, g key stakeholders.



### National Semiconductor Technology Center

CHIPS for AMERICA

**Elements:**

- Core of centrally operated, in-house research, engineering, and program capabilities combined with a network of directly funded and affiliated entities.
- Includes applied research, prototyping of devices and processes in a real-world environment, challenges related to scaling, start-up company support, or development of advanced manufacturing tools and processes.
- Focus research and engineering on challenging projects with a time horizon beyond 5 years.
- The NSTC will serve as a key convening body for the ecosystem.

**Process:**

- NSTC white paper, 1st quarter 2023.
- Summarize the results of a landscape analysis, outline a governance structure, and describe a preliminary operating and financial model.

Pacific Northwest
NATIONAL LABORATORY

# NIST R&D Elements:
## National Advanced Packaging Manufacturing Program

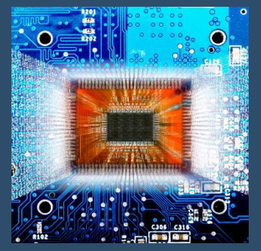

### National Advanced Packaging Manufacturing Program

CHIPS for AMERICA

- Strengthen semiconductor advanced test, assembly, and packaging capability in the domestic ecosystem
- Leverage public-private partnerships, that can include support for facilities managed by the NSTC and MUSA
- Broad range of technologies:
  - Heterogeneous integration
  - Wafer and panel-based approaches
  - Tooling and automation
  - Substrate technology

Pacific Northwest NATIONAL LABORATORY

# NIST R&D Elements:
## National Advanced Packaging Manufacturing Program

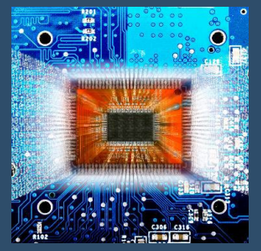

### National Advanced Packaging Manufacturing Program

CHIPS for AMERICA

- Strengthen semiconductor advanced test, assembly, and packaging capability in the domestic ecosystem
- Leverage public-private partnerships, th
  support for facilities managed by the NS
- Broad range of technologies:
  - Heterogeneous integration
  - Wafer and panel-based approache
  - Tooling and automation
  - Substrate technology



### NAPMP Target Areas

CHIPS for AMERICA

| | |
|---|---|
| Technology innovation | Create an R&D environment advancing the state-of-the art in advanced packaging. |
| Ecosystem support | Investments to bolster the growth in domestic capacity and enhance capabilities for competitive edge. |

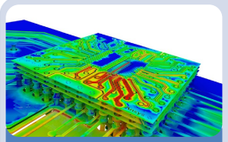

Co-design and simulation

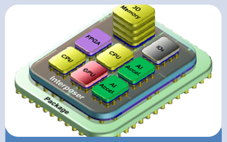

Chiplets

Pilot packaging facilities

Tooling and automation

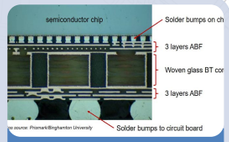


Materials and substrates

# NIST R&D Elements:
## Industrial Advisory Committee

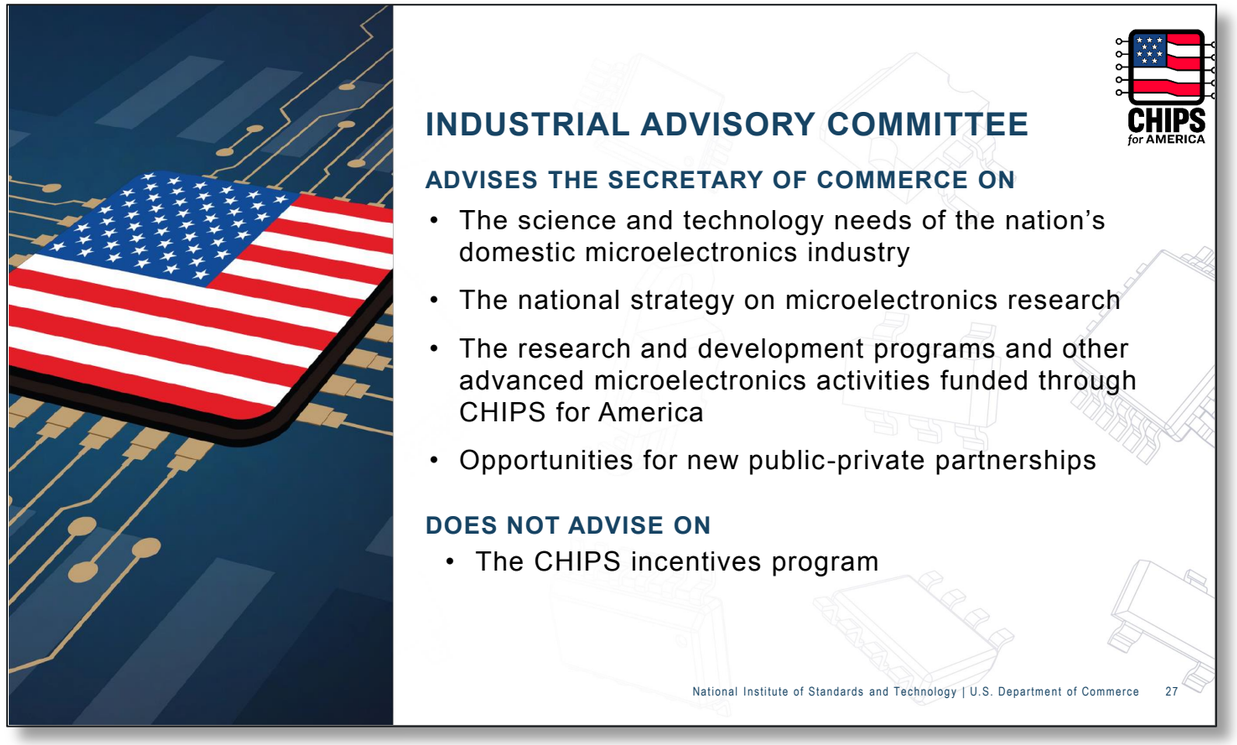

### INDUSTRIAL ADVISORY COMMITTEE

**ADVISES THE SECRETARY OF COMMERCE ON**

- The science and technology needs of the nation's domestic microelectronics industry
- The national strategy on microelectronics research
- The research and development programs and other advanced microelectronics activities funded through CHIPS for America
- Opportunities for new public-private partnerships

**DOES NOT ADVISE ON**

- The CHIPS incentives program

# NIST R&D Elements:
## Industrial Advisory Committee

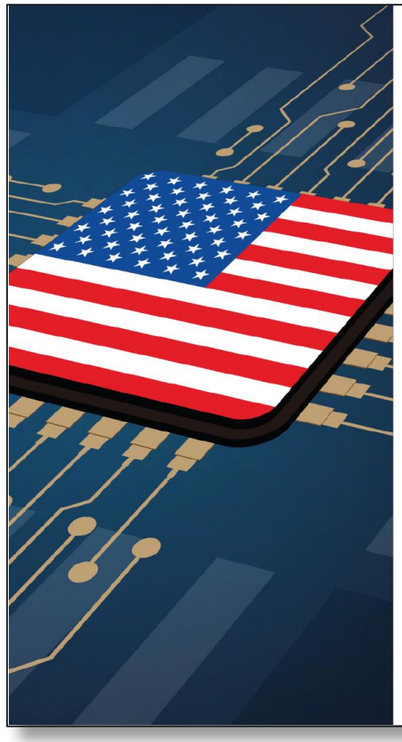

### INDUSTRIAL ADVISORY COMMITTEE

**ADVISES THE SECRETARY OF COMMERCE ON**

- The science and technology needs of the nation's domestic microelectronics industry
- The national strategy on microelectronics research
- The research and development programs and advanced microelectronics activities funded t CHIPS for America
- Opportunities for new public-private partners

**DOES NOT ADVISE ON**

- The CHIPS incentives program



### IAC Members

**Industry**


Mike Splinter
Chair
MRS Business and Technology Advisors


Susan Feindt
Vice-Chair
Analog Devices


Susie Armstrong
Qualcomm


Ahmad Bahai
Texas Instruments


Deirdre Hanford
Synopsys


Ken Joyce
Brewer Science


Ann Kelleher
Intel Corporation


Mukesh Khare
IBM Research


Meredith LaBeau
Calumet Electronics


Om Nalamasu
Applied Materials


Debo Olaosebikan
Kepler Computing


Alex Oscilowski
TEL America


Anthony Yen
ASML Technology Center

**Ecosystem/ Customers**


Daniel Armbrust
Silicon Catalyst


Bill Chappell
Microsoft


Charles Gray
Ford Motor Company

**Academia/Other Stakeholders**


James Ang
Pacific Northwest National Laboratory


Michael Fritze
Potomac Institute for Policy Studies


Carol Handwerker
Purdue University


Rajarao Jammy
MITRE Engenuity


Tsu-Jae King Liu
University of California Berkley


Willy Shih
Harvard Business School


Brandon Tucker
Washtenaw Community College


H.S. Philip Wong
Stanford University

# Data-Model Convergence (DMC) Initiative
## Co-design for Heterogeneous Computing

- DMC drives co-design with converged physical science and data science applications
- High Level Synthesis tools to develop hardware designs for reconfigurable heterogeneous architectures, or ASIC accelerator designs
- Software stack to support mapping of *converged* applications to DMC architecture concepts
  - Sparse linear algebra support, sparse data structures
  - Runtime systems
  - Compilers & Tools

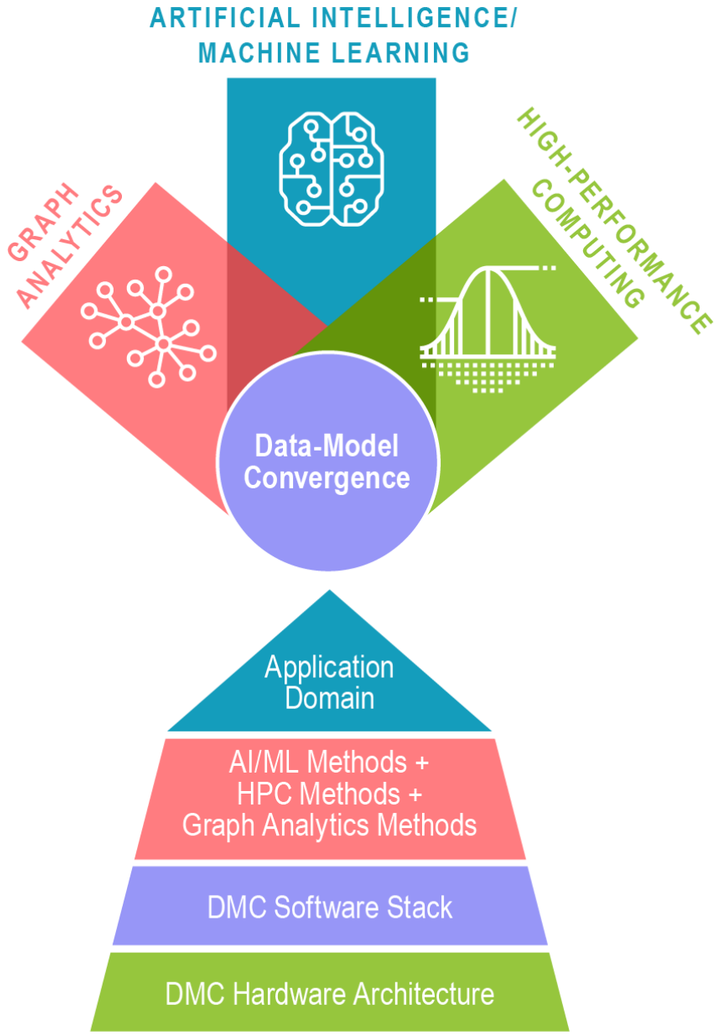


*Codesign for Extreme Heterogeneity: Integrating Custom Hardware with Commodity Computing Technology to Support Next-Generation HPC Converged Workloads*,
IEEE Internet Computing: Special Issue on Future of HPC, DOI: 10.1109/MIC.2022.3217423

# Open Source Co-Design Capabilities

- **Proxy Applications**
  - ExaGraph: miniVite, Ripples
  - TSPRD – Time Series Prediction
  - Sparse DNN
- **System Software Capabilities**
  - COMET – Compiler Framework
  - MCL – Heterogeneous Runtime
  - Fallacy – Memory Centric Tools
- **Hardware Design Tools**
  - SODA Synthesizer
  - PandA / Bambu
  - OpenCGRA

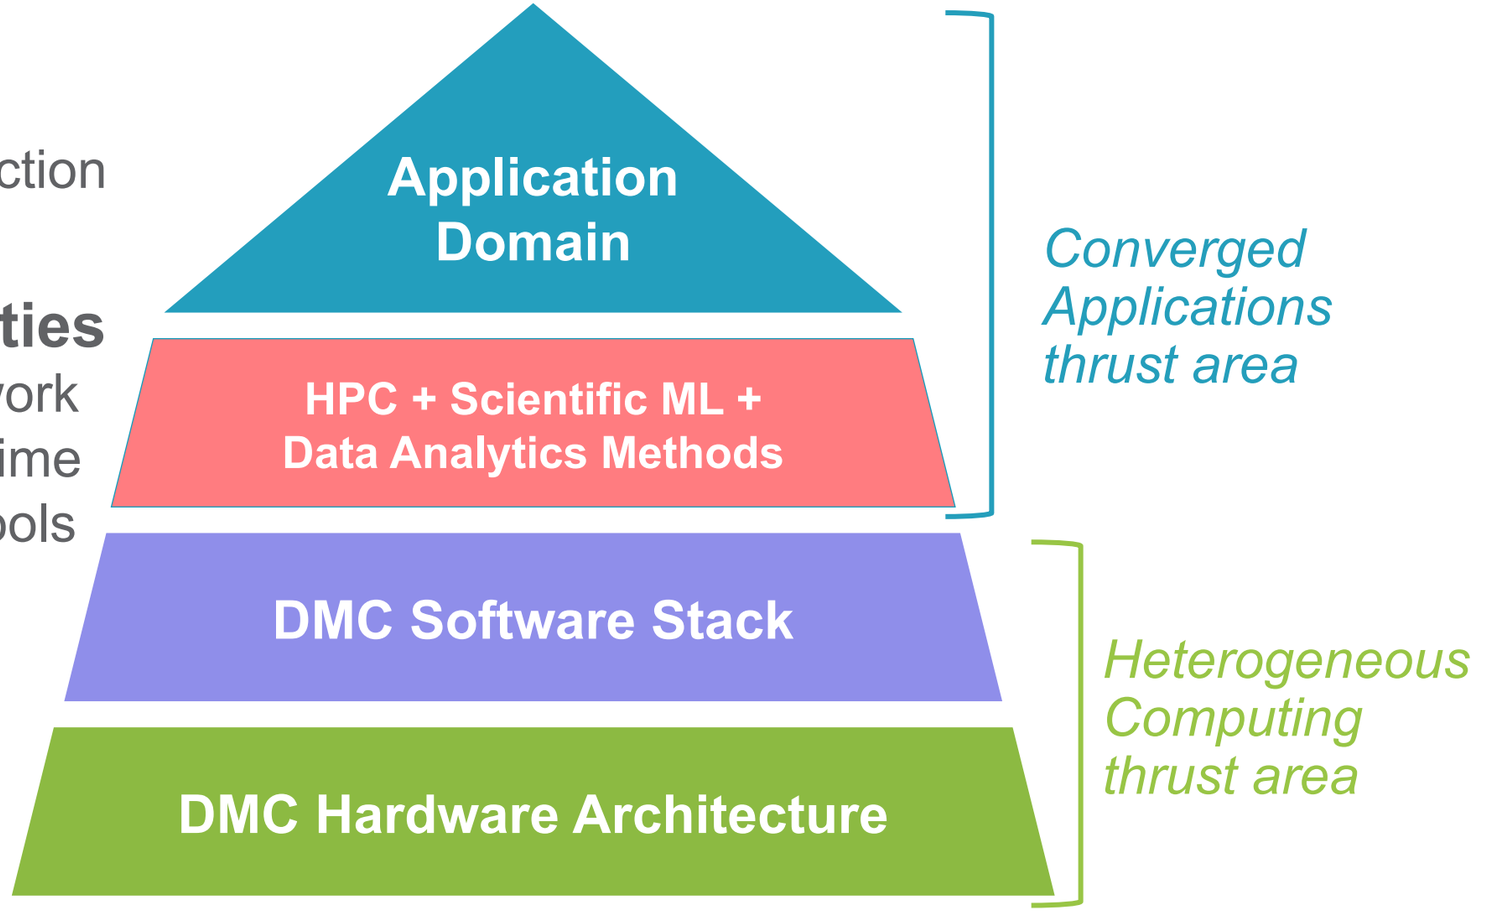


https://www.pnnl.gov/projects/dmc

# COMET: Domain Specific Compilation in Multi-level IR

PI: Gokcen Kestor

- COMET is a compiler and domain-specific language (DSL) for target heterogeneous systems that support dense and sparse tensor operators and storage formats
  - Focus on computational chemistry and graph analytics application domains
  - Supports various frontends, including DSL, numPy, and RUST
- COMET compiler infrastructure
  - **Abstraction** for dense/sparse storage formats to support wide range range of sparse storage formats
  - **Multi-level** code optimizations, including domain-specific and architecture specific
  - Support for **semiring** operations to represent graph algorithms
  - **Kernel Fusion** to avoid temporaries and redundant computation
  - Targets **heterogeneous architectures**, including CPUs, GPU, FPGAs, emerging dataflow architectures (SambaNova and Xilinx Versal)

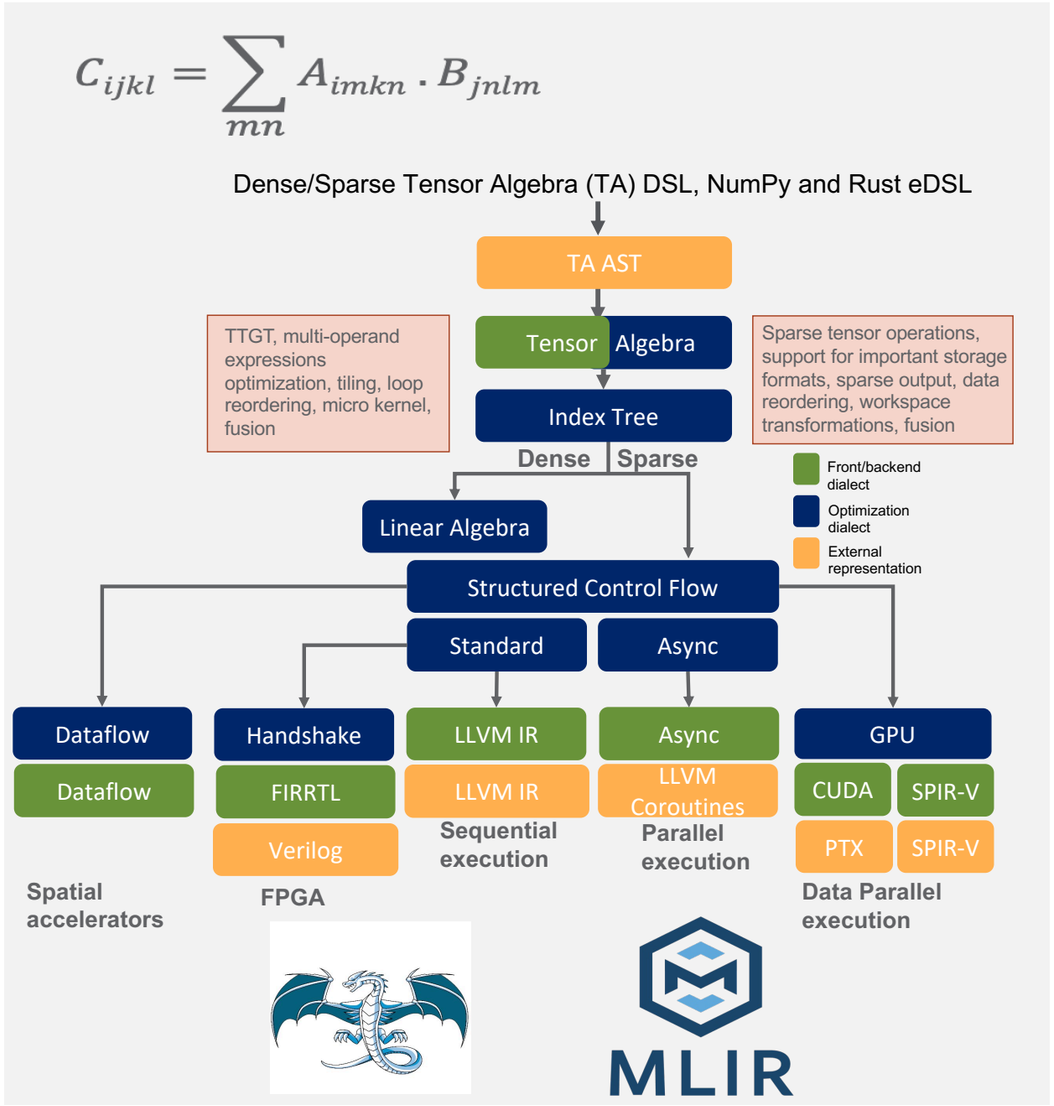

# SODA Synthesizer

PI: Antonino Tumeo

- A modular, multi-level, interoperable, extensible, open source ***hardware compiler*** from ***high-level programming frameworks to silicon***
- ***Compiler-based frontend***, leveraging the MultiLevel Intermediate Representation (MLIR)
- ***Compiler-based backend***, leveraging optimized architectural templates to match computing patterns as well as conventional High-Level Synthesis (HLS) techniques
- Generates ***synthesizable Verilog*** for a variety of targets, from Field Programmable Gate Arrays (FPGAs) to Application Specific Integrated Circuits (ASICs)
- Optimizations at all levels are performed as ***compiler optimization*** passes

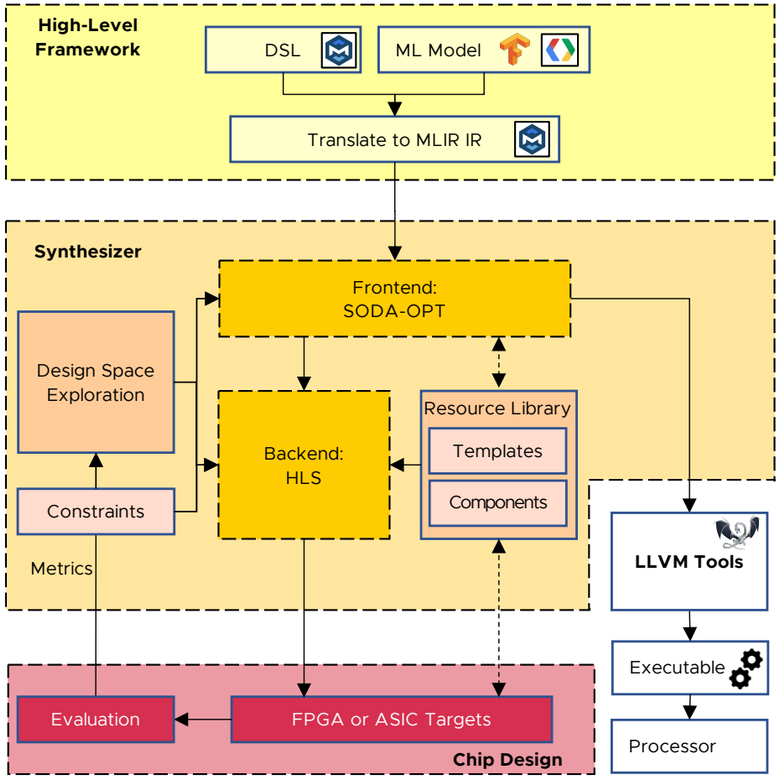


https://gitlab.pnnl.gov/sodalite/soda-frontend
https://gitlab.pnnl.gov/sodalite/soda-opt

# Open Hardware Technology Commons

- NIST RFI Submission: PNNL, SNL and LBNL
- University Collaborators
  - Columbia
  - UCSD
  - Northeastern
  - Harvard
  - Utah
- Industry Collaborators
  - Zero ASIC Corp.
  - Precision Innovations, Inc.
- Proposed On-ramp to NIST NSTC
  - Lower Barriers to Entry
  - Produce test hardware at leading edge
  - Support workforce development

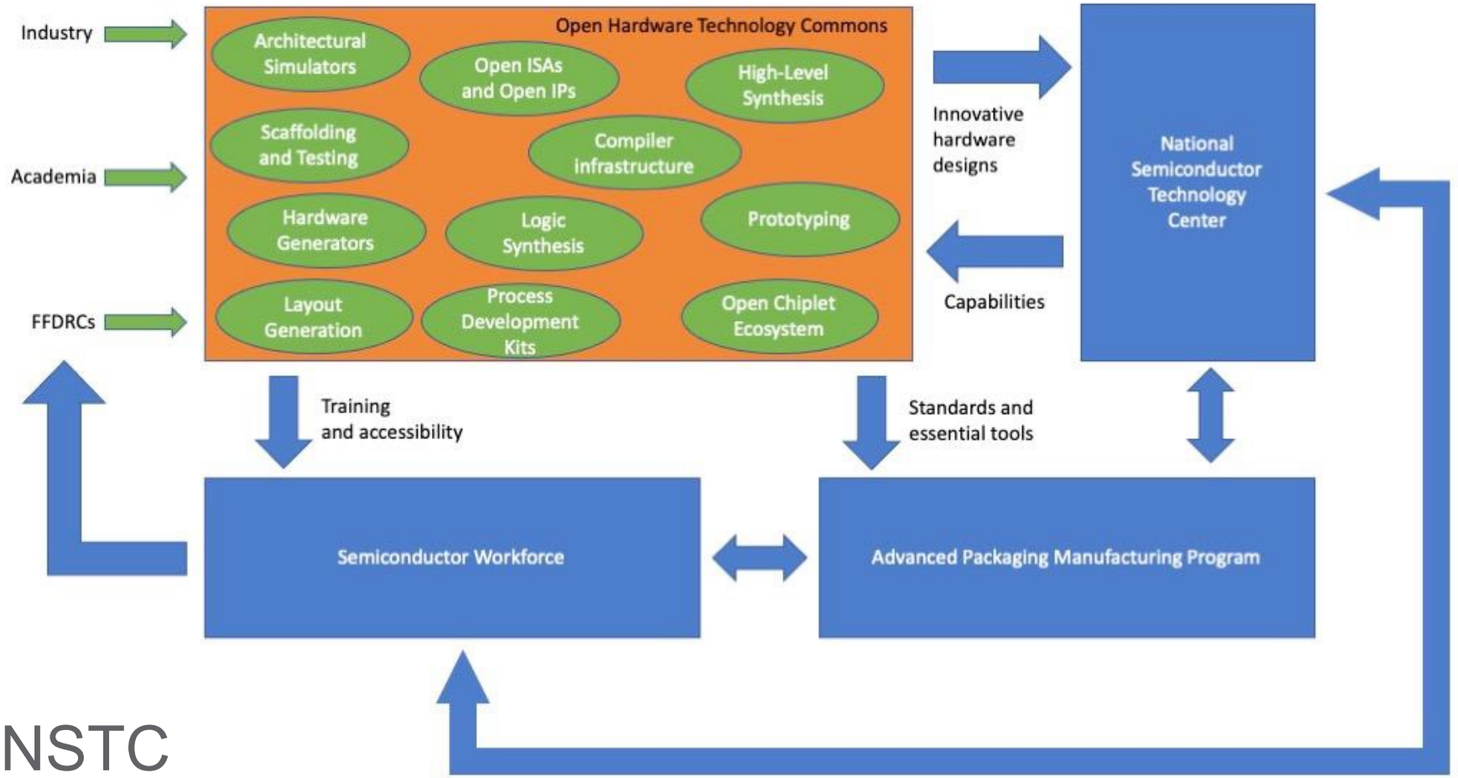

# Continuum Computing Use Case

## E3SM and ARM User Facility Data → Distributed HPC + Data

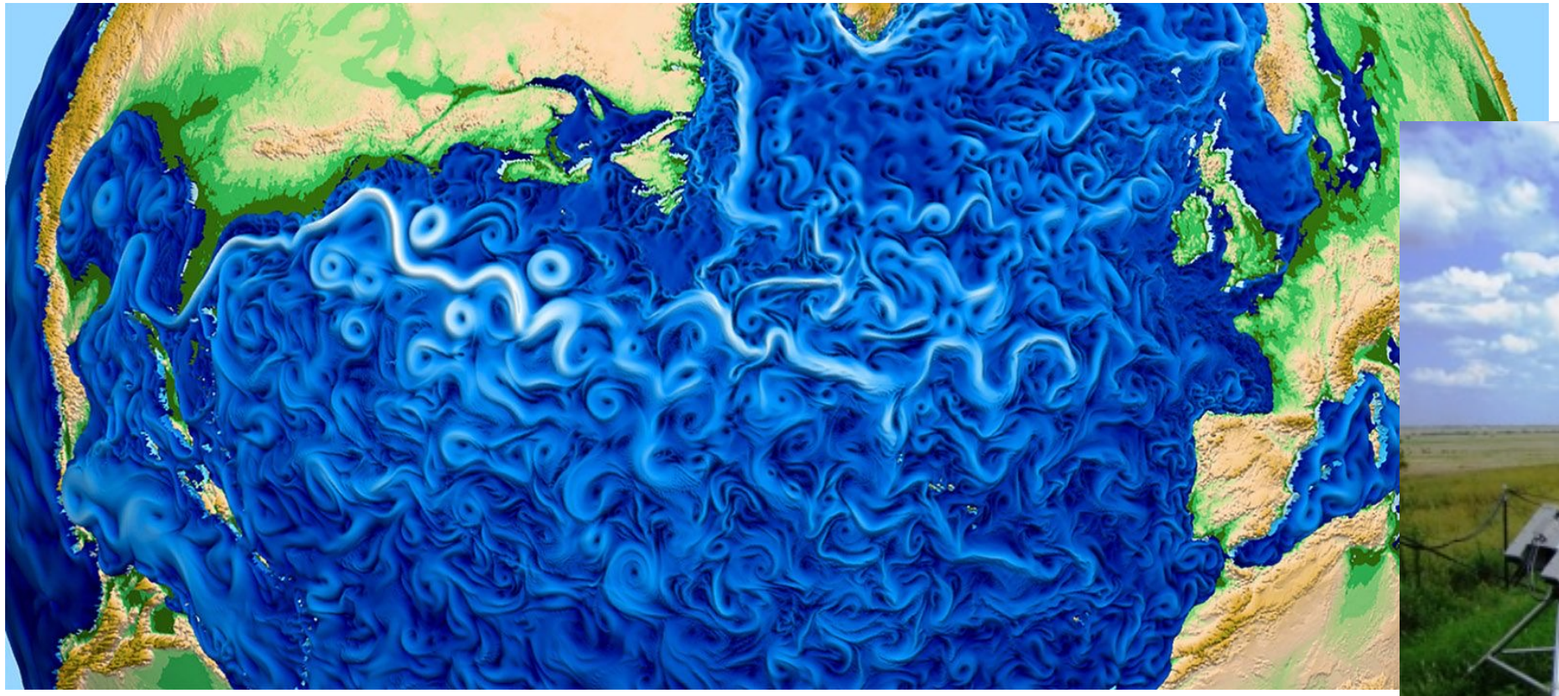

Radiometers at ARM's Southern Great Plains atmospheric observatory. Image courtesy of the U.S. Department of Energy's Atmospheric Radiation Measurement (ARM) user facility.

https://arm.gov

E3SM Energy Exascale Earth System Model

https://e3sm.org

AI4ESP AI for Earth System Predictability

https://ai4esp.org

# Metrics for Continuum Computing

- What is the Application?
  - Distributed scientific discovery workflows
  - Control of complex distributed engineering systems, e.g., Smart Grid
  - Digital Twins!
- Computing performance tracks application operations across the continuum
- Performance is time to solution for Applications
- Energy Efficiency is energy consumed to complete Applications
- What is more important in Continuum Computing
  - Time to Solution? or Energy to Solution?
- What is the programming model?
- What are the strategies for fault tolerance?

Pacific Northwest
NATIONAL LABORATORY

DMC
DATA-MODEL
CONVERGENCE
INITIATIVE
@PNNL

# Questions

*ang@pnnl.gov*