OAK RIDGE National Laboratory

Deep Codesign in the Post-Exascale Computing Era

Jeffrey S. Vetter

With many contributions from ACSR Section, Colleagues, and Collaborators

Multicore World
Wellington
14 Feb 2023

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

https://www.ornl.gov/section/advanced-computing-systems-research (https://j.mp/acsrs)
vetter@computer.org

ORNL is managed by UT-Battelle
for the US Department of Energy

Highlights

15 years to go from Exascale ideation to deployment

Many reports and predictions

Exascale status

What did we get right, get wrong, overlook?

Post Exascale

Heterogeneous systems enabled by heterogeneous integration and Chiplets

Codesign becomes CRITICAL

Abikso

New microelectronics codesign project attempting deep codesign from algorithms to materials

Exascale Reports (and predictions) from 2007 to 2014

Modeling and Simulation at the Exascale for Energy and the Environment

Report on the Advanced
Town Ha
Simulation and Modeling at the Exa
and Globa

Co-Chairs: Lawrence Berkeley Natio
Oak Ridge National Labor
Argonne National Laborat

Office of Advanced Scientific Computing Research Contact: Michael Strayer

Special Assistance
Technical: Lawrence Berkeley Natio
Deb Agarwal, David Bailey,
William Collins, Nikos Kyrpi
Peter Nugent, Leonid Olike
Lin-Wang Wang, Michael W

Oak Ridge National Labor
Eduardo D'Azevedo, David
James Hack, Victor Hazlew
Bronson Messer, Anthony M
B. (Rad) Radhakrishnan, N
Jeffrey Vetter, Gilbert Weiga

Argonne National Laborat
Raymond Bair, Pete Beckm
Ed Frank, Ian Foster, Willia
Robert Jacob, Kenneth Ker
Jorge Moré, Lois McInnes,
Michael Papka, Robert Ros

Administrative: Lawrence Berkeley Natio
Oak Ridge National Labor
Argonne National Laborat

Publication: Oak Ridge National Labor
Argonne National Laborat

Editorial: Oak Ridge National Labor
Argonne National Laborat

This report is available on the web at http://www.sc.d

ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems

Peter Kogge, Editor & Study Lead
Keren Bergman
Shekhar Borkar
Dan Campbell
William Carlson
William Dally
Monty Denneau
Paul Franzon
William Harrod
Kerry Hill
Jon Hiller
Sherman Karp
Stephen Keckler
Dean Klein
Robert Lucas
Mark Richards
Al Scarpelli
Steven Scott
Allan Snavely
Thomas Sterling
R. Stanley Williams
Katherine Yelick

September 28, 2008

This work was sponsored by DARPA IPTO i
as Program Manager; AFRL contract number
interest of scientific and technical informatio
Government's approval or disapproval of its

SCIENTIFIC GRAND CHALLENGES: CROSSCUTTING TECHNOLOGIES FOR COMPUTING AT THE EXASCALE

Report from the Workshop Held February 2-4, 2010

**Sponsored by the U.S. Department of Energy, O
Research, Office of Science; and the Office of A
National Nuclear Security Administration**

Chair, **David L. Brown**
Lawrence Livermore National Laboratory

Chair, **Paul Messina**
Argonne National Laboratory

Theme I: Domain Science and System Architecture

Principal Lead, **David Keyes**
King Abdullah University of Science and Technology a

Co-Lead, **John Morrison**
Los Alamos National Laboratory

Co-Lead, **Robert Lucas**
University of Southern California

Co-Lead, **John Shalf**
Lawrence Berkeley National Laboratory

Theme II: System Software

Principal Lead, **Pete Beckman**
Argonne National Laboratory

Co-Lead, **Ron Brightwell**
Sandia National Laboratories

Co-Lead, **Al Geist**
Oak Ridge National Laboratory

Theme III: Programming Models and Environment

Principal Lead, **Jeffrey Vetter**
Oak Ridge National Laboratory and Georgia Institute of Technology

ASCAC Subcommittee for the Top Ten Exascale Research Challenges

Subcommittee Chair
Robert Lucas (University of Southern California

Subcommittee Members
James Ang (Sandia National Laboratories)
Keren Bergman (Columbia University)
Shekhar Borkar (Intel)
William Carlson (Institute for Defense Analyses)
Laura Carrington (UC, San Diego)
George Chiu (IBM)
Robert Colwell (DARPA)
William Dally (NVIDIA)
Jack Dongarra (U. Tennessee)
Al Geist (ORNL)
Gary Grider (LANL)
Rud Haring (IBM)
Jeffrey Hittinger (LLNL)
Adolfy Hoisie (PNLL)
Dean Klein (Micron)
Peter Kogge (U. Notre Dame)
Richard Lethin (Reservoir Labs)
Vivek Sarkar (Rice U.)
Robert Schreiber (Hewlett Packard)
John Shalf (LBNL)
Thomas Sterling (Indiana U.)
Rick Stevens (ANL)

The International Exascale Software Project roadmap

The International Journal of High Performance Computing Applications
25(1) 3–60
© The Author(s) 2011
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/1094342010391989
hpc.sagepub.com
SAGE

Jack Dongarra, Pete Beckman, Terry Moore, Patrick Aerts, Giovanni Aloisio, Jean-Claude Andre, David Barkai, Jean-Yves Berthou, Taisuke Boku, Bertrand Braunschweig, Franck Cappello, Barbara Chapman, Xuebin Chi, Alok Choudhary, Sudip Dosanjh, Thom Dunning, Sandro Fiore, Al Geist, Bill Gropp, Robert Harrison, Mark Hereld, Michael Heroux, Adolfy Hoisie, Koh Hotta, Zhong Jin, Yutaka Ishikawa, Fred Johnson, Sanjay Kale, Richard Kenway, David Keyes, Bill Kramer, Jesus Labarta, Alain Lichnewsky, Thomas Lippert, Bob Lucas, Barney Maccabe, Satoshi Matsuoka, Paul Messina, Peter Michielse, Bernd Mohr, Matthias S. Mueller, Wolfgang E. Nagel, Hiroshi Nakashima, Michael E Papka, Dan Reed, Mitsuhisa Sato, Ed Seidel, John Shalf, David Skinner, Marc Snir, Thomas Sterling, Rick Stevens, Fred Streitz, Bob Sugar, Shinji Sumimoto, William Tang, John Taylor, Rajeev Thakur, Anne Trefethen, Mateo Valero, Aad van der Steen, Jeffrey Vetter, Peg Williams, Robert Wisniewski and Kathy Yelick

Abstract
Over the last 20 years, the open-source community has provided more and more software on which the world's high-performance computing systems depend for performance and productivity. The community has invested millions of dollars and years of effort to build key components. However, although the investments in these separate software elements have been tremendously valuable, a great deal of productivity has also been lost because of the lack of planning, coordination, and key integration of technologies necessary to make them work together smoothly and efficiently, both within individual petascale systems and between different systems. It seems clear that this completely uncoordinated development model will not provide the software needed to support the unprecedented parallelism required for peta/exascale computation on millions of cores, or the flexibility required to exploit new hardware models and features, such as transactional memory, speculative execution, and graphics processing units. This report describes the work of the community to prepare for the challenges of exascale computing, ultimately combing their efforts in a coordinated International Exascale Software Project.

Keywords
exascale computing, high-performance computing, software stack

Table of Contents

DOE HPC Roadmap to Exascale Systems

FY 2012	FY 2016	FY 2018	FY 2021	FY 2022	FY 2023

decommissioned

ORNL
Cray/AMD/NVIDIA

ANL
IBM BG/Q

ANL
Cray/Intel KNL

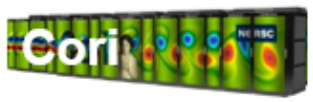

LBNL
Cray/Intel Xeon/KNL

LLNL
IBM BG/Q

LANL/SNL
Cray/Intel Xeon/KNL

ORNL
IBM/NVIDIA

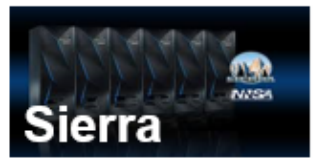

LLNL
IBM/NVIDIA

ORNL
HPE/AMD

Exascale Systems

ANL
HPE/AMD/NVIDIA

LBNL
HPE/AMD/NVIDIA

ANL
Intel/HPE

LANL/SNL
HPE/Intel

LLNL
HPE/AMD

More heterogeneity in processing, memory, storage

Frontier System

System
- 74 compute racks
- 29 MW Power Consumption
- 9,408 nodes
- 9.2 PB memory
  (4.6 PB HBM, 4.6 PB DDR4)
- Cray Slingshot network with dragonfly topology
- 37 PB Node Local Storage
- 716 PB Center-wide storage
- 4000 ft^2 foot print

Frontier Cabinet

Olympus rack
- 128 AMD nodes
- 8,000 lbs
- Supports 400 KW

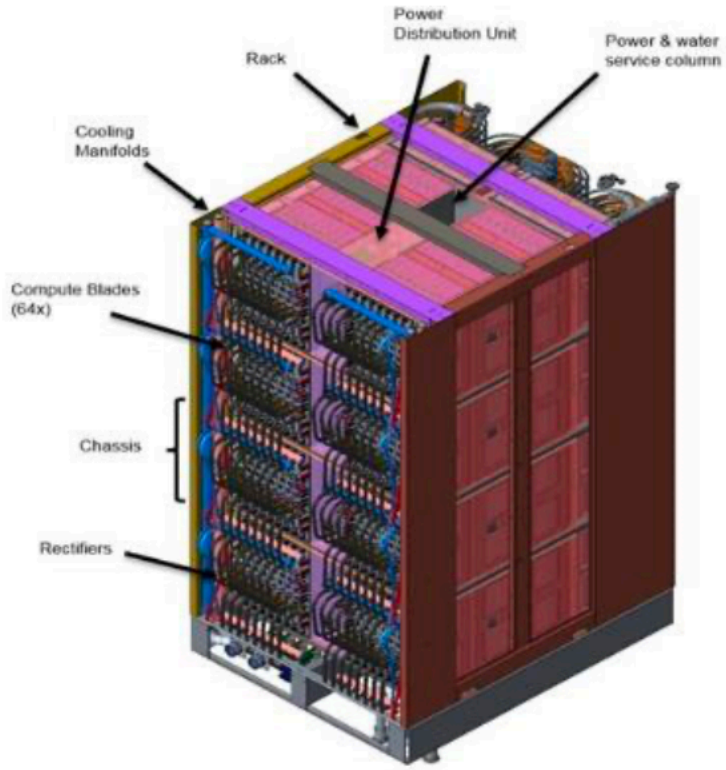

Frontier Node

AMD extraordinary engineering
- 1 AMD "Trento" CPU (optimized Milan)
- 4 AMD MI250X GPUs
- 512 GiB DDR4 memory on CPU
- 512 GiB HBM2e total per node
- 4 Cassini NICs connected to the 4 GPUs

Compute blade
- 2 AMD nodes

All water cooled, even DIMMS and NICs

OAK RIDGE NATIONAL LABORATORY'S FRONTIER SUPERCOMPUTER

- **74 HPE Cray EX cabinets**
- **9,408 AMD EPYC CPUs, 37,632 AMD GPUs**
- **700 petabytes of storage capacity, peak write speeds of 5 terabytes per second using Cray Clusterstor Storage System**
- **90 miles of HPE Slingshot networking cables**

TOP500

#1

1.1 exaflops of performance on the May 2022 Top500.

GREEN500

#1, #2

62.04 gigaflops/watt power efficiency on a single cabinet.
52.23 gigaflops/watt power efficiency on the full system.

HPL-AI

#1

6.88 exaflops on the HPL-AI benchmark.

Sources: May 30, 2022 Top500 release

Exascale Computing Project has three technical areas to meet national goals

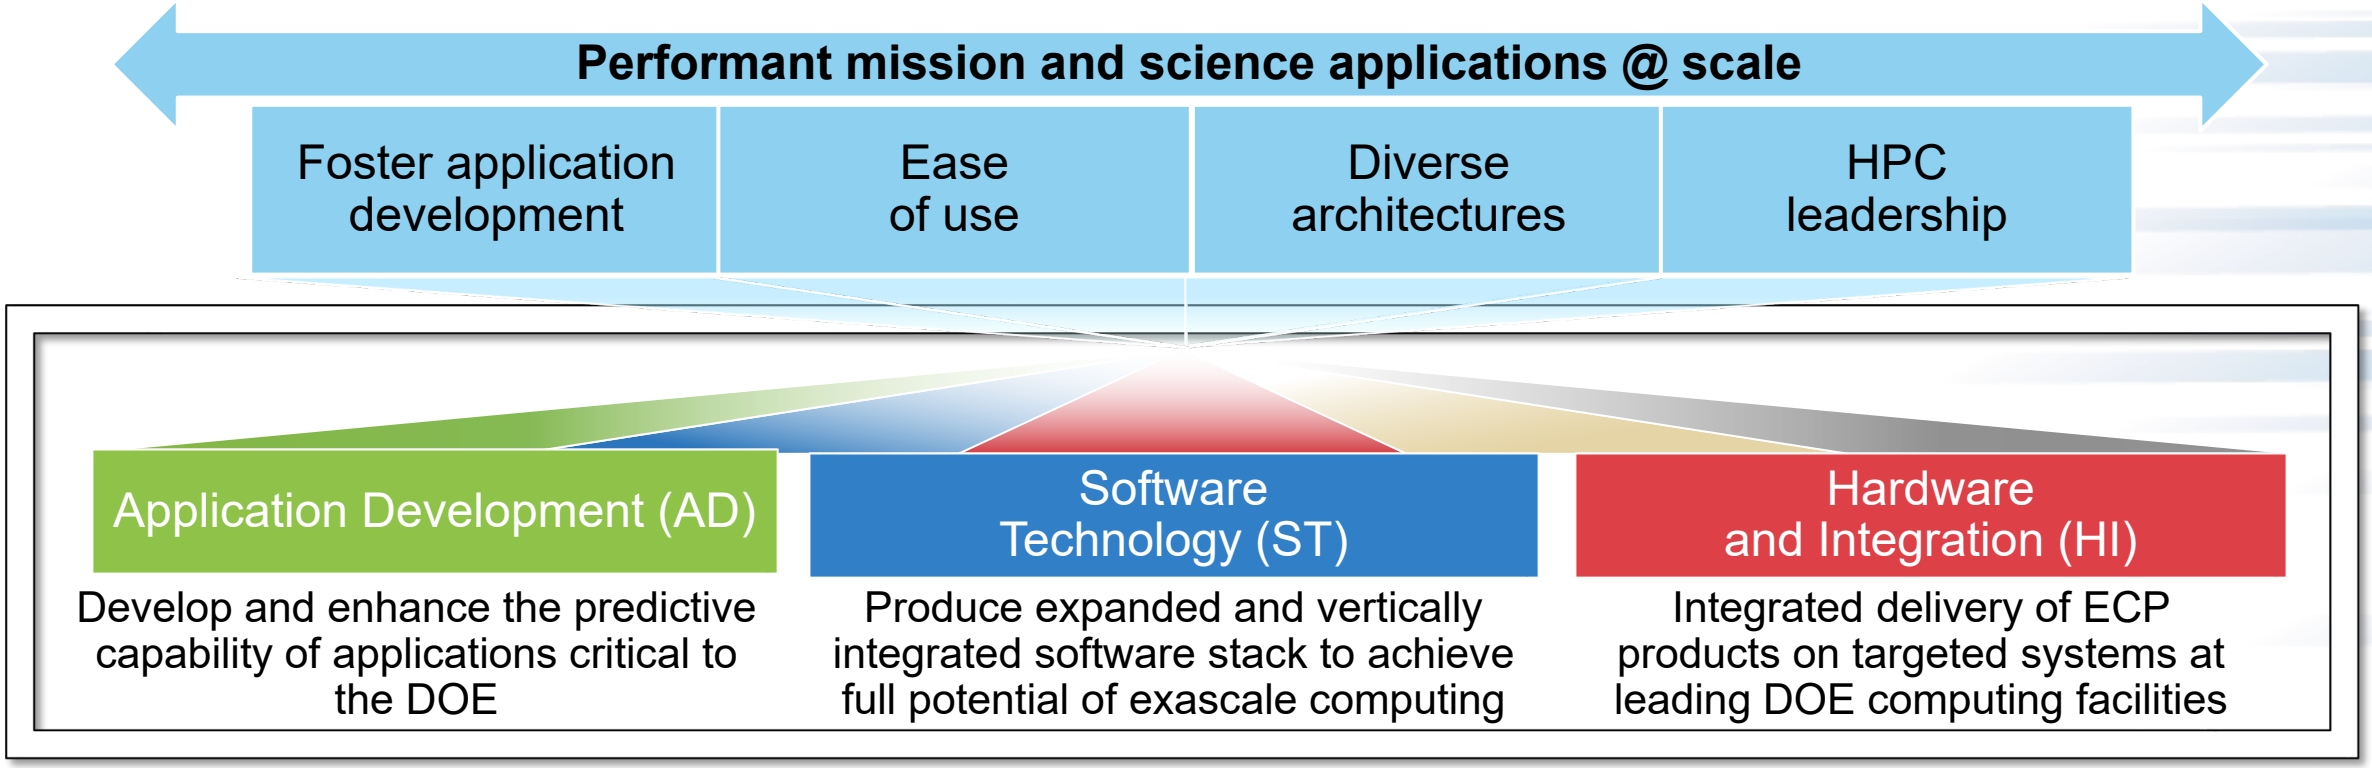

25 applications ranging from national security, to energy, earth systems, economic security, materials, and data

80+ unique software products spanning programming models and run times, math libraries, data and visualization

6 vendors supported by PathForward focused on memory, node, connectivity advancements; deployment to facilities

https://www.exascaleproject.org/

EXASCALE COMPUTING PROJECT

Extreme-scale Scientific Software Stack (E4S)

- E4S: HPC software ecosystem – a curated software portfolio
- A **Spack-based** distribution of software tested for interoperability and portability to multiple architectures
- Available from **source**, **containers**, **cloud, binary caches**
- Leverages and enhances SDK interoperability thrust
- Not a commercial product – an open resource for all
- Growing functionality: May 2022: E4S 22.05 – 100+ full release products

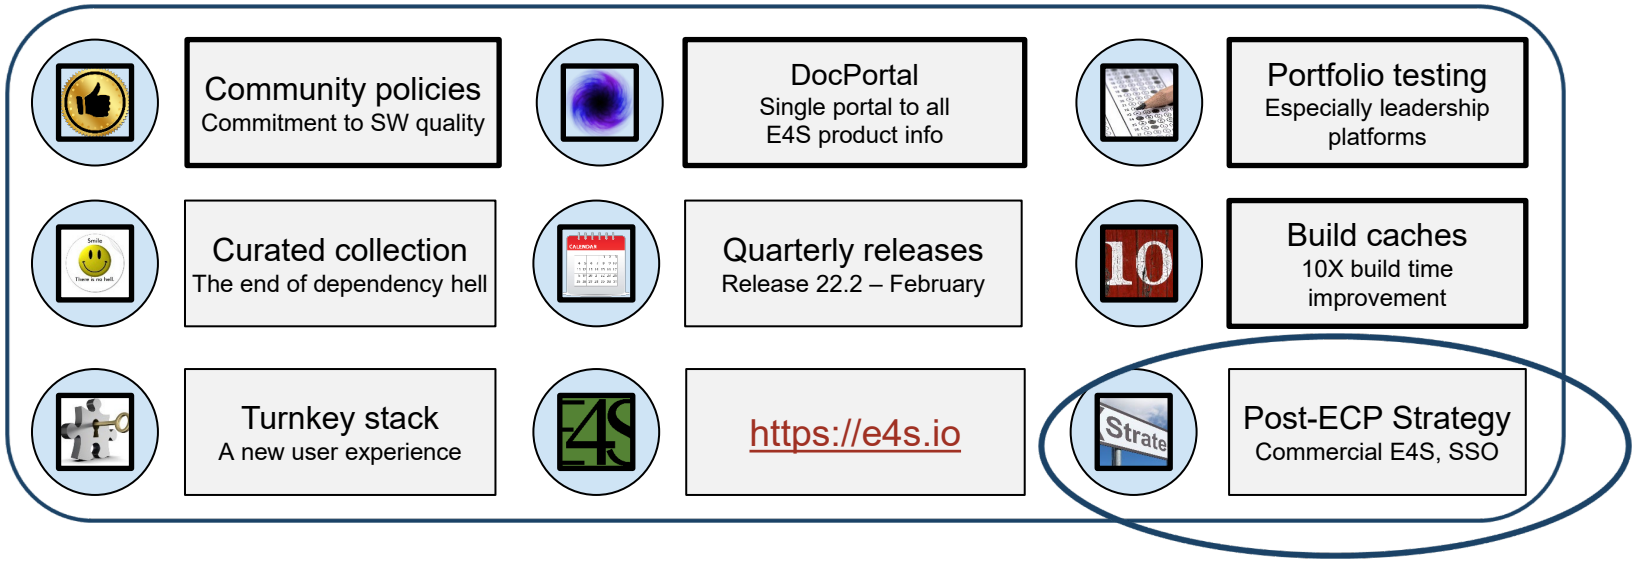

Spack

https://spack.io

Spack lead: Todd Gamblin (LLNL)

E4S

https://e4s.io

E4S lead: Sameer Shende (U Oregon)

Also includes other products, e.g.,
AI: PyTorch, TensorFlow, Horovod
Co-Design: AMReX, Cabana, MFEM

ECP EXASCALE COMPUTING PROJECT

ECP is Improving the LLVM Compiler Ecosystem

LLVM

- Very popular open-source **compiler infrastructure**
- **Permissive** license
- **Modular**, well-defined IR allows use by a lot of different languages, ML frameworks, etc.
- **Backend infrastructure** allowing the efficient creation of backends for new (heterogeneous) hardware.
- A **state-of-the-art C++ frontend**, CUDA support, scalable LTO, sanitizers and other debugging capabilities, and more.

+ SOLLVE

- Enhancing the implementation of OpenMP in LLVM
- Unified memory
- Prototype OMP features for LLVM
- OMP Optimizations
- OMP test suite
- Tracking OMP implementation quality
- Training

+ PROTEAS-TUNE

- Core optimization improvements to LLVM
  - OpenMP offload
- OpenACC capability for LLVM
  - Clacc
  - Flacc
- Autotuning for OpenACC and OpenMP in LLVM
- Integration with Tau performance tools
- SYCL characterizing and benchmarking
- Kokkos offloading
- Leading LLVM-DOE fork

+ FLANG

- Developing an open-source, production Fortran frontend
- Upstream to LLVM public release
- Support for OpenMP and OpenACC
- Recently approved by LLVM
- Initial implementation of serial F77 compiler for CPUs under review

+ HPCToolkit

- Improvements to OpenMP profiling interface OMPT
- OMPT specification improvements
- Refine HPCT for OMPT improvements

+ NNSA

- Enhancing LLVM to optimize template expansion for FlexCSI, Kokkos, RAJA, etc.
- Flang testing and evaluation
- Kitsune and Tapir

Vendors

- Increasing dependence on LLVM
- Many vendors import and redistribute LLVM
- Contributions and collaborations with many vendors through LLVM
  - AMD
  - ARM
  - Cray
  - HPE
  - IBM
  - Intel
  - NVIDIA

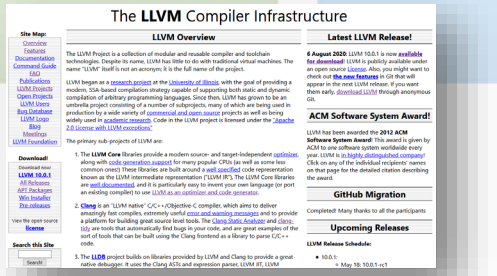

Other ECP activities with LLVM emerging organically.

Active involvement with broad LLVM community: LLVM Dev, EuroLLVM
ECP personnel had 10+ presentations at the 2020 Dev Meeting

OAK RIDGE National Laboratory

KokkACC: Enhancing Kokkos with OpenACC

The Science

This study set out to improve the programming productivity and capacity of Kokkos programing model, the reference programing model for performance portability, by using OpenACC. Due to the descriptive nature of OpenACC, the new OpenACC back end of Kokkos is not only simpler to implement, but also cheaper to maintain than other back ends such as CUDA or HIP, which are typically not designed to deal with Kokkos high-level programing constructs.

The Impact

This study was awarded best paper by the Ninth Workshop on Accelerator Programming Using Directives, at the International Conference for High Performance Computing, Networking, Storage, and Analysis 2022.

The OpenACC's integration with Kokkos has provided new capacities that enhance Kokkos with new capacities in terms of performance, portability, heterogeneity and autotuning. The result, KokkACC, has proven faster than OpenMP-Target and CUDA back ends by as much as an order of magnitude on some of the most important ECP applications, such as LAMMPS or Lulesh.

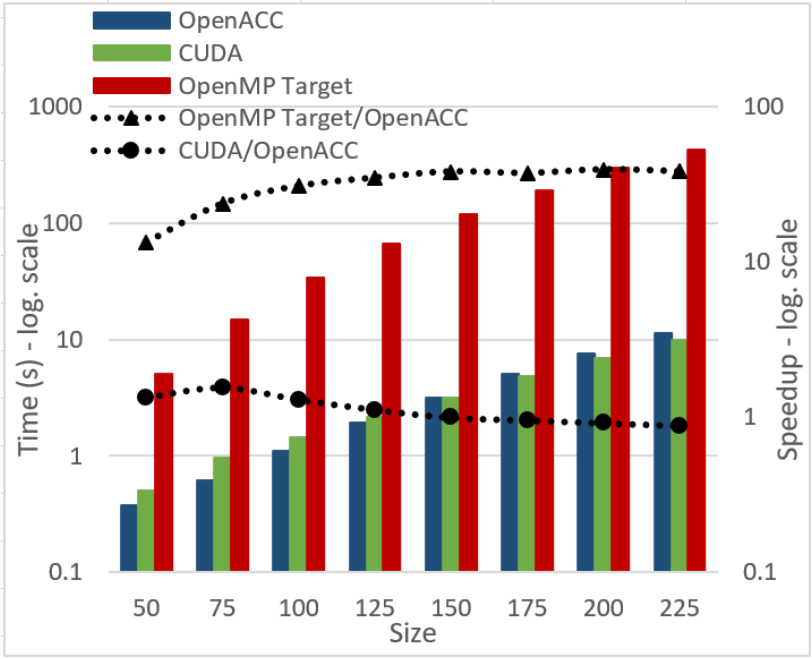

This graphic illustrates the run time of different Kokkos back ends; OpenACC (KokkACC), OpenMP-Target and CUDA on high-performance computing runs of the ECP LULESH code on one ORNL's SUMMIT supercomputer.

PI/Facility Lead(s): Jeffrey S. Vetter, Pedro Valero-Lara
ASCR Program/Facility: ECRP
Funding: ASCR, ECP
Publication(s) for this work: Pedro Valero-Lara, et al., "KokkACC: Enhancing Kokkos with OpenACC," Ninth Workshop on Accelerator Programming Using Directives, *International Conference for High Performance Computing, Networking, Storage, and Analysis 2022*.
DOI: 10.1109/WACCPD56842.2022.00009

Predictions

- "It's tough to make predictions, especially about the future" – Yogi Berra
- "Prediction is very difficult, especially about the future" -- Niels Bohr

How did our predictions play out?

ECP Panel at ECP AHM in May 2022

Jeffrey Vetter (ORNL), Moderator
Pete Beckman (ANL)
Jack Dongarra (UTK, ORNL)
Bob Lucas (Ansys)
Kathy Yelick (UCB)

Hits

- System power came in at O(20MW) not O(1GW)
- Few major software rewrites / evolution
  - So far, FORTRAN -> C++ is the main conversion
- ECP included applications, software, and hardware
  - ~70 teams, ~1000 researchers
  - IESP
- Concurrency (1B-way parallelism)
- Open-source software

Misses

- Systems deployed 4+ years later than expected (of 2018)
- Programming systems are multiplying and immature/incomplete
- Hardware diversity
- Resiliency and fault tolerance
- Dwindling number of vendors capable of $100M+ procurements

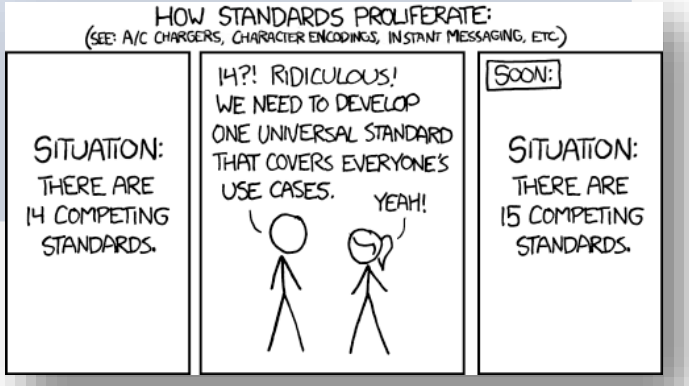

Overlooked

- Productive programming models (ala AI/ML): Python, Jupyter, Julia
- Cost of ECP + NRE + Procurements approaches ~$4B USD
- AI/ML is not predicted (or even mentioned)
- Cloud deployment models
- Green/sustainable computing

OAK RIDGE National Laboratory

Pondering Post-Exascale Computing

- Thinking about the next 10-15 years

TOP500 Macro View

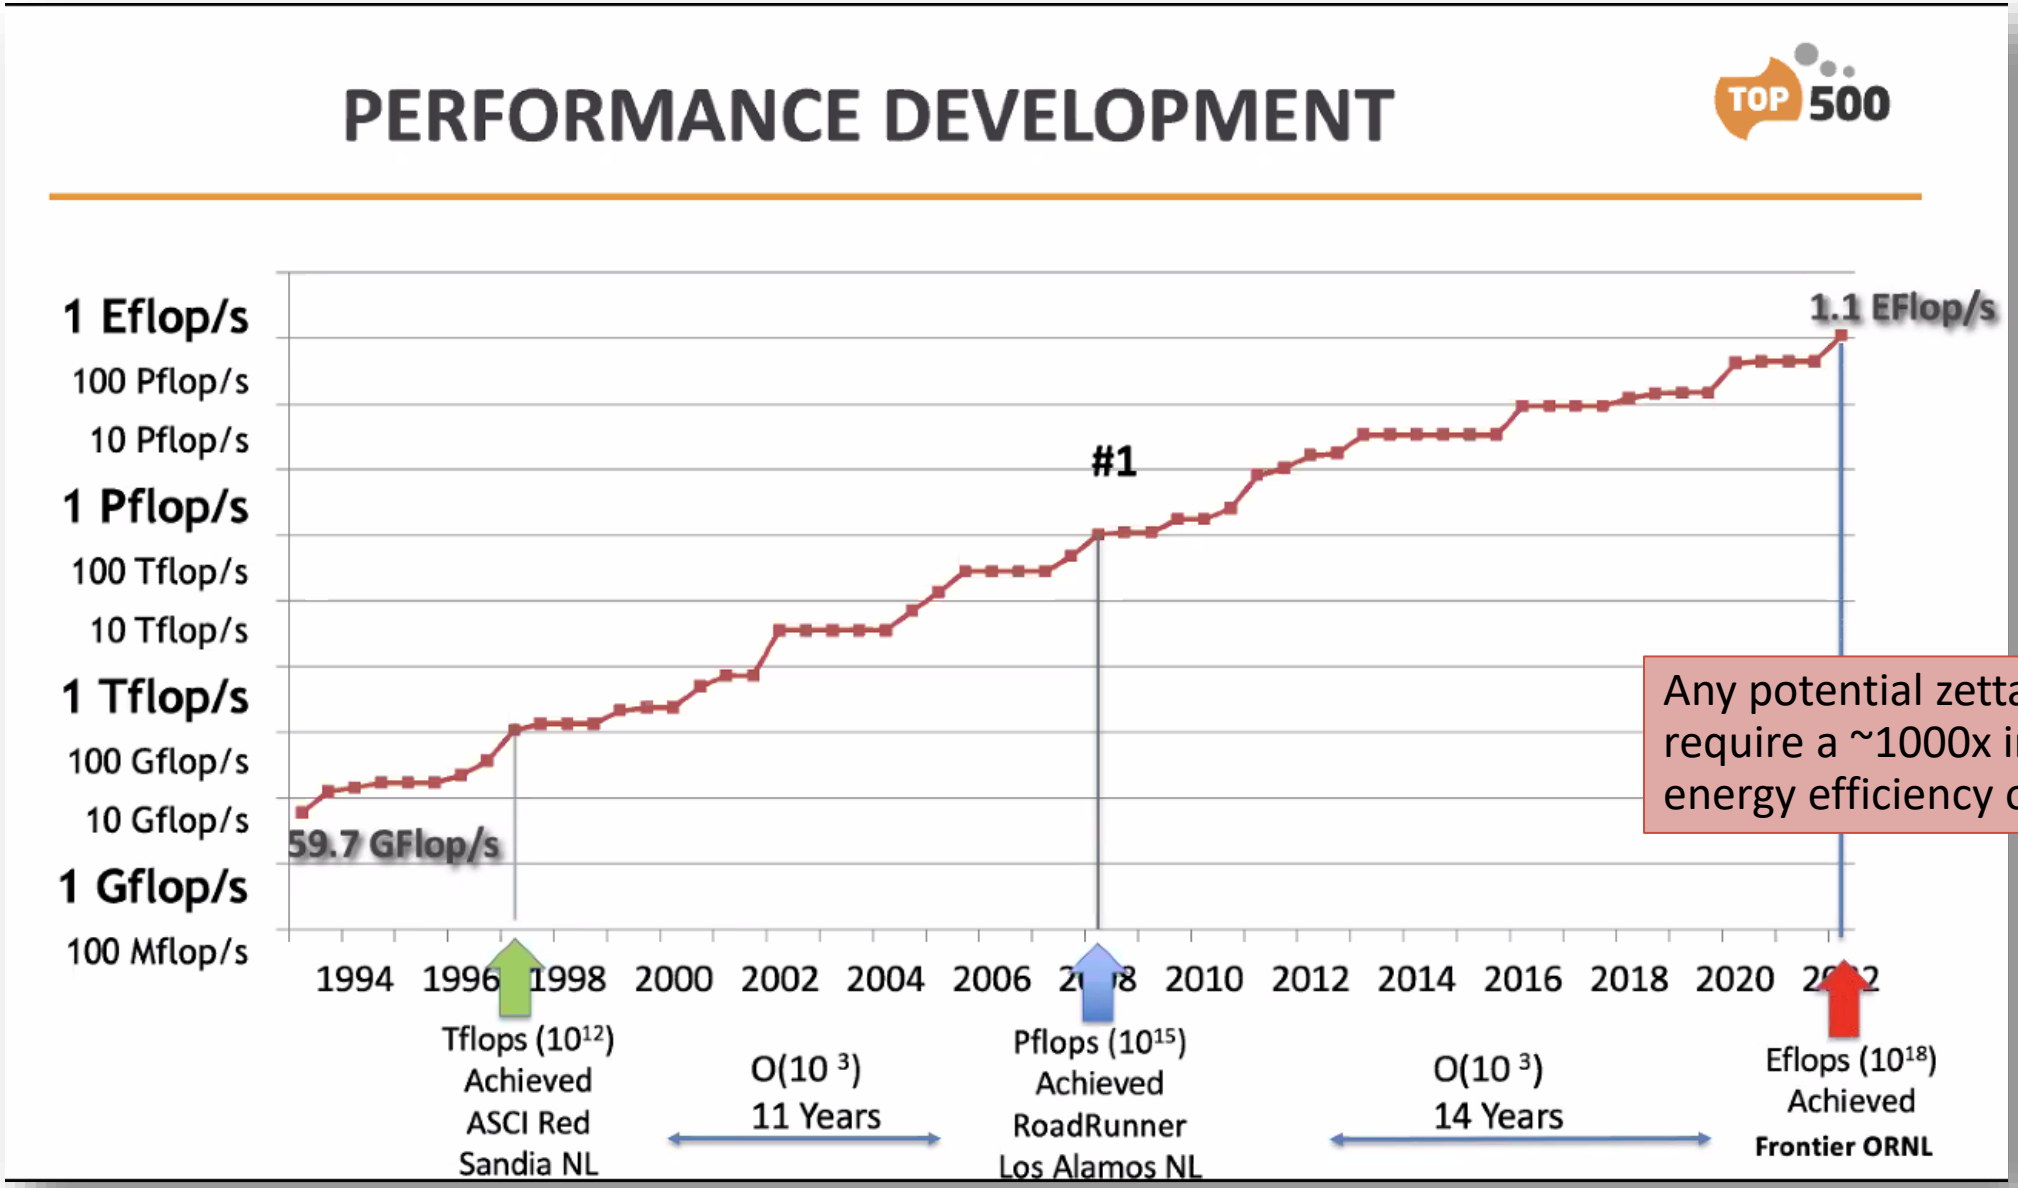

Any potential zettascale system will require a ~1000x improvement in energy efficiency over Frontier.

Dongarra, 2023

Important Architectural Trends

- Chiplets
  - Improved yields
- Heterogeneous integration
  - Integration of *many* different technologies into the same package
- **CHIPS and Science Act**

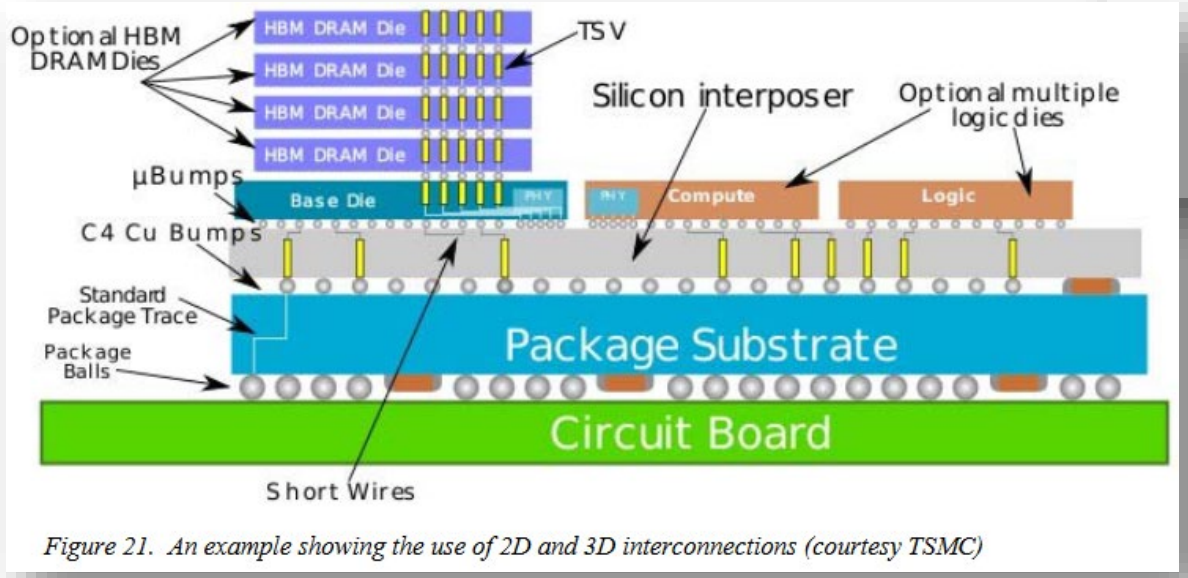

Figure 21. An example showing the use of 2D and 3D interconnections (courtesy TSMC)

[IEEE HIR 2021]

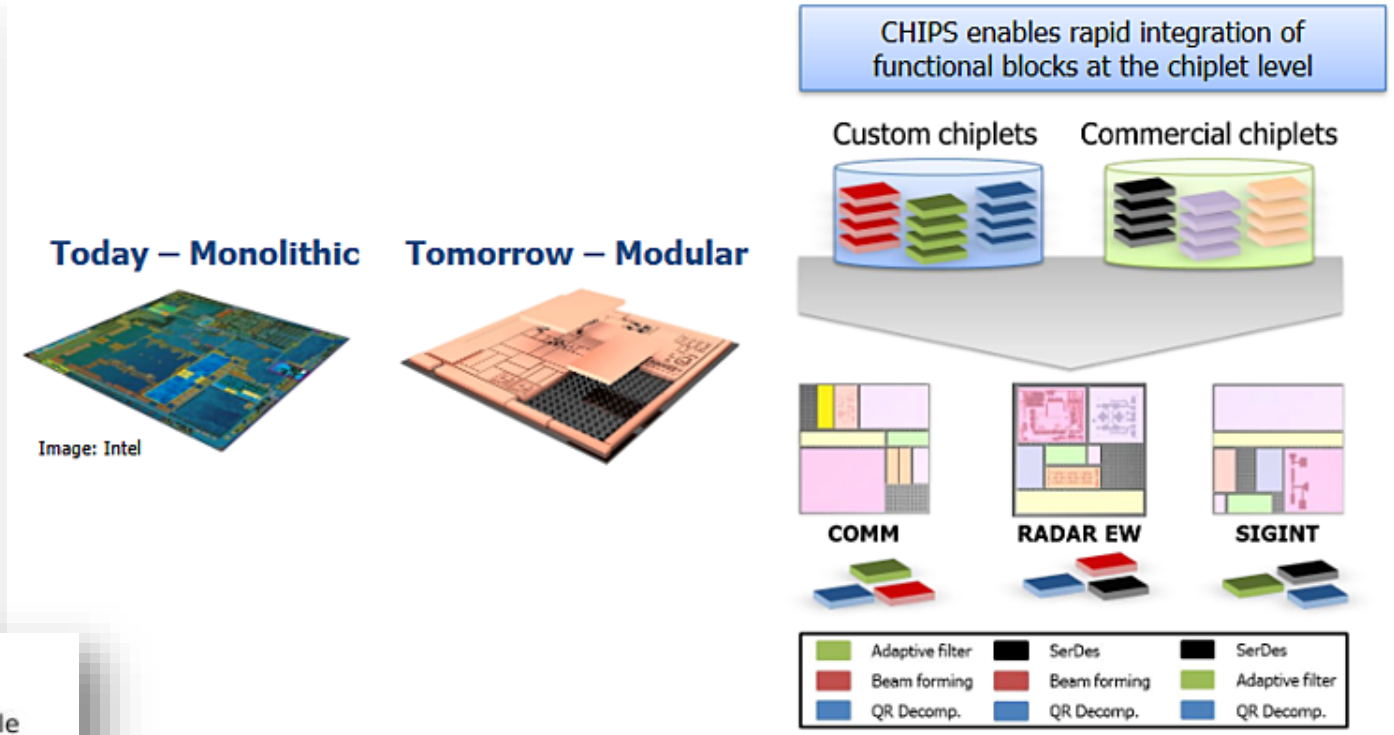

Figure 1. CHIPS Vision (DARPA)

[DARPA ERI Summit 2018]

Specialization using Standards

- Proprietary chiplets already here
- Ecosystems and standards are underway
  - CXL, UCIe, BoW, …
- Open-source Tools and IP
  - RISC-V, OpenLane, Silicon Compiler, etc
- Open foundries

- ***Codesign will be more important than ever***
- ***Extreme Heterogeneity***

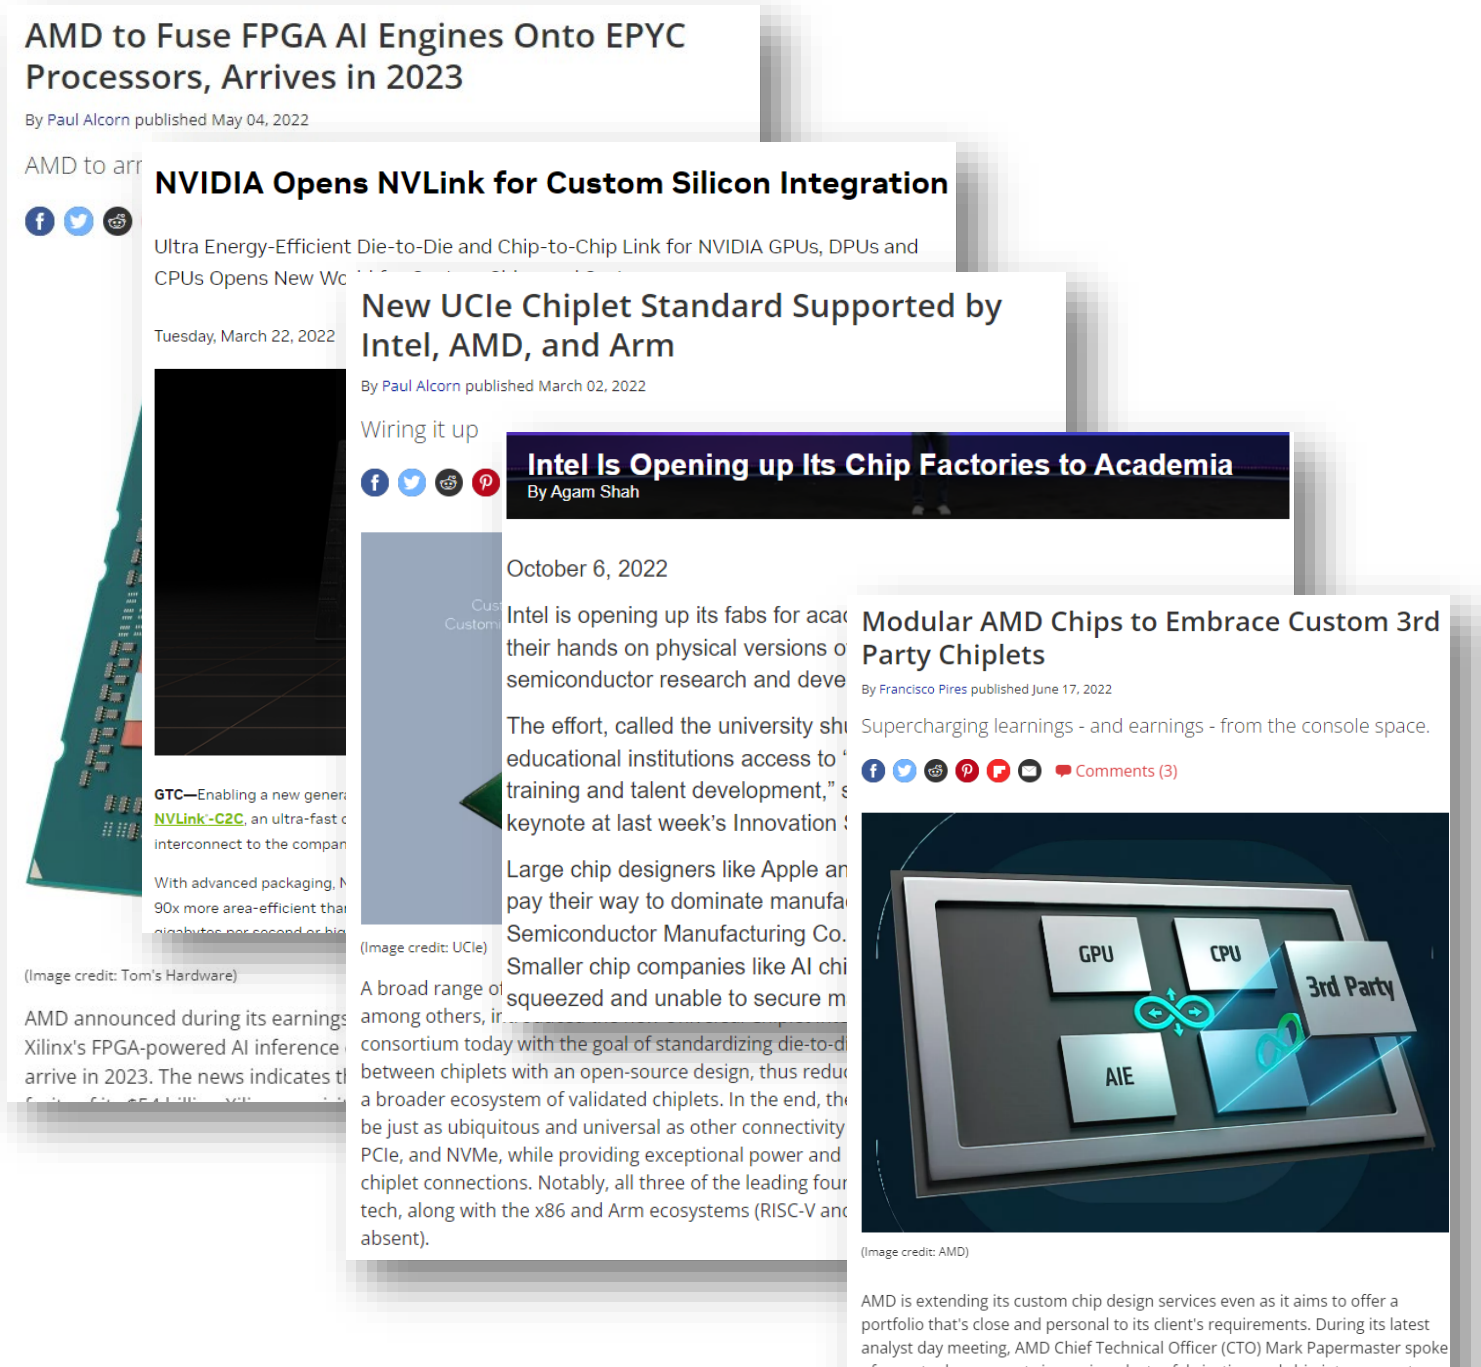

Reimagining Codesign

- Four priority research directions
  - Drive Breakthrough Computing Capabilities with Targeted Heterogeneity and Rapid Design
  - Software and Applications that Embrace Radical Architecture Diversity
  - Engineered Security and Integrity from Transistors to Applications
  - Design with Data-Rich Processes
- We must make codesign agile, more accurate, and use real workloads

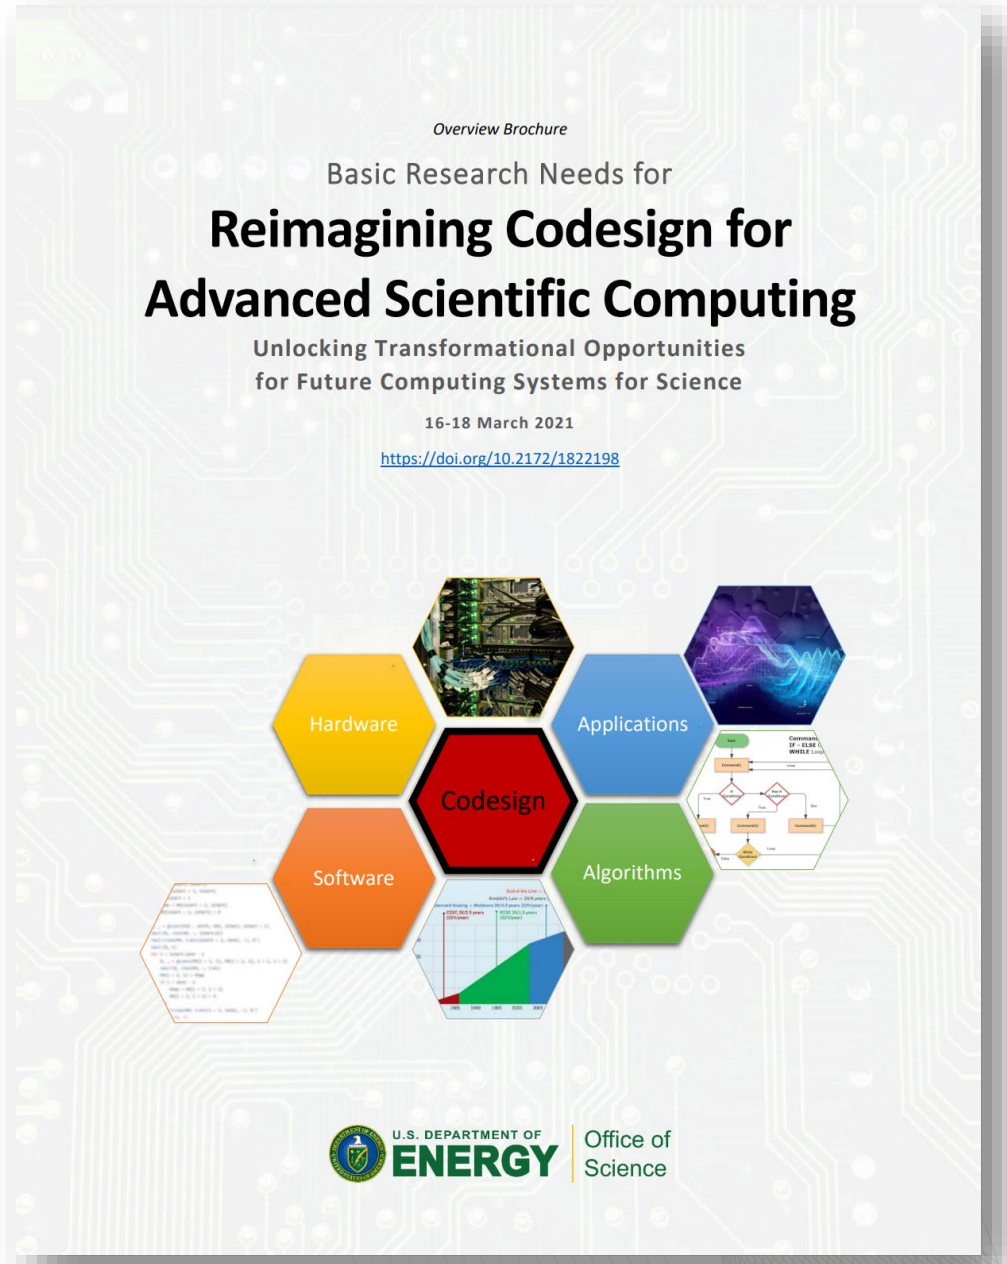

2021 Workshop
https://www.osti.gov/biblio/1822198-reimagining-codesign-advanced-scientific-computing-unlocking-transformational-opportunities-future-computing-systems-science

OAK RIDGE National Laboratory

Abisko: Microelectronics Codesign Project

- Aaron Young, Prasanna Date, David Brooks, Farah Fahim, Frank Liu, Gu-Yeon Wei, Holland Hysmith, Anton Ievlev, Kevin Cao, Shruti Kulkarni, Sung-Kyu Lim, Petro Maksymovych, Marc Gonzales Tallada, Matthew Marinella, Narasinga Rao Miniskar, Nhan Tran, Pruek Vanna-Iampikul, Catherine Schuman, Bobby Sumpter, Alec Talin, Jeffrey Vetter

OAK RIDGE National Laboratory
Sandia National Laboratories
Arizona State University
Georgia Tech
HARVARD UNIVERSITY
Collaborator
Fermilab

This research is funded by the DOE Office of Science Research Program for Microelectronics Codesign (sponsored by ASCR, BES, HEP, NP, and FES) through the Abisko Project with program managers Robinson Pino (ASCR), Hal Finkel (ASCR), and Andrew Schwartz (BES).

Abisko Microelectronics Codesign Overview

OAK RIDGE National Laboratory | Sandia National Laboratories | ASU Arizona State University | Georgia Tech | HARVARD UNIVERSITY | Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

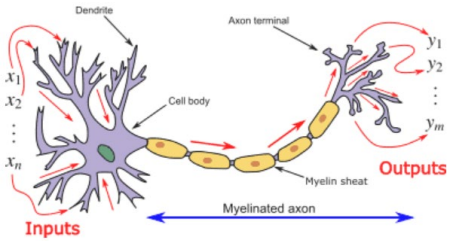

Source: Wikipedia

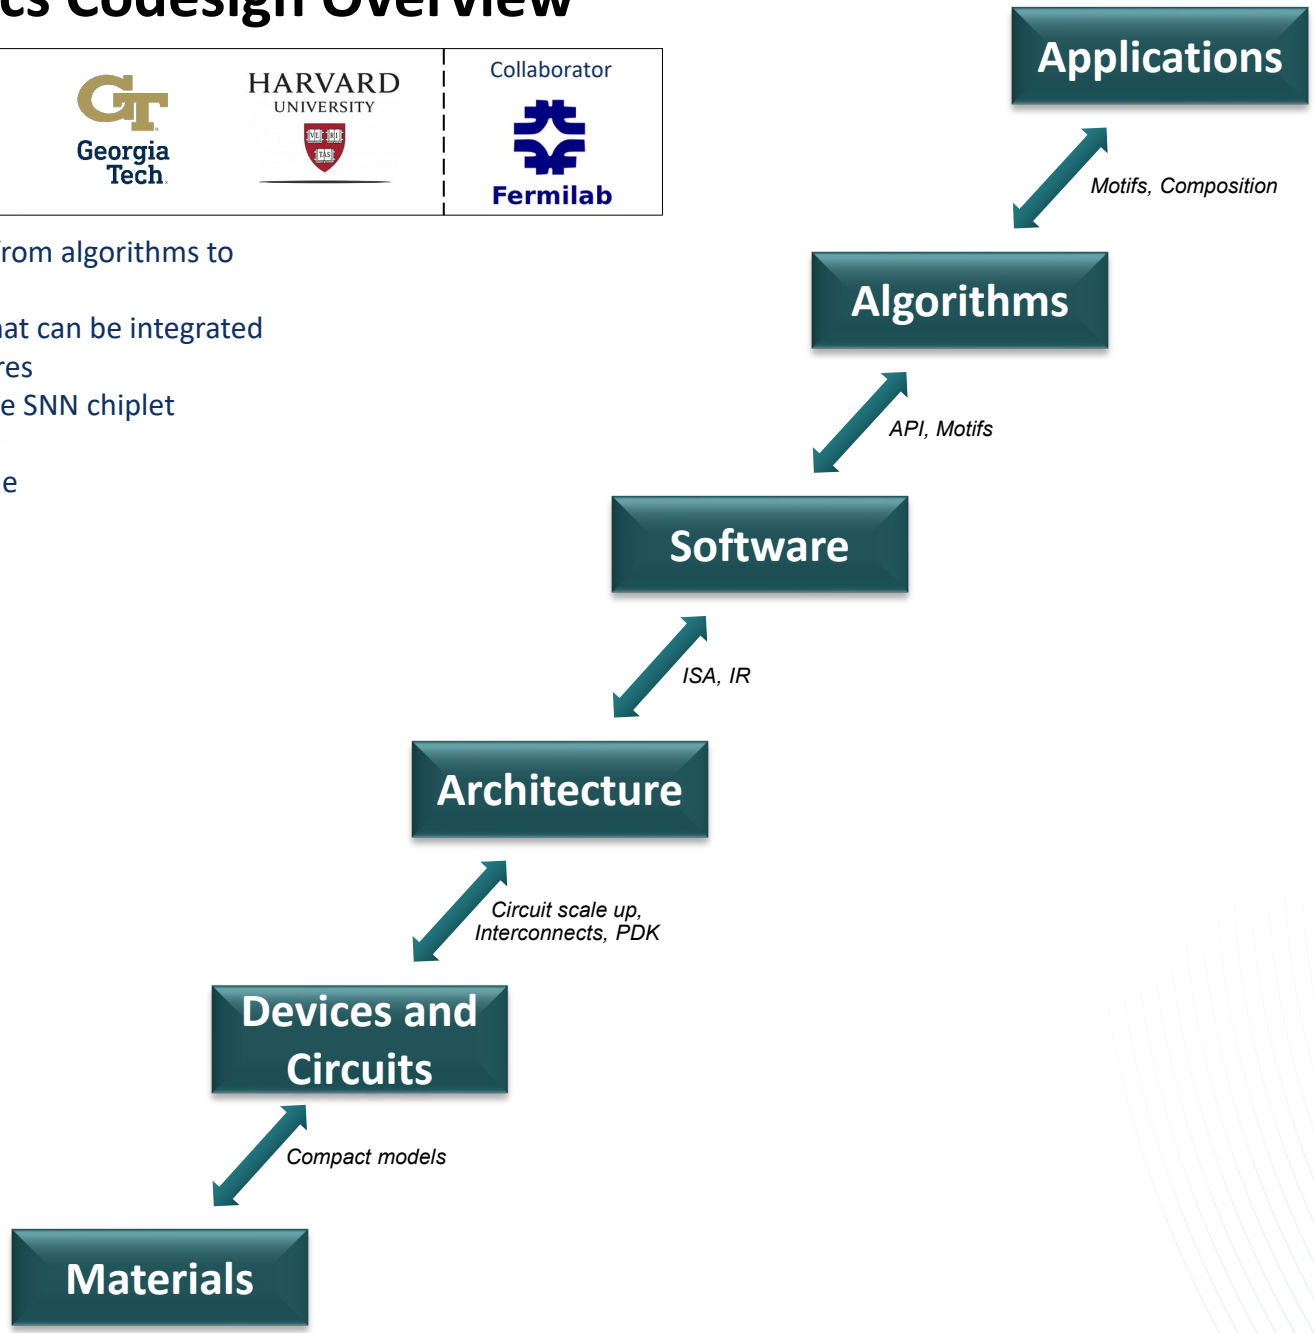

Abisko Microelectronics Codesign Overview

Oak Ridge National Laboratory | Sandia National Laboratories | Arizona State University | Georgia Tech | Harvard University

Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

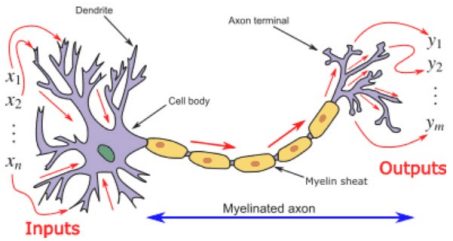

Source: Wikipedia

Motivation

- Transportation
- CMS Sensors

Applications

Motifs, Composition

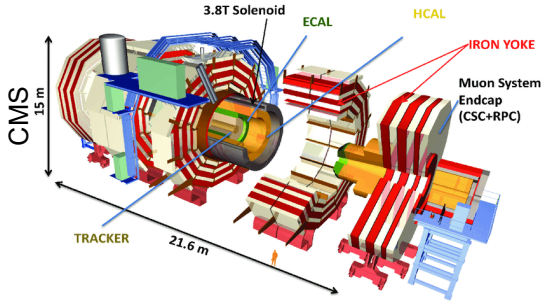

Slide courtesy of F. Fahim (FermiLab)

CMS Experiment

40MHz collision rate

~1B detector channels

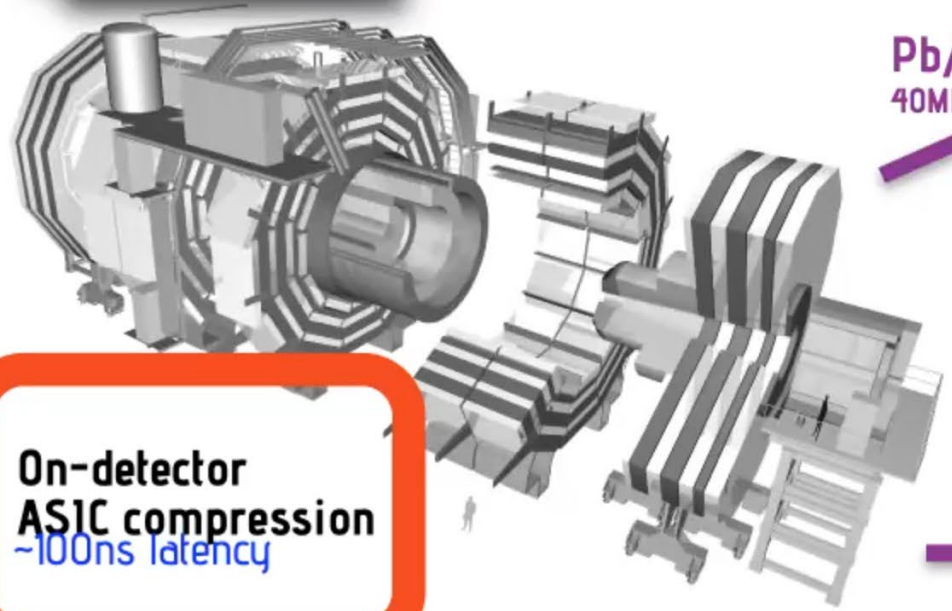

Pb/s

40MHz

FPGA filter stack

~µs latency

On-detector ASIC compression

~100ns latency

10s Tb/s

100s kHz

10s Gb/s

~5 kHz

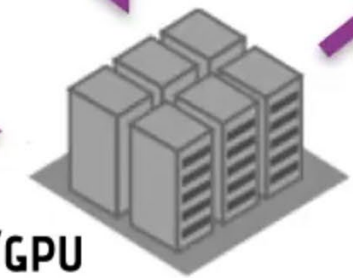

On-prem CPU/GPU filter farm

~100 ms latency

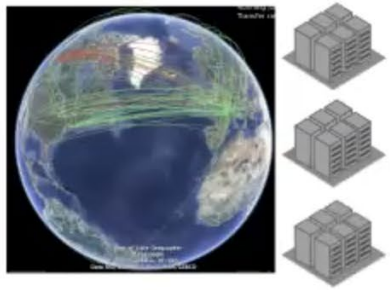

Worldwide computing grid

Exabyte-scale datasets

1 Billion channels →

10x the average internet traffic in all of North America

Slide courtesy of F. Fahim (FermiLab)

Pixel Detector: Proposed ML implementation

Digital neuromorphic implementation

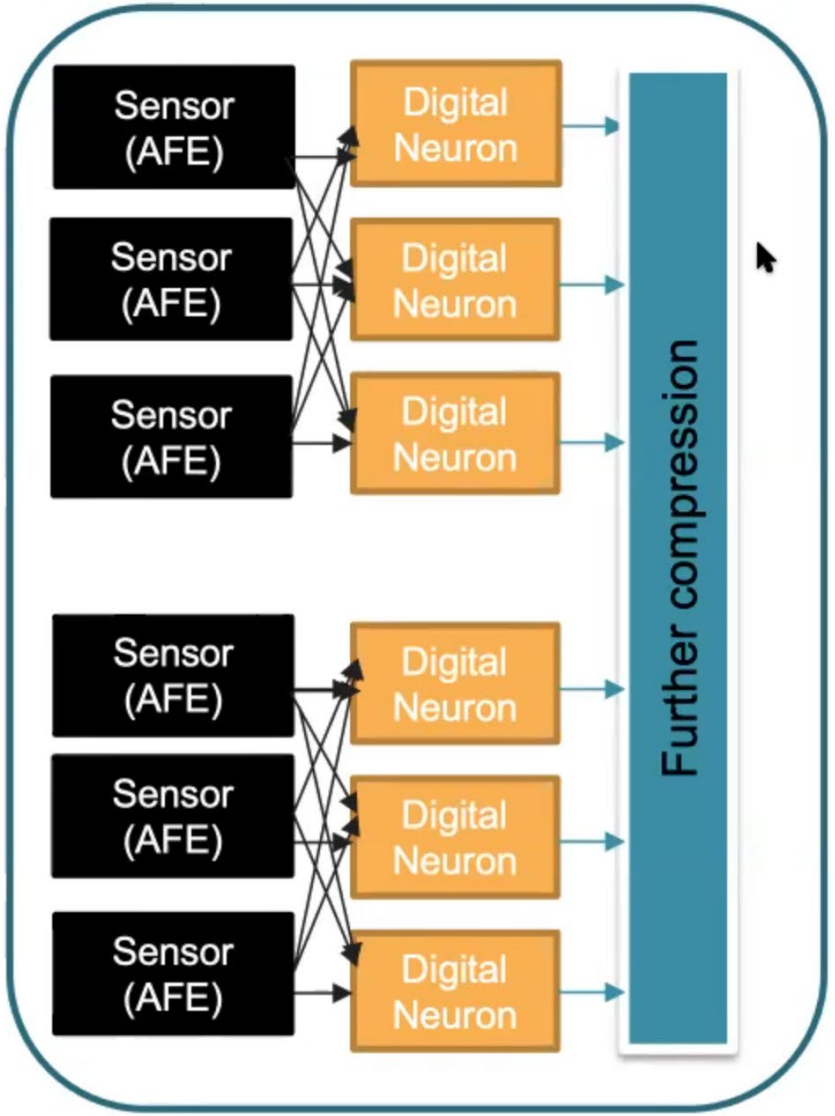

Analog – Mixed Signal implementation using floating gates or memristive cross-bar arrays

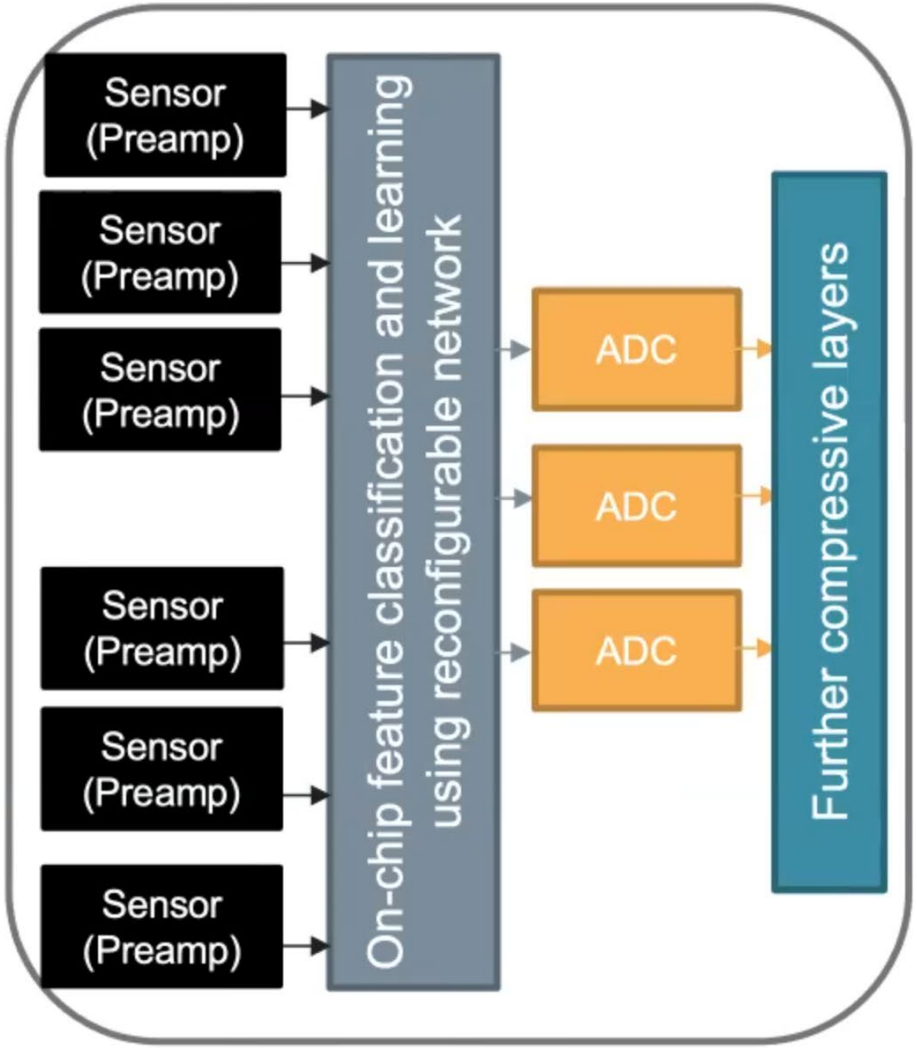

- Ability to work in the latent space (downstream resources)
- Reconfigurability vs. pruning?
- On-chip inference vs. on-chip training?
- Light weight models?
- Can lead to self calibrating detectors?

Abisko Microelectronics Codesign Overview

Oak Ridge National Laboratory | Sandia National Laboratories | Arizona State University | Georgia Tech | Harvard University | Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

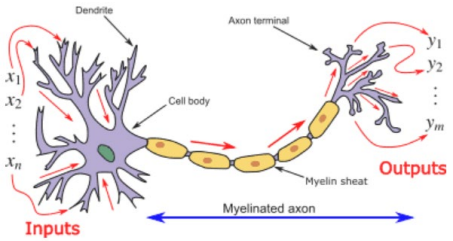

Source: Wikipedia

Motivation

- Transportation
- CMS Sensors

Applications

Motifs, Composition

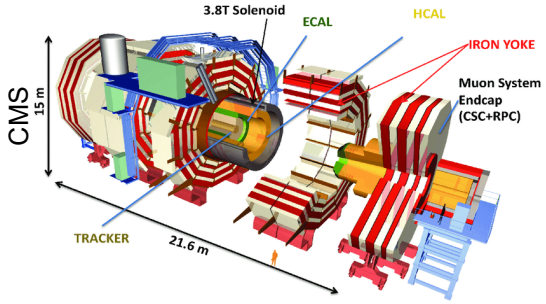

Algorithms

- ML: SLAYER, Whetstone, EONS, eProp, STDP
- Non-ML: Graph algorithms, CSP
- Simulators: NEST, Brian2

Algorithms

API, Motifs

nest :: BRIAN

Algorithms: Developing SNN Encoding and Configuration

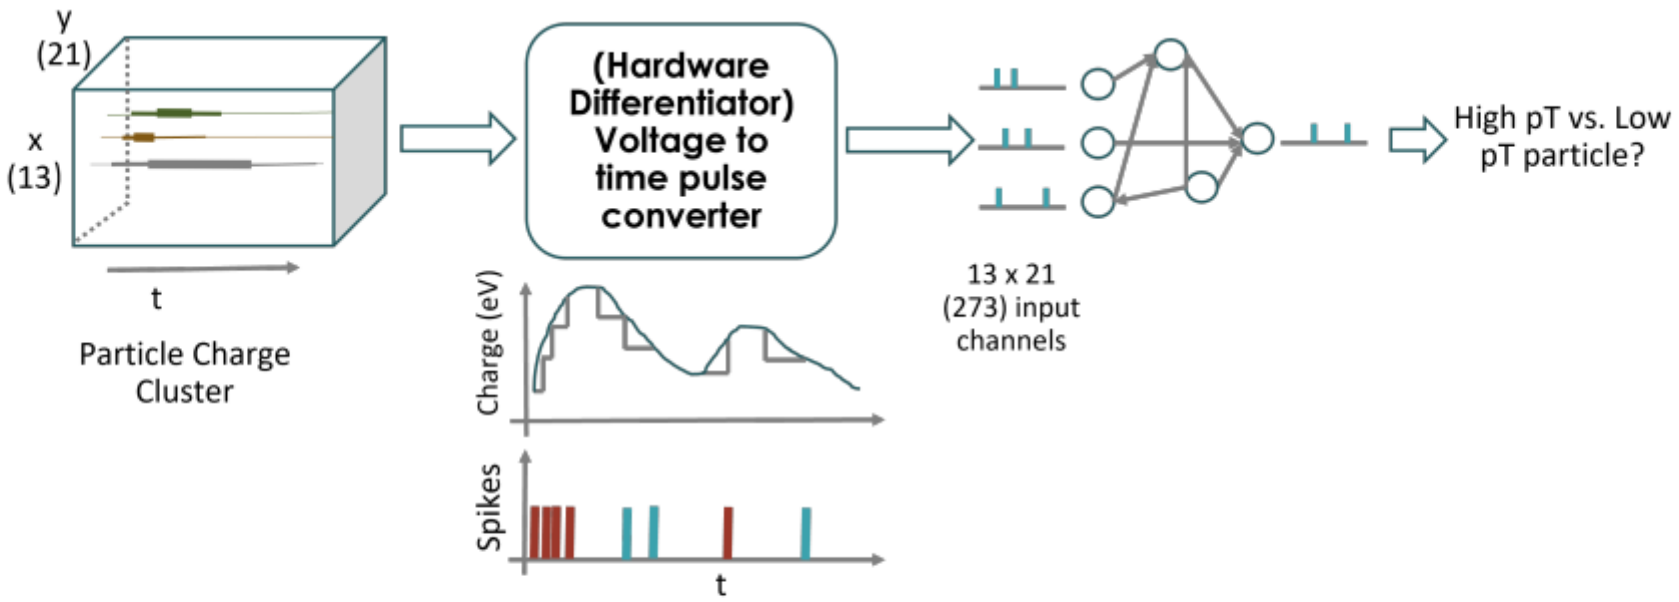

Figure 1. End-to-end in-pixel filtering of particle charge clusters into high pT or low pT samples. Each real-valued incoming signal from the 13×21 array is converted into spike streams. The inter-spike times are related to the rise and fall time of the signal waveform. There are two input spike channels per sensor waveform: one corresponding to the rising edge (in brown) and another for the falling edge (in cyan).

Abisko Microelectronics Codesign Overview

Oak Ridge National Laboratory | Sandia National Laboratories | Arizona State University | Georgia Tech | Harvard University | Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

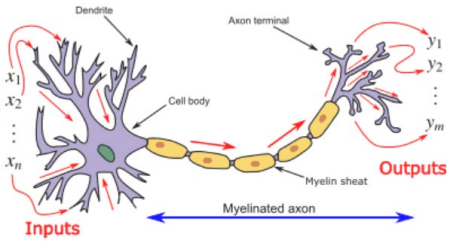

Source: Wikipedia

Applications

Motivation
- Transportation
- CMS Sensors

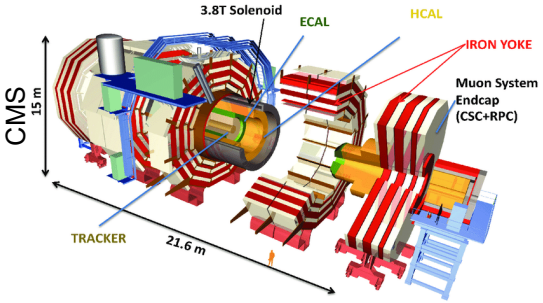

Motifs, Composition

Algorithms

Algorithms
- ML: SLAYER, Whetstone, EONS, eProp, STDP
- Non-ML: Graph algorithms, CSP
- Simulators: NEST, Brian2

nest :: BRIAN

API, Motifs

Software

Software
- DSL and API for neuromorphic co-processing
- Built on LLVM and MLIR
- Portable across Abisko chiplet, GPUs, etc.

LLVM COMPILER INFRASTRUCTURE | MLIR | XACC

ISA, IR

Example Virtual Neuron in Aurora eDSL

```
(1)   def x_pos("x_pos", Layer); // positive number X

(2)   def bits_pos("bits_pos", Layer); // positive bit neurons

(3)   use create_neuron("create_neuron", (Real=0, Real=-1.0) ->* LIF);

(4)   x_pos[0,positive_precision-1] = create_neuron(0);

(5)   bits_pos[0,positive_precision][0] = create_neuron(0);
(6)   bits_pos[0,positive_precision][1] = create_neuron(1);
(7)   bits_pos[1,positive_precision][2] = create_neuron(2);

(8)   var range(Range);
(9)   range = (0, positive_precision-1);

(10)  Connect(x_pos[range],
                bits_pos[range][0,1] = Synapse("weight"_m = 1.0,
                                               "delay"_m = Real(range+1));
(11)  Connect(x_pos[1, positive_precision-1],
                bits_pos[range][2]) = Synapse("weight"_m = 1.0,
                                              "delay"_m = Real(range+1));
```

Collapse Section

Figure 6. Extract of Virtual Neuron graph specification in *Aurora*. Layers are defined and combined with *Range* type variables, which generates *Views* of the nodes in the graph layers. Usage of *Connect* operator to connect elements in the graph. Usage of *Synapse* type, which is a derived type from *EdgeType* native type in *Aurora*.

```
def("create_neuron", (Real=0, Real=-1.0)  ->* LIF,

  Begin(in V_th, in internal_state) {

    var neuron(LIF);

    neuron["V_th"] =  V_th;
    neuron["V_m"] = internal_state;
    neuron["tau_m"] = -1e-6;
    neuron["t_ref"] = 0.0;
    neuron["E_L"] = neuron["V_m"];

    Return(neuron);

  } ); // def
```

Figure 7. *Aurora* definition of a function to generate a LIF instance and update internal parameters of the LIF data model.

Abisko Microelectronics Codesign Overview

Oak Ridge National Laboratory | Sandia National Laboratories | Arizona State University | Georgia Tech | Harvard University | Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

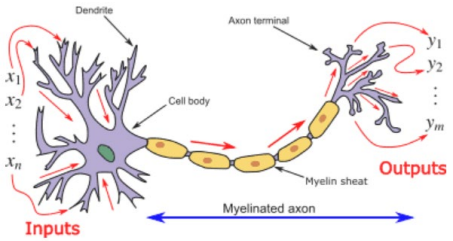

Source: Wikipedia

Motivation
- Transportation
- CMS Sensors

Applications

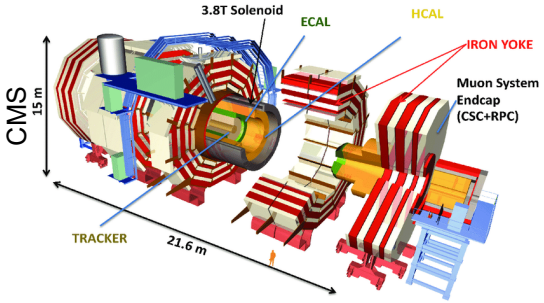

Motifs, Composition

Algorithms
- ML: SLAYER, Whetstone, EONS, eProp, STDP
- Non-ML: Graph algorithms, CSP
- Simulators: NEST, Brian2

Algorithms — nest :: — BRIAN

API, Motifs

Software
- DSL and API for neuromorphic co-processing
- Built on LLVM and MLIR
- Portable across Abisko chiplet, GPUs, etc.

Software — LLVM COMPILER INFRASTRUCTURE — MLIR — XACC

ISA, IR

Architecture
- Design neuromorphic chiplet
- RISC-V neuromorphic extensions
- Heterogeneous integration with contemporary technologies

Architecture — RISC-V

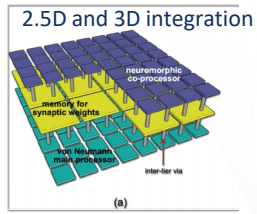

Circuit scale up, Interconnects, PDK

Simulation/Emulation: ALADDIN — gem5

Chiplet Architectures

- Design an (analog) SNN chiplet that can be easily integrated with contemporary technologies
  - Heterogeneous integration with mixed processes
  - Compatible with existing processes
- Extensive advances in chiplets, packaging, and heterogeneous integration recently
  - Open Domain-Specific Architecture
  - UCIe, BoW, TSMC SoIC-CoW, Intel Foveros
- Using open architecture to explore chiplet designs: RISC-V, OpenLane

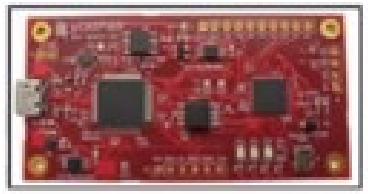

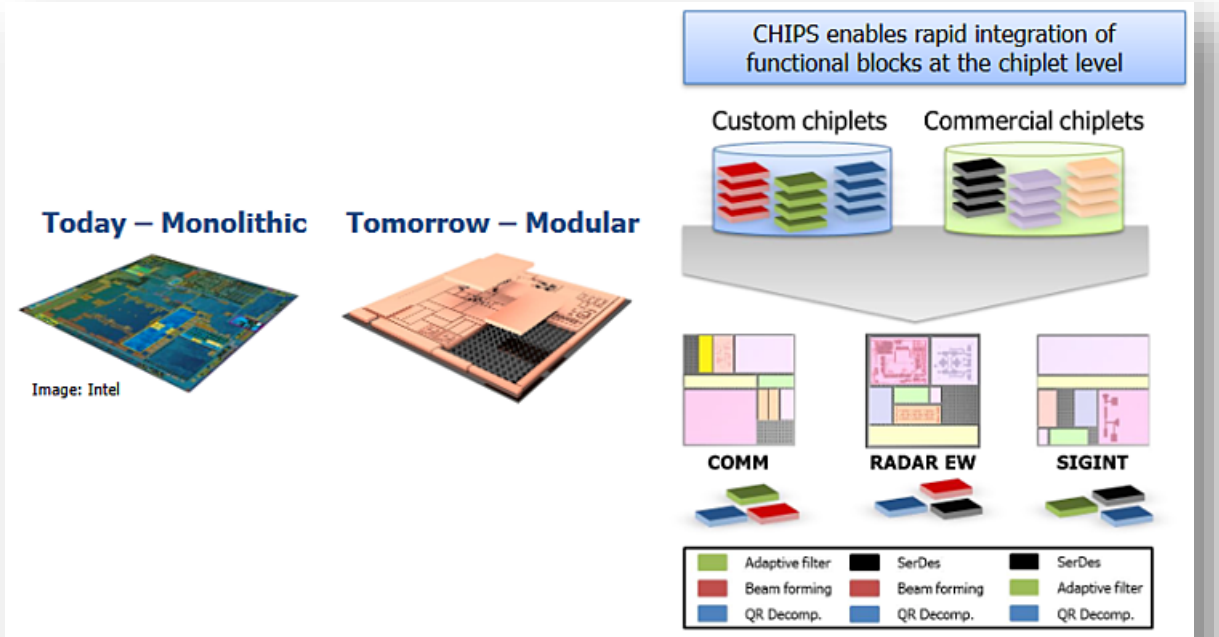

Figure 1. CHIPS Vision (DARPA)

[DARPA ERI Summit 2018]

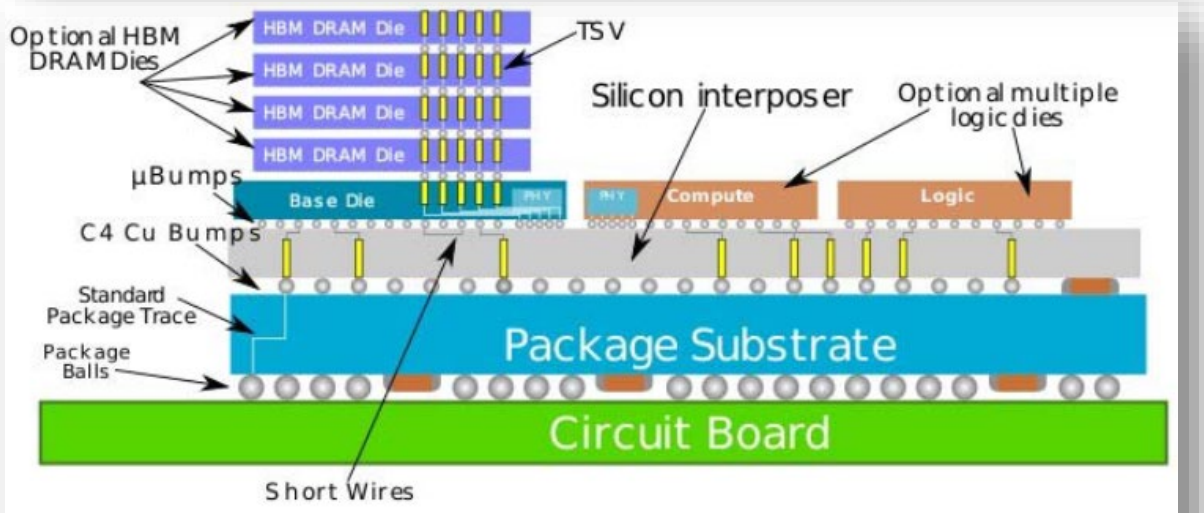

Figure 21. An example showing the use of 2D and 3D interconnections (courtesy TSMC)

[IEEE HIR 2021]

Design of 2.5D Chiplet for Neuromorphic Computing (1)

Element Types:

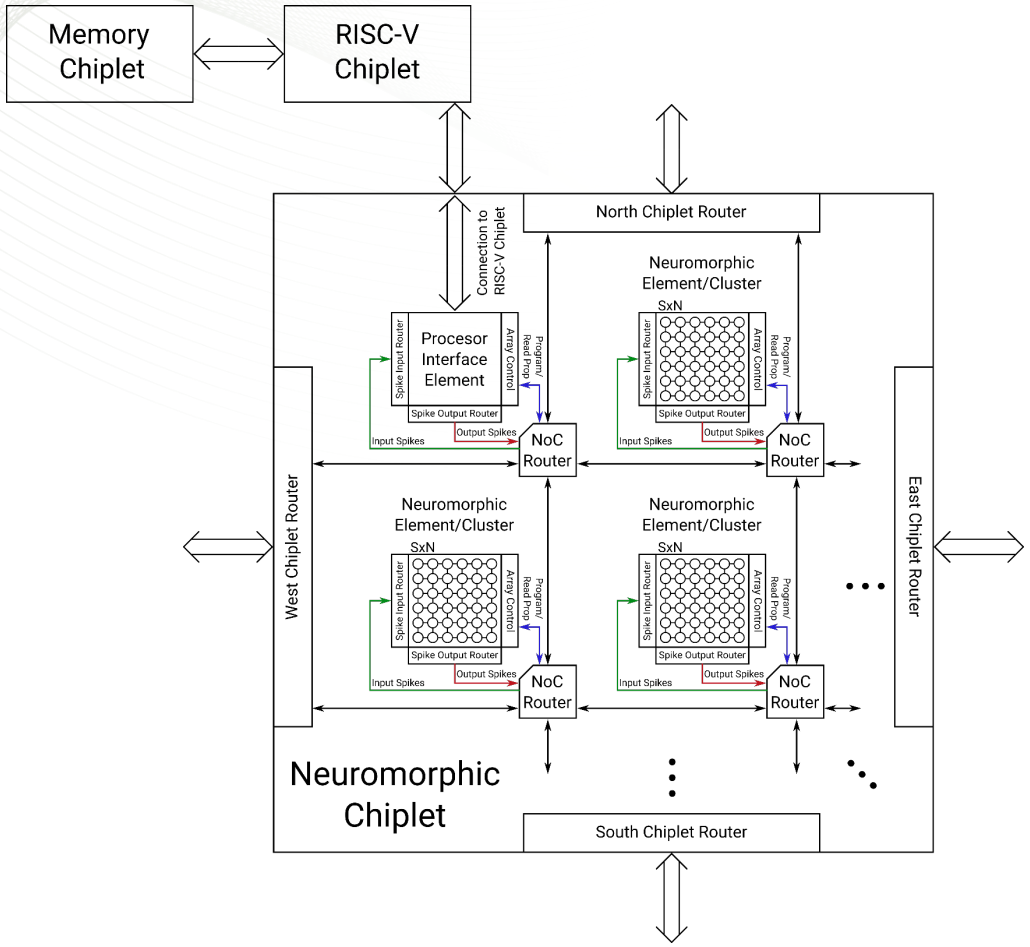

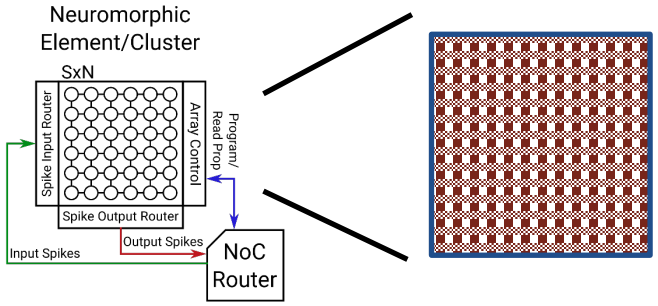

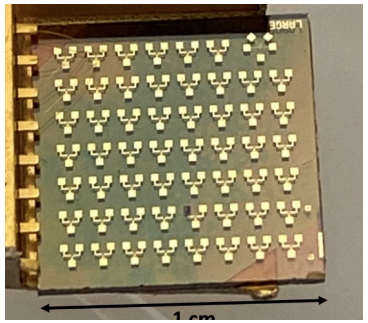

Sandia VO_2 ECRAM

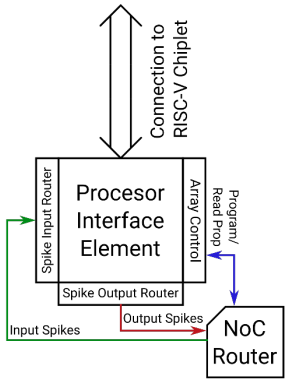

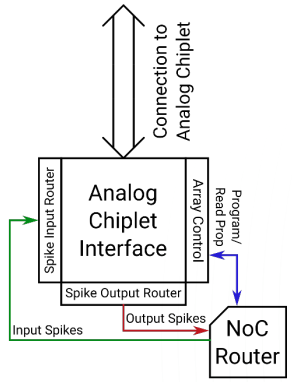

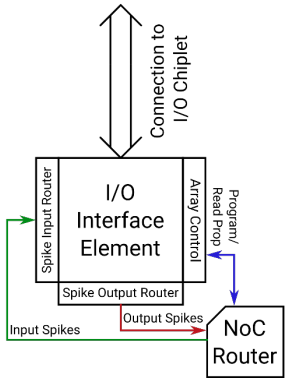

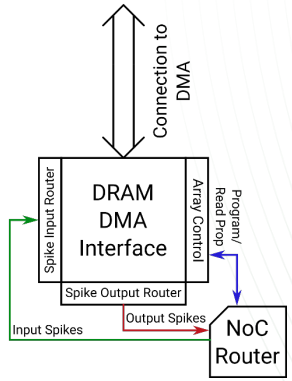

Design of 2.5D Chiplet for Neuromorphic Computing (2)

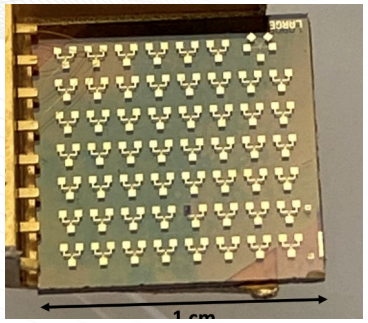

Sandia VO_2 ECRAM

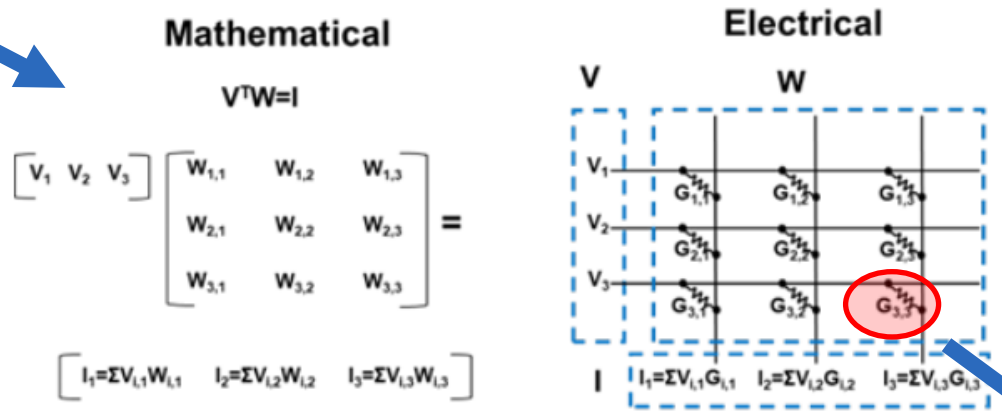

Figure 17. (a) Mathematical and (b) electrical vector matrix multiplication (Marinella et al. 2022).

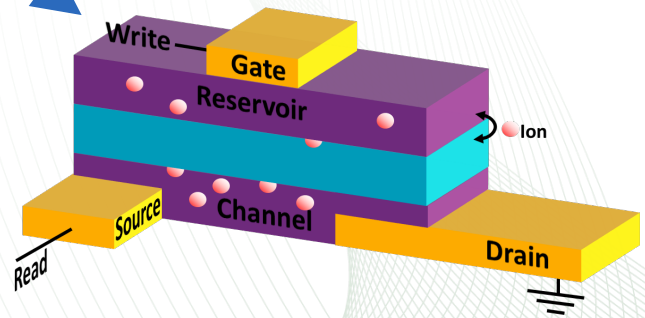

ASIC Flow for Digital NN Inference (baseline)

- Investigate the performance of fully customized ASIC design for ultra-fast NN inference
- Model details:
  - Fixed NN architecture with quantized weights
  - Experimented with 2bit or 3bit of inputs (limited by FermiLab implementation)
- Flow:
  - Vitis HLS to generate RTL
  - Catapult logic synthesis
  - Customized backend layout tool (incl. tech mapping, placement and routing)
- Achieved clock frequency of 1~ 2GHz in a 28nm technology

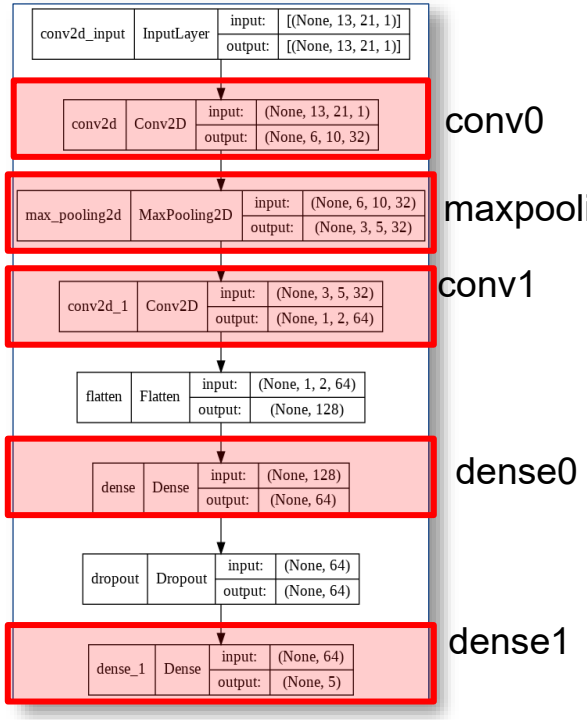

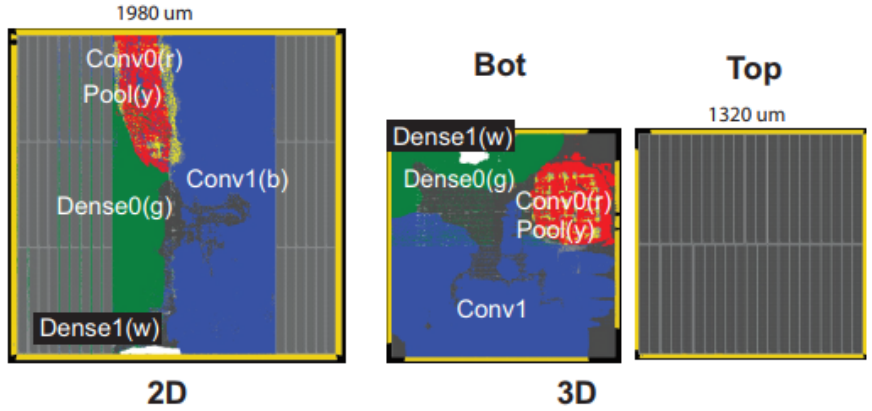

Figure 15. Placement comparison between 2D and 3D designs. SRAM arrays are placed on the sides in 2D and on the top tier in the 3D case. The modules are colored as conv0 (red), pool (yellow), conv1 (blue), dense0 (green), and dense1 (white).

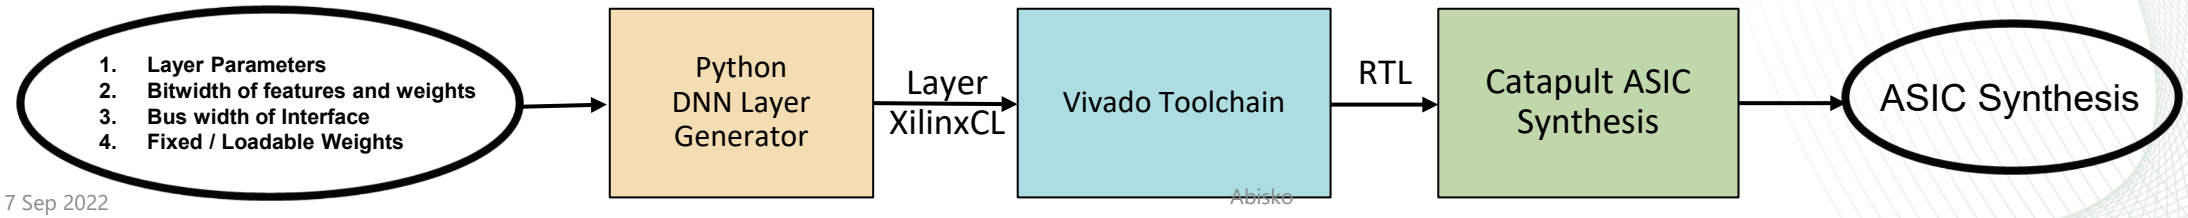

Abisko Microelectronics Codesign Overview

OAK RIDGE National Laboratory · Sandia National Laboratories · Arizona State University · Georgia Tech · HARVARD UNIVERSITY

Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

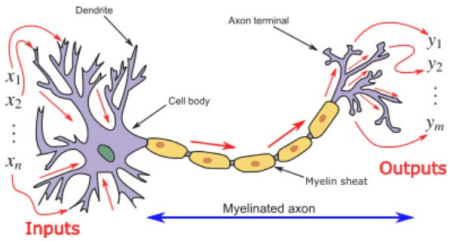

Source: Wikipedia

Applications

Motivation
- Transportation
- CMS Sensors

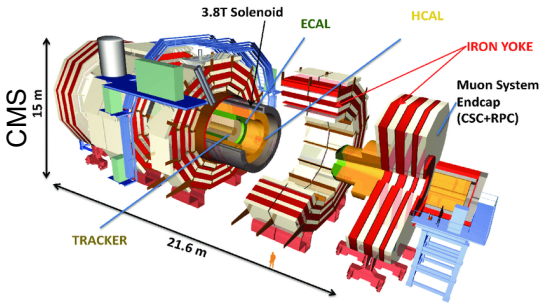

Motifs, Composition

Algorithms

Algorithms
- ML: SLAYER, Whetstone, EONS, eProp, STDP
- Non-ML: Graph algorithms, CSP
- Simulators: NEST, Brian2

nest :: · BRIAN

API, Motifs

Software

Software
- DSL and API for neuromorphic co-processing
- Built on LLVM and MLIR
- Portable across Abisko chiplet, GPUs, etc.

LLVM COMPILER INFRASTRUCTURE · MLIR · XACC

ISA, IR

Architecture

Architecture
- Design neuromorphic chiplet
- RISC-V neuromorphic extensions
- Heterogeneous integration with contemporary technologies

RISC-V

2.5D and 3D integration

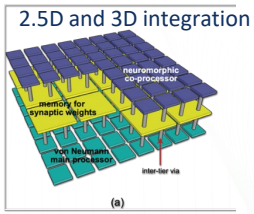

Simulation/Emulation: ALADDIN · gem5

MesaFAB ReRAM

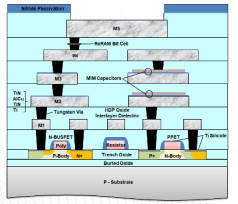

Circuit scale up, Interconnects, PDK

Devices and Circuits

Devices and Circuits
- ion insertion (reversible doping) sets analog states
- mRaman captures transition linear, non-linear switching
- Will extend to 36x36 x-bar array
- Electronic and other optical spectroscopies

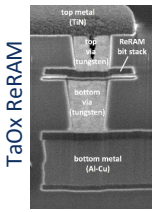

TaOx ReRAM

ECRAM

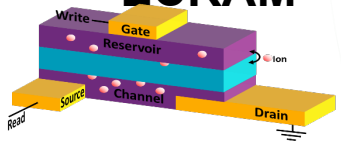

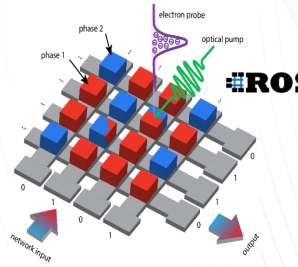

ROSS SIM

Compact models

Devices and Circuits

- Goals
  - Harness the interplay between mobile defects (ions and vacancies) and electronic properties to realize functional elements for spiking and non-spiking analog neuromorphic networks
  - Create and validate small network models; generate device and network data for co-design
  - Understand and mitigate radiation induced degradation mechanisms at the device and circuit level

Devices

1) ECRAM

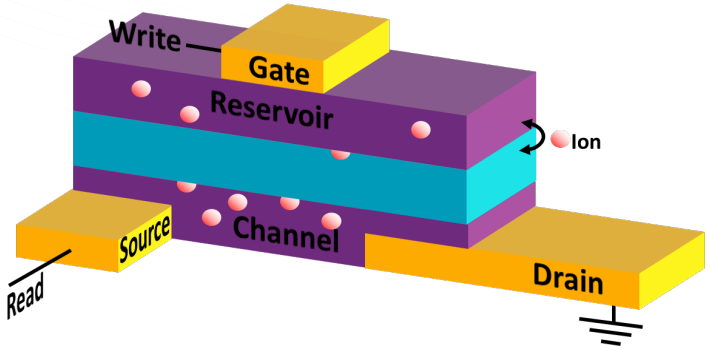

2) ReRAM

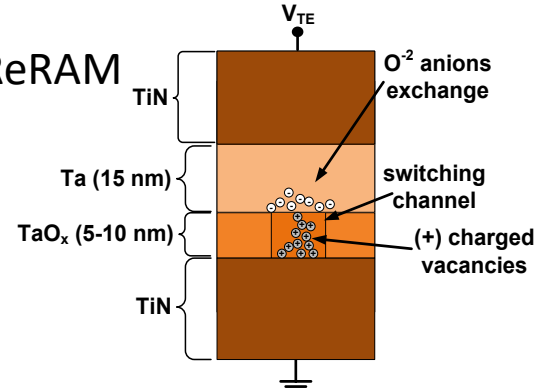

Circuits

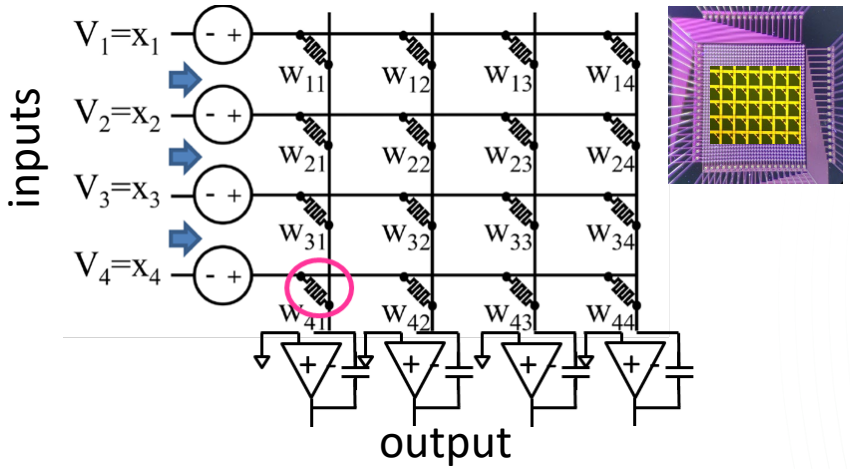

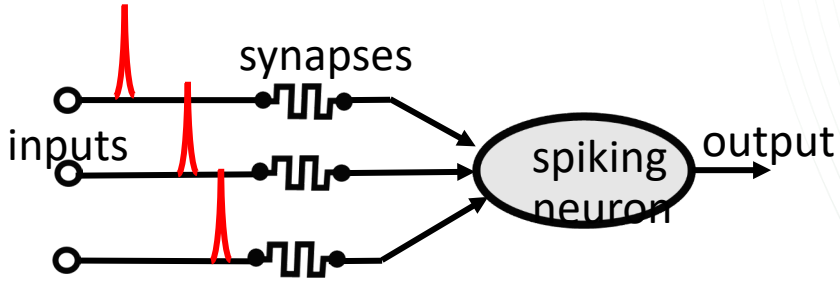

Prototypes

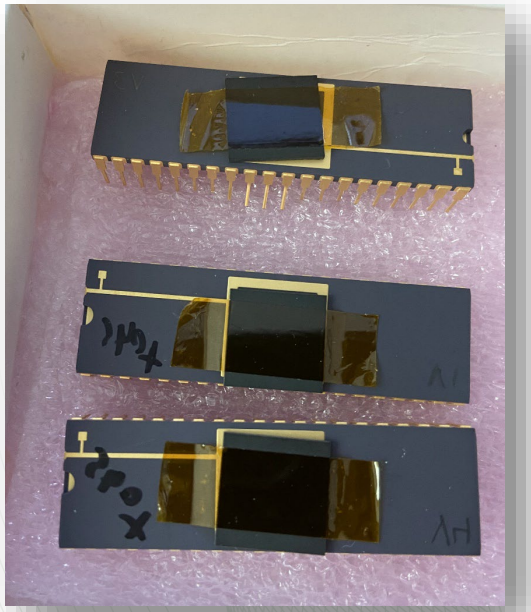

Experimental TaOx ReRAM Conductance Distributions

Developed TaOx weight mapping and programming routine for optimizing inference accuracy

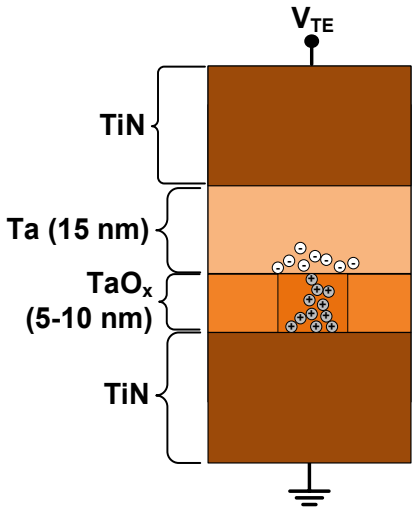

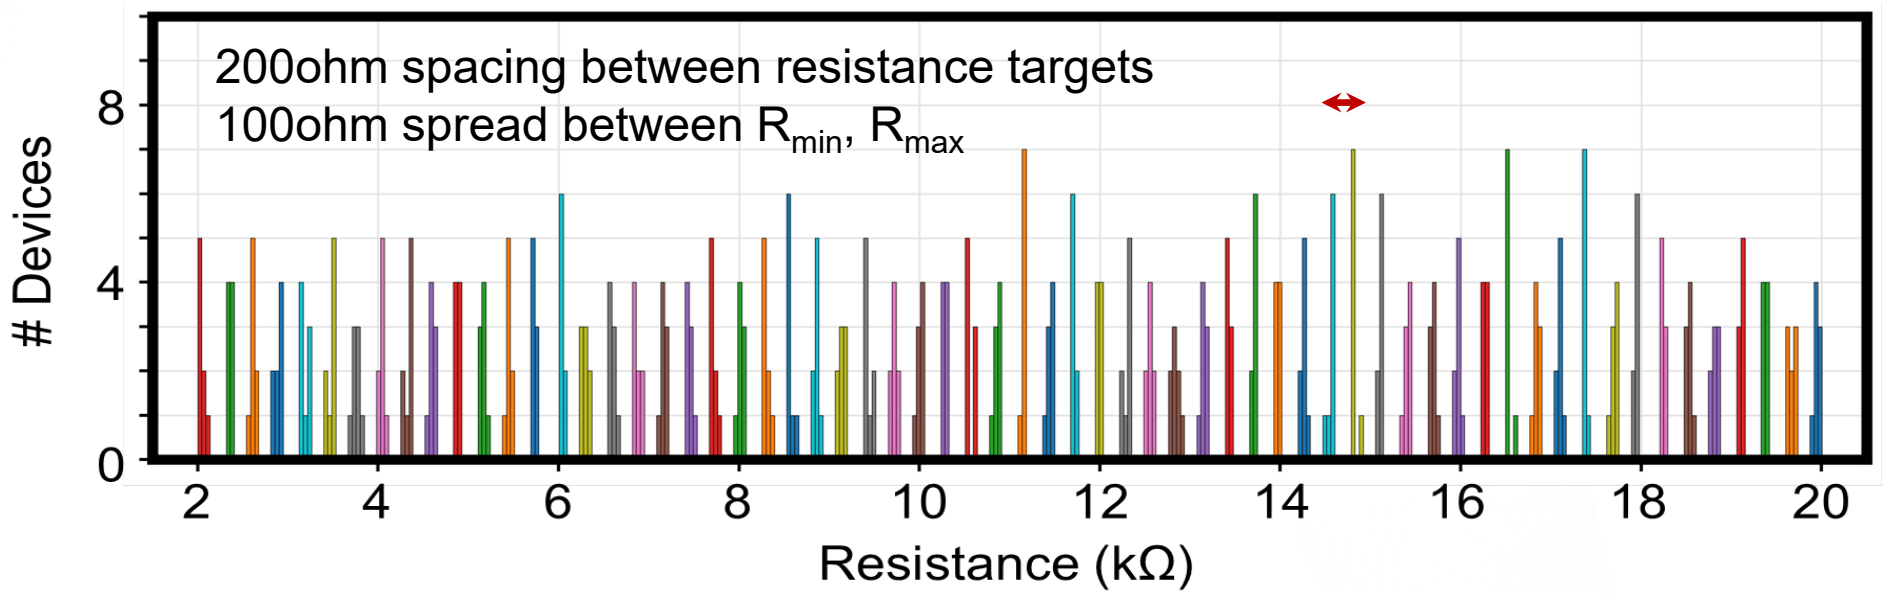

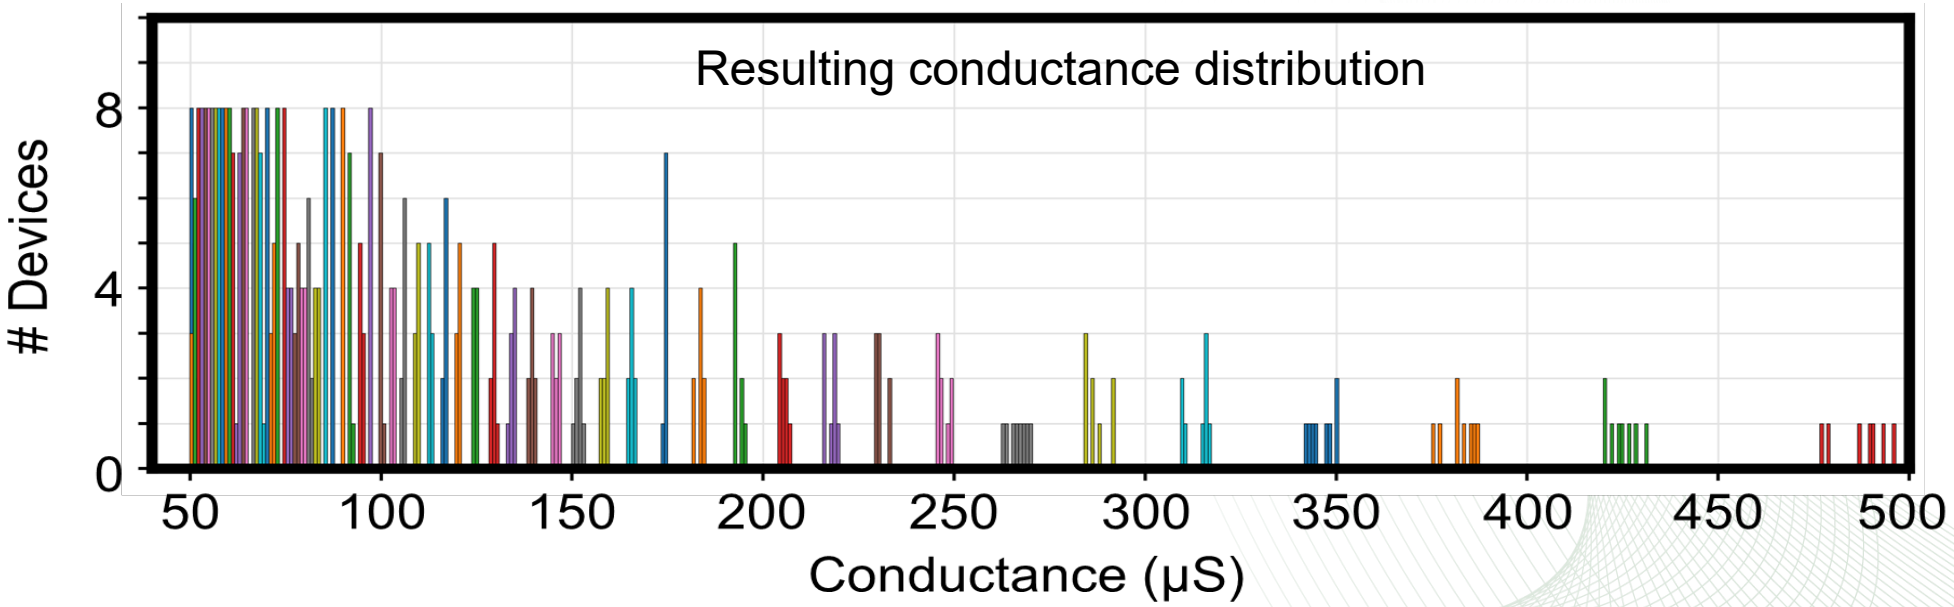

Abisko Microelectronics Codesign Overview

OAK RIDGE National Laboratory · Sandia National Laboratories · Arizona State University · Georgia Tech · HARVARD UNIVERSITY

Collaborator: Fermilab

1. Develop better techniques for codesign from algorithms to devices and materials
2. Design Spiking Neural Network chiplet that can be integrated with contemporary computer architectures
3. Explore new devices and materials for the SNN chiplet (neuron, synapse, plasticity, etc.)
4. Design language abstractions and runtime support for SNN chiplet

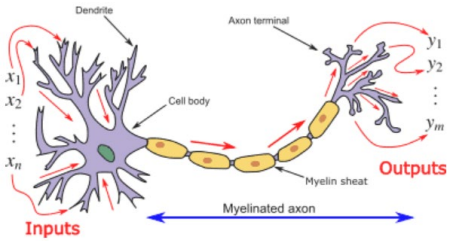

Source: Wikipedia

Applications

Motivation
- Transportation
- CMS Sensors

Motifs, Composition

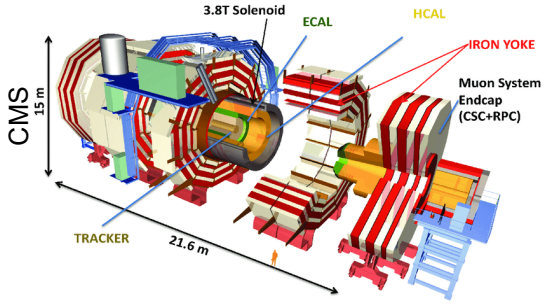

Algorithms

Algorithms
- ML: SLAYER, Whetstone, EONS, eProp, STDP
- Non-ML: Graph algorithms, CSP
- Simulators: NEST, Brian2

nest :: · BRIAN

API, Motifs

Software

Software
- DSL and API for neuromorphic co-processing
- Built on LLVM and MLIR
- Portable across Abisko chiplet, GPUs, etc.

LLVM COMPILER INFRASTRUCTURE · MLIR · XACC

ISA, IR

Architecture

Architecture
- Design neuromorphic chiplet
- RISC-V neuromorphic extensions
- Heterogeneous integration with contemporary technologies

RISC-V

2.5D and 3D integration

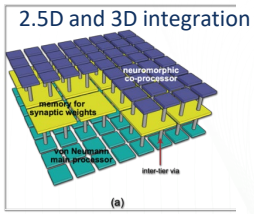

Simulation/Emulation: ALADDIN · gem5

Circuit scale up, Interconnects, PDK

Devices and Circuits

Devices and Circuits
- ion insertion (reversible doping) sets analog states
- mRaman captures transition linear, non-linear switching
- Will extend to 36x36 x-bar array
- Electronic and other optical spectroscopies

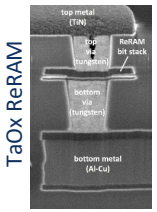

TaOx ReRAM

ECRAM

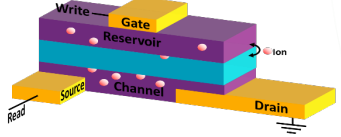

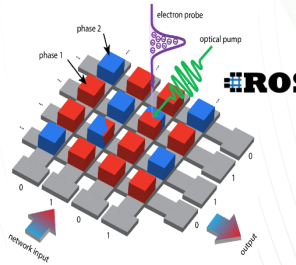

ROSS SIM

MesaFAB ReRAM

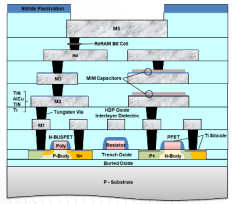

Compact models

Materials

Materials
- Non-equilibrium probes to few nm
- Data-driven modeling
- On-demand neuromorphism

Domain wall memristor

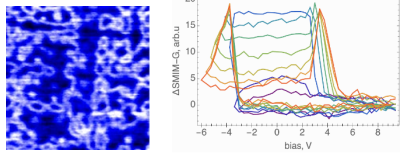

Computational data mining

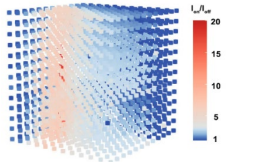

CNMS scanning probe microscopy and chemical imaging

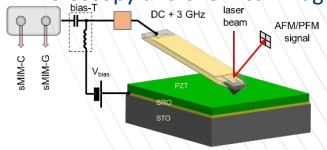

Computing Discovery Platform

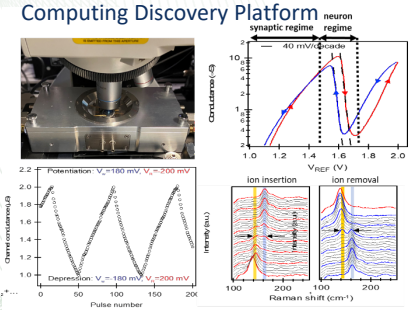

Characterizing Candidate Materials for Neuromorphic Computing

Microwave Microscopy of VO_{2-x}

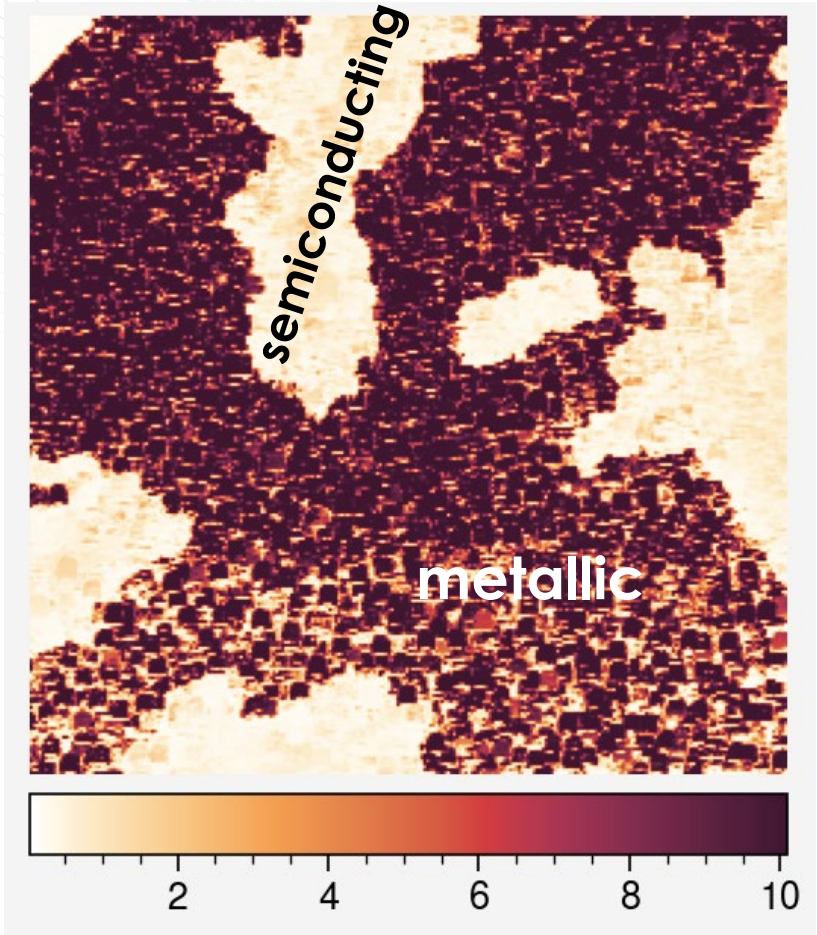

3 GHz conductance of ECRAM channel

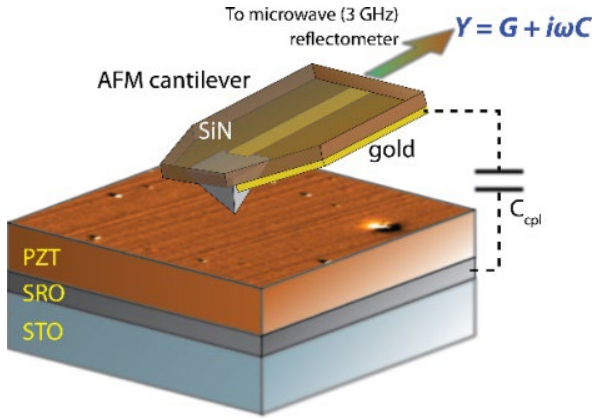

Capacitance+ conductance at 3GHz

- Quantitative
- Interpretable
- Non-invasive

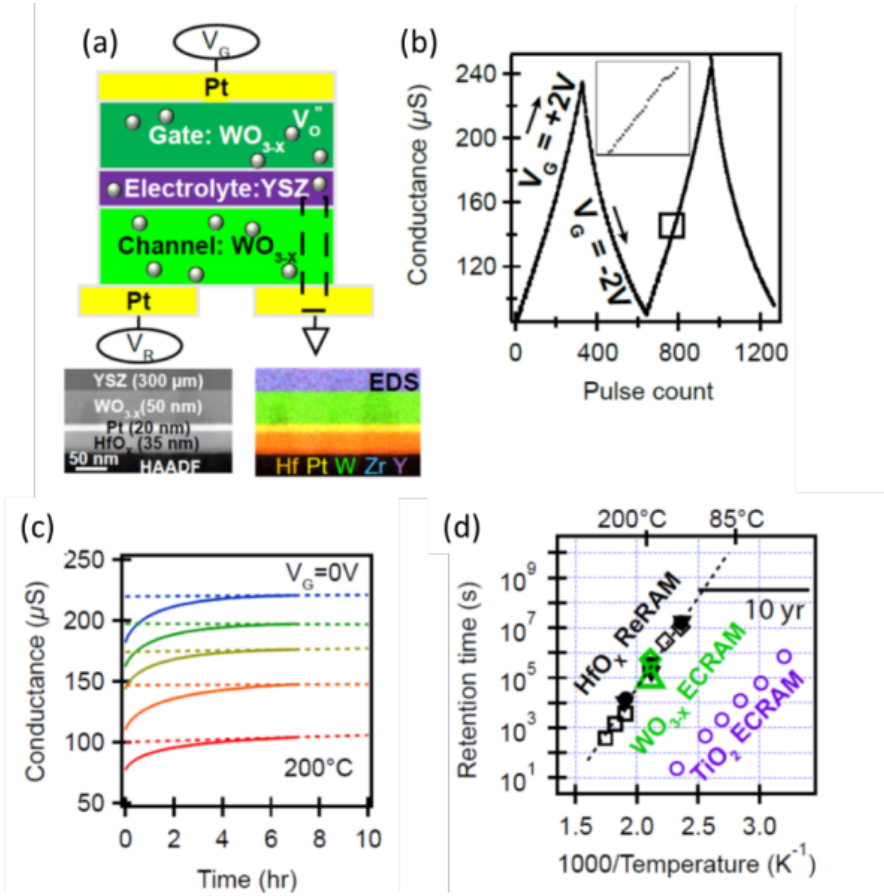

Figure 20. ECRAM device. (a) Schematic and cross-section TEM of WO_{3-x} ECRAM cell. (b) Analog switching characteristics that demonstrate high state density. (c) ECRAM retention characteristics when the gate and channel are shorted at 200°C. (d) Comparison of retention times of WO_{3-x} ECRAM with filament-based ReRAM and past TiO_x-based ECRAM. (Yi et al. 2022) Copyright 2022, Wiley.

OAK RIDGE National Laboratory

Conclusions

Recap

- Exascale is here!
- Our predictions were reasonably accurate, but we completely missed some
  - AI/ML
  - Programming systems remain major challenge
- Post-exascale
  - Heterogeneous integration and Chiplet architectures are vastly diversifying the architectural landscape
  - Post exascale will be accelerated by recent major semiconductor investments
- Deep codesign is critical moving beyond Exascale
  - Abisko is a new microelectronics codesign project developing a chiplet for analog SNN

- Visit us (post COVID ☺)
  - We host interns and other visitors year round
    - Faculty, grad, undergrad, high school, industry
- Jobs at ORNL
  - Visit https://jobs.ornl.gov
- Contact me vetter@ornl.gov
- Experimental Computing Lab
  - Lots of emerging archs
  - https://excl.ornl.gov