Neuromorphic Computing: How the Brain can inspire computing from HPC to the Edge

2023

Salishan Conference on High Speed Computing

sgcardw@sandia.gov

Suma George Cardwell

Cognitive and Emerging Computing

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.SAND2023 -02582O

FUNDAMENTAL CHALLENGES IN COMPUTING

- Limits of scaling have ushered in the "Golden Age of Computer Architecture"
  Hennessy & Patterson 2019
- Inefficiency of generality
- Performance saturation

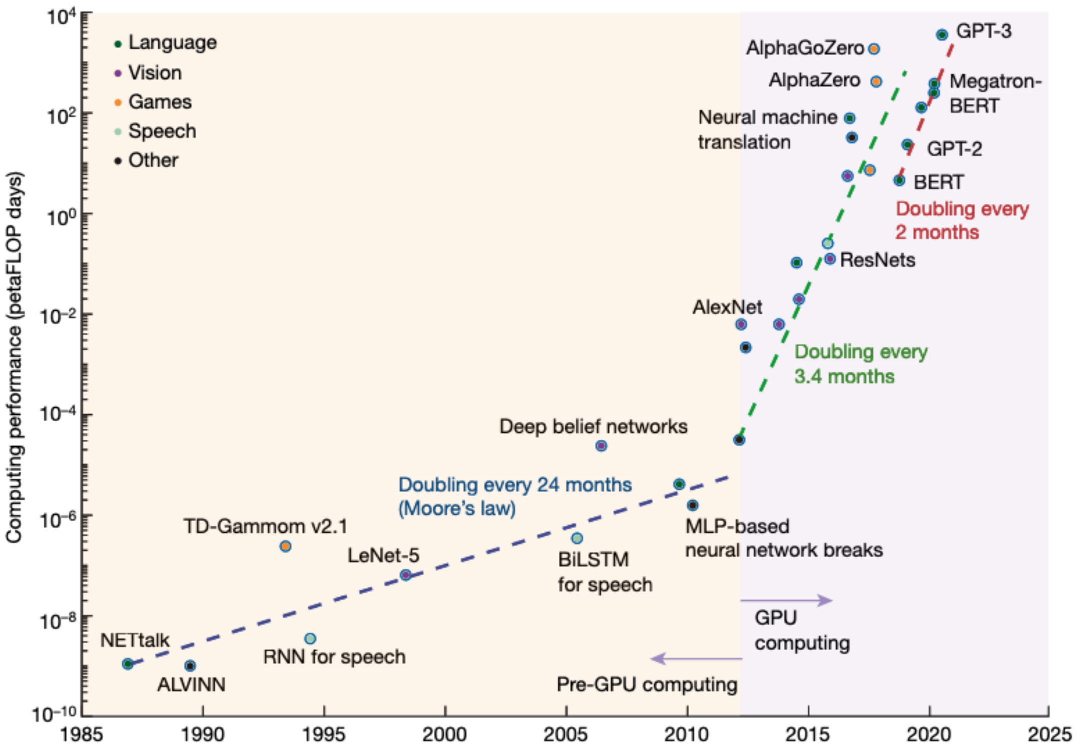

AI compute demands are increasing

Mehonic & Kenyon 2022, Open AI Research Blog

Neuromorphic Computing gives a path forward for power efficiency scaling and meeting future computing needs.

COMPUTING LANDSCAPE

Sensors

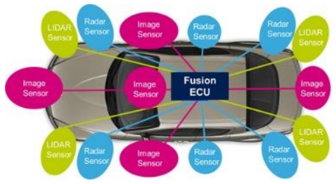

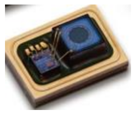

50 billion IoT devices by 2030

Algorithms

- Scientific Computation
- Machine Learning
- Brain-derived algorithms
- Signal Processing

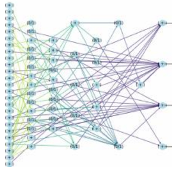

Applications

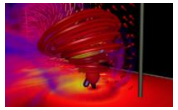

Modern Computing

Conventional Digital

Novel Computing Paradigms

CPUs

GPUs

FPGAs

TPUs

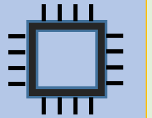

ASICs

Neuromorphic

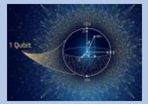

Quantum

Digital Neuromorphic

Intel Loihi
Davies 2018

SpiNNaker
Furber 2016

Analog/Mixed-signal Neuromorphic

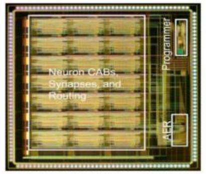

GT Neuron
Brink 2013

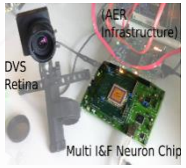

DAVIS 240C, DYNAPSEL

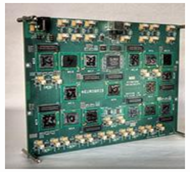

NeuroGrid
Benjamin 2014

Beyond CMOS devices

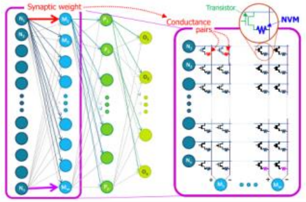

RERAM
Marinella et al., 2016

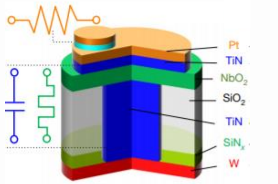

Mott- Memristor
Kumar et al., 2020

4

FUTURE OF COMPUTING: HETEROGENEOUS ARCHITECTURES

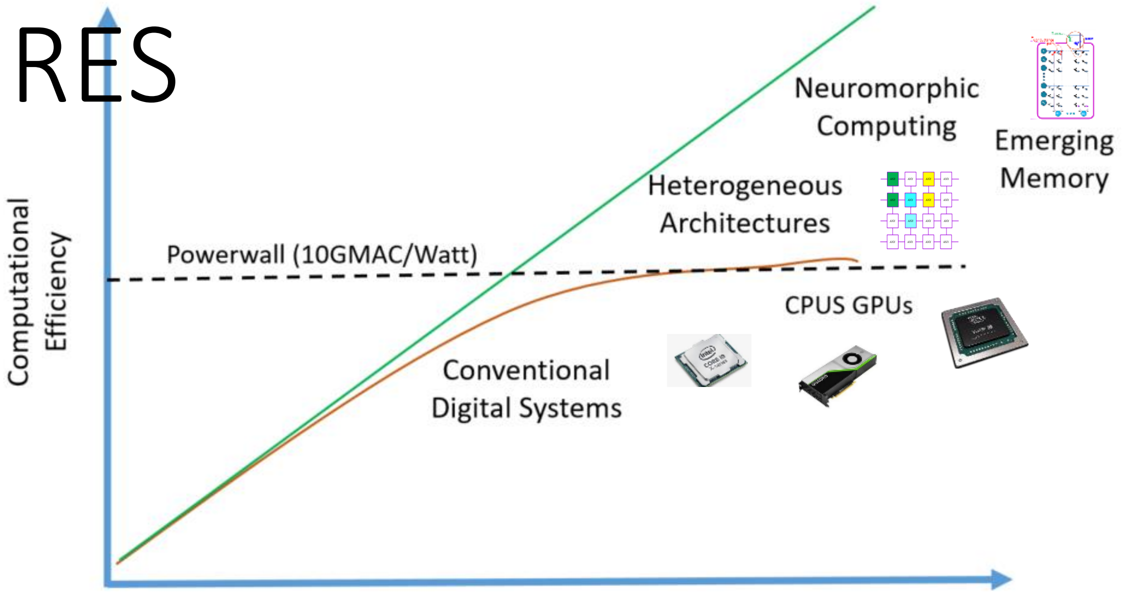

Co-Design is critical to build the next-generation heterogeneous systems

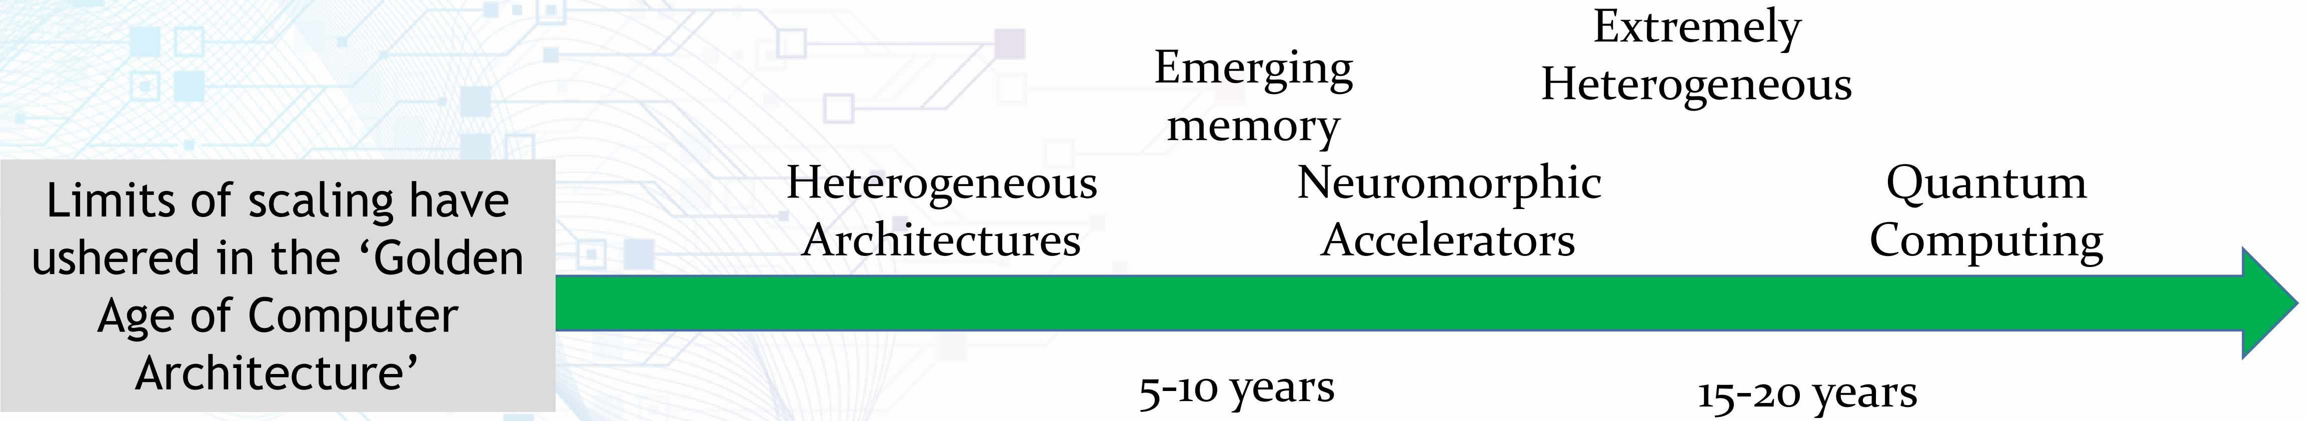

['Truly Heterogeneous Computing', Cardwell et al., SMC 2020]

NEUROMORPHIC COMPUTING: INSPIRED BY THE BRAIN

Brain and Computing: Why make the connection ?

- High computational efficiency, Single neuron ~1MMAC/pW
- Processing and memory operations performed by the same components
- Self-organizing system
- Online learning
- Solving ill-structured problems
- Transfer learning
- Spiking/event driven communication, subthreshold computation

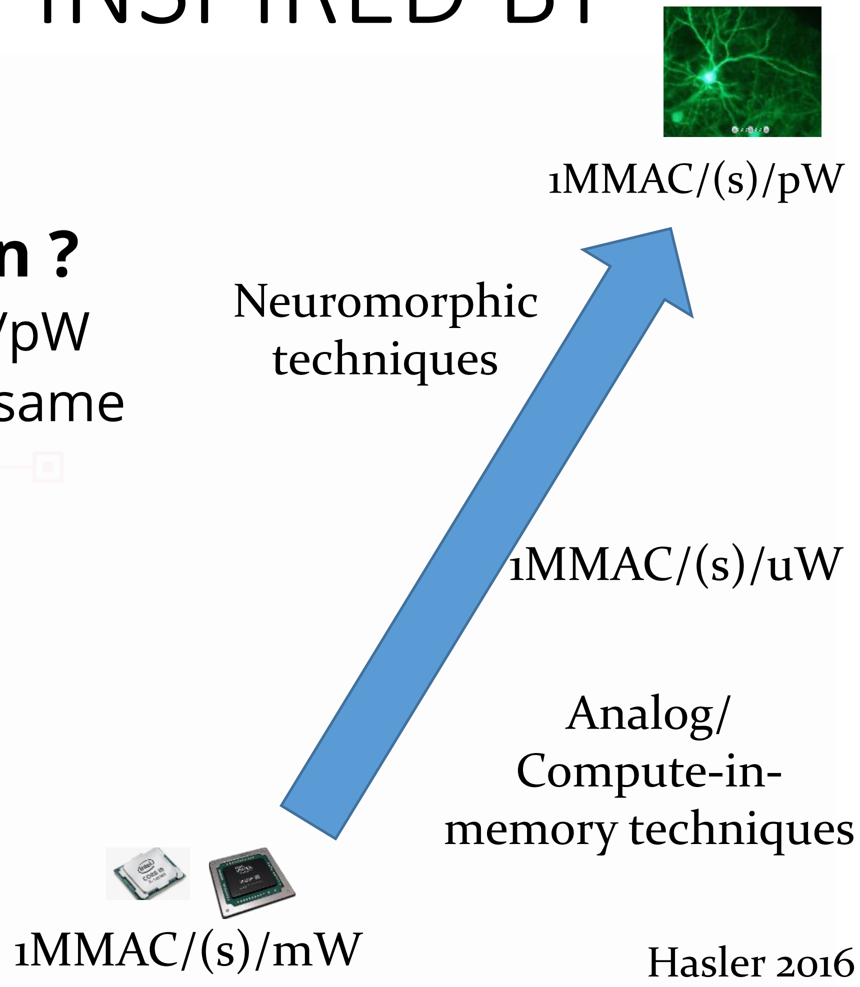

Neuromorphic techniques will be disruptive to how we develop our computing systems

MMAC: Million Multiply Accumulates

NEUROMORPHIC COMPUTING

Digital Neuromorphic

Intel Loihi/ Loihi 2.0

SNL hosts Intel's 50 million neural supercomputer

→ Scaled to a billion neurons

SpiNNaker/ SpiNNaker 2

IBM TrueNorth

ODIN (Open-source)

Analog/Mixed-Signal

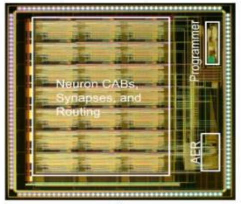

GT Neuron

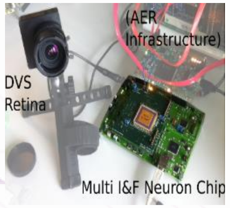

INI, ETH Zurich

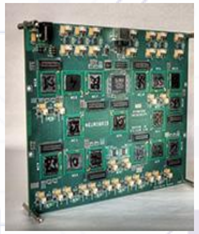

Stanford Neurogrid

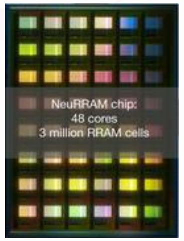

NeuRRAM UCSD/Tsinghua

Beyond CMOS Devices

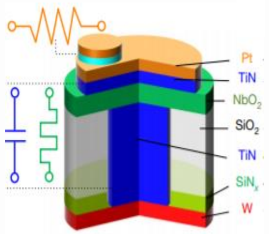

Mott- Memristor

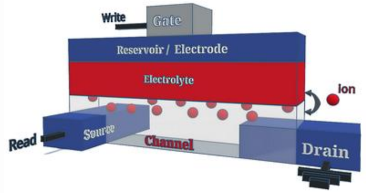

ECRAM

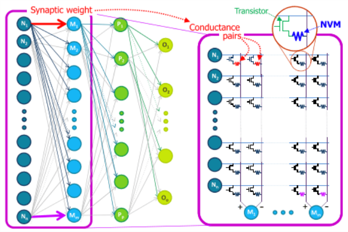

RRAM Crossbar

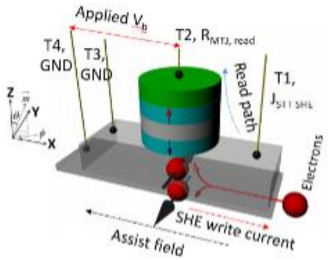

MTJ

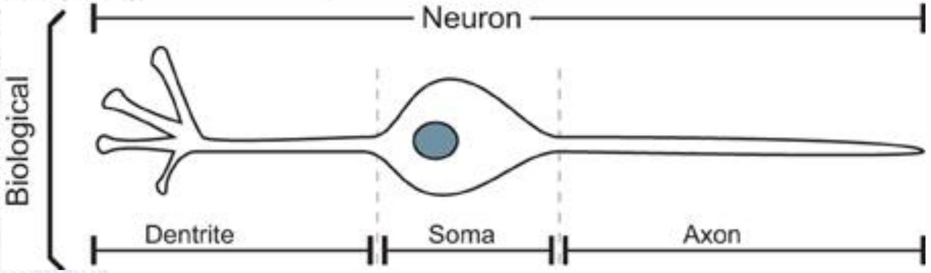

Brink et al., 2013

NEUROMORPHIC BUILDING BLOCKS

Analog Crossbars using NVMs

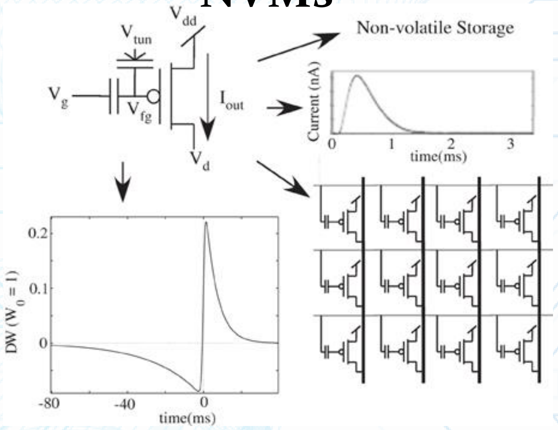

Neuromorphic offers computational richness we can leverage, to move beyond today's computational limitations.

Winner-Take-All

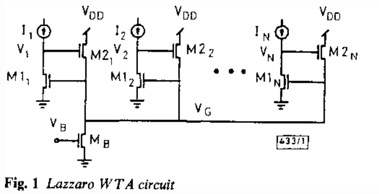

Lazzaro et al. 1988

Dendritic Processing

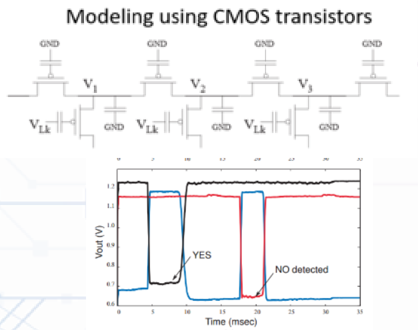

George Cardwell et al. 2013

Neural Path Planning

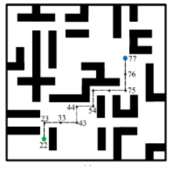

Koziol et al. 2013

Random Walks

Smith et al. 2021

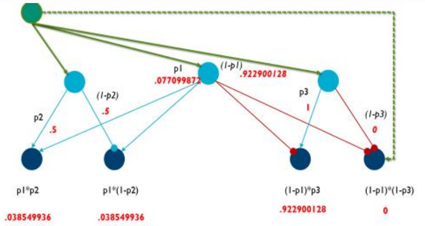

Learning Synapses

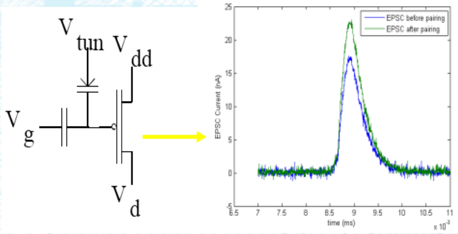

Many different models for neurons, synapses, online learning and dendrites.

Silicon Cochlea

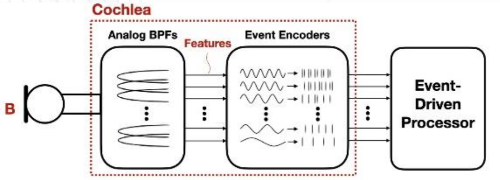

Liu et al. 2020

Silicon Retina/ Event Sensor

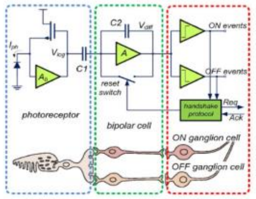

Posch et al. 2014

Delbruck et al. 2020

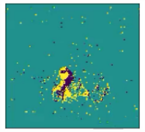

APPLICATIONS

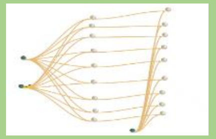

AI/ML (ANN, SNN)

Brain-inspired algorithms

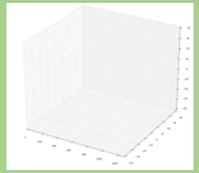

Scientific Computing

Edge Computing

8

NEUROMORPHIC EDGE COMPUTING

ALGORITHM

Supervised Machine Learning Approach- SLAYER SpikeMS for Background Subtraction

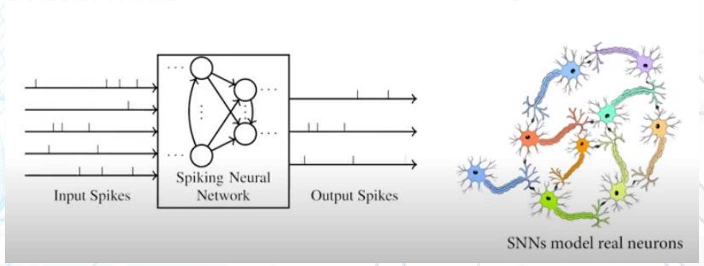

FalconNeuro on ISS, McHarg et al. 2022

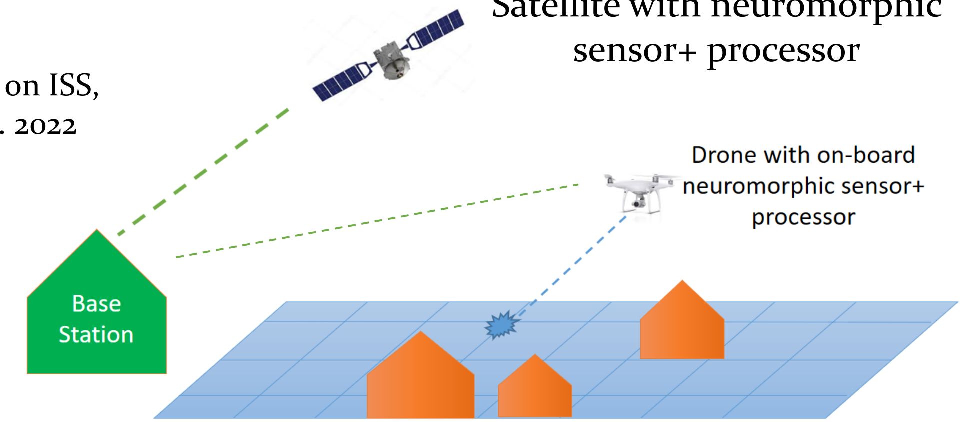

Event Sensor Input → Background Subtraction → Target Classification →

HARDWARE

Neuromorphic Sensor Node (Edge Classification)

Prophesee Gen4

DAVIS 240C Event-based Sensor

Intel's Loihi Neuromorphic Processor

Low power, SWaP constrained

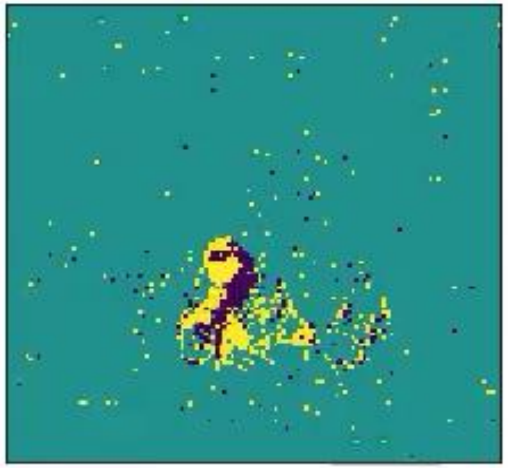

IBM Gesture Dataset

Applications: Space and Remote Sensing, Robotics

External Partnerships: Intel (Intel Neuromorphic Research Community), Event-Based sensing group (AFRL, MITRE, Space Force, AFA)

9

SpikeMS: Prophesee Driving Data

Input

Prophesee Event Sensor Output

Prediction

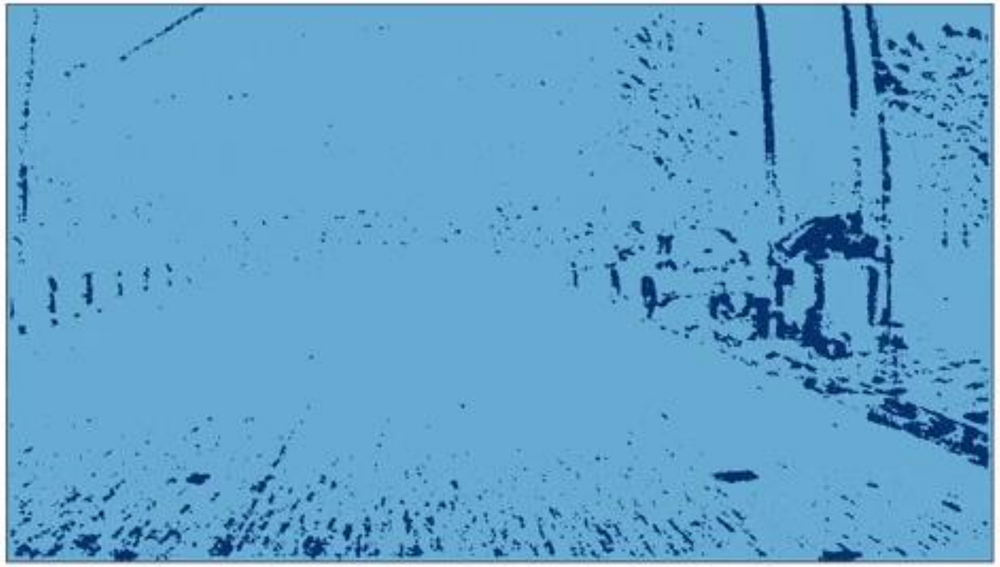

Post-SpikeMS removing most background information

SpikeMS: Parameshwara et al. , IROS 2021

DIRECTIONALLY SELECTIVE

Loihi

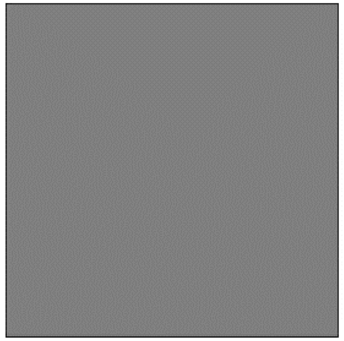

Led Motion Down

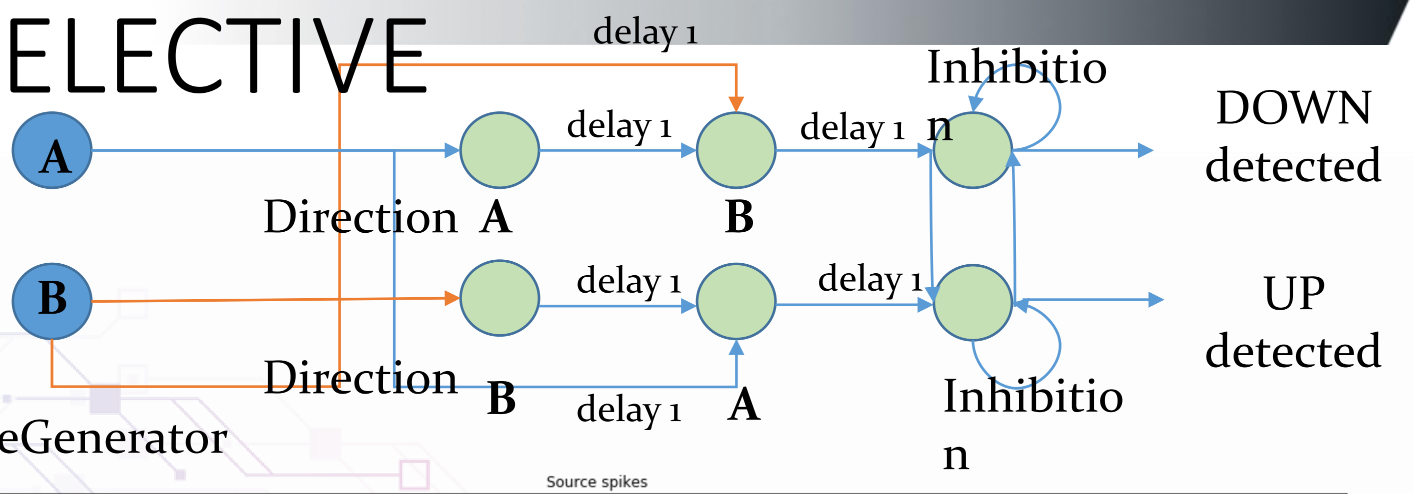

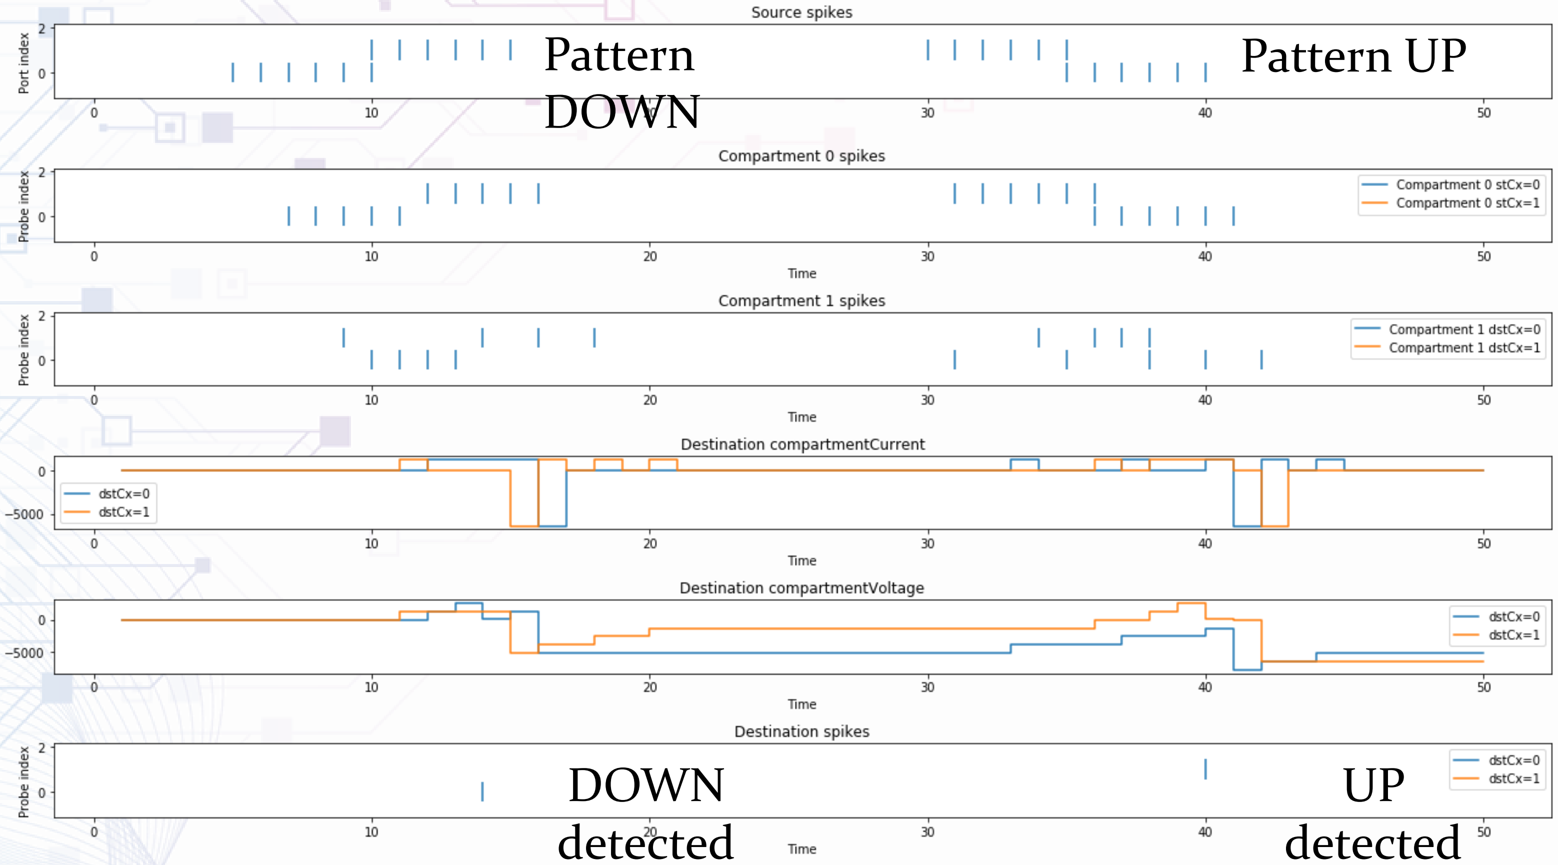

Cardwell et al. 2023 (In review)

NEUROMORPHIC COMPUTING CHALLENGE: SCALABILITY VS. COMPLEXITY

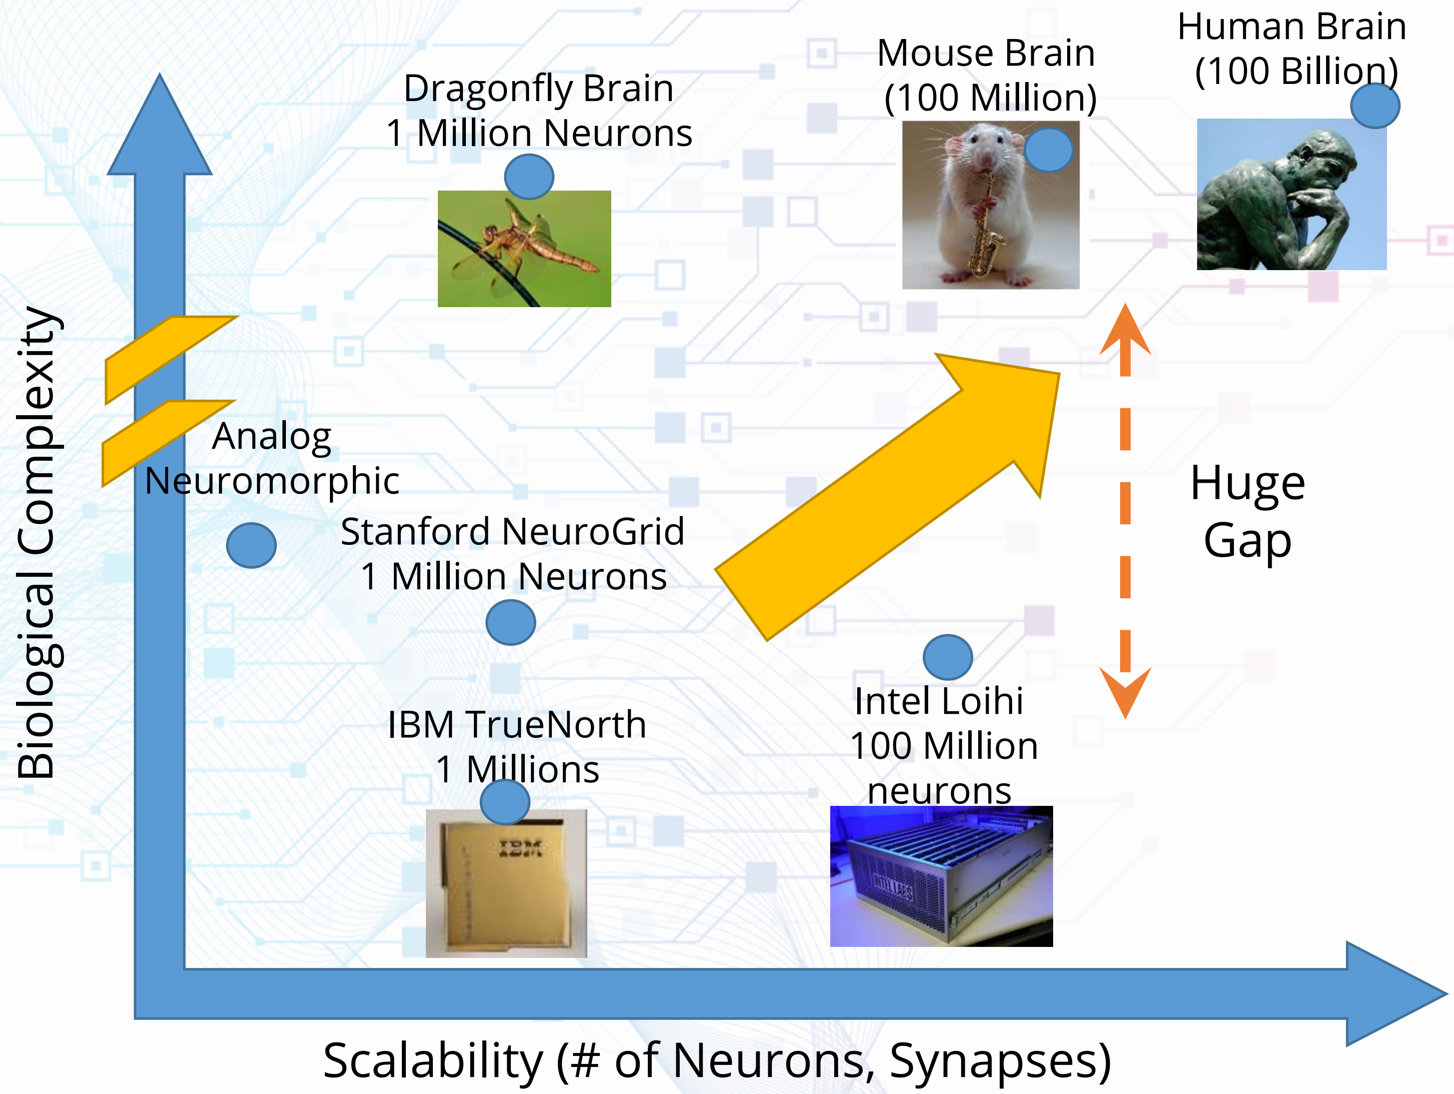

- However, to achieve brain-like complexity we need both scaling and rich dynamics.
  - Solving ill-structured problems
  - Online learning
  - Transfer learning

Understanding fundamental mechanisms in neuroscience, translated to algorithms and models will influence next-generation devices, architectures and intelligent computing systems

INCREASING "BIOLOGICAL COMPLEXITY"

Increase computational efficiency
and
Increase computational density

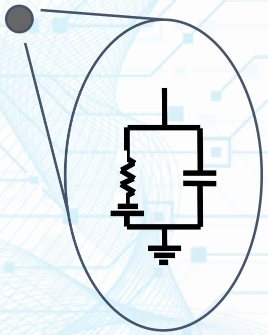

Novel devices and materials can help bridge this gap.

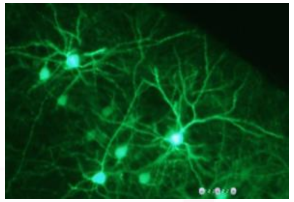

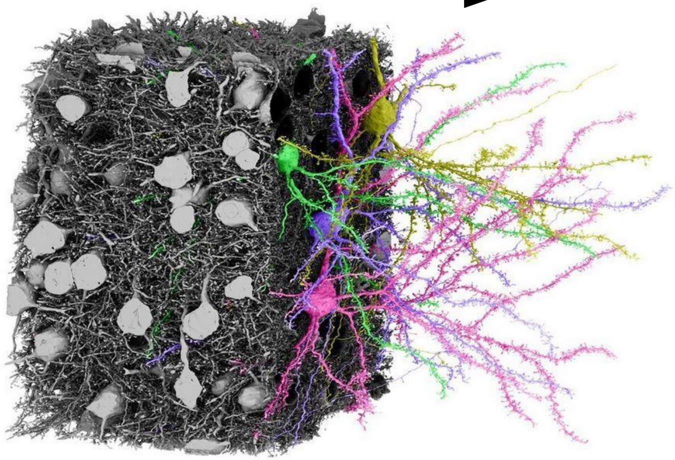

LIF neuron
- Single passive compartment
- Spikes
- Limited dynamics
- Relatively easy to scale

Biological neuron
- Dendrites = intricate structure and dense connectivity
- Complex pattern of active conductances
- Rich dynamics , multiple patterns of spiking, subthreshold computation
- More computational power, not compact

LIF: leaky Integrate and Fire

MODELING DENDRITES

- Dendrites are tree-like structures that connect neurons synapses to its soma.
- **Dendrites are not wires!**
- They can perform interesting computation like
  - Coincidence Detection
  - Current Summation
  - Directional selectivity
  - Non-linear filtering
  - Amplification of Synaptic inputs

London 2005, Hausser 2003

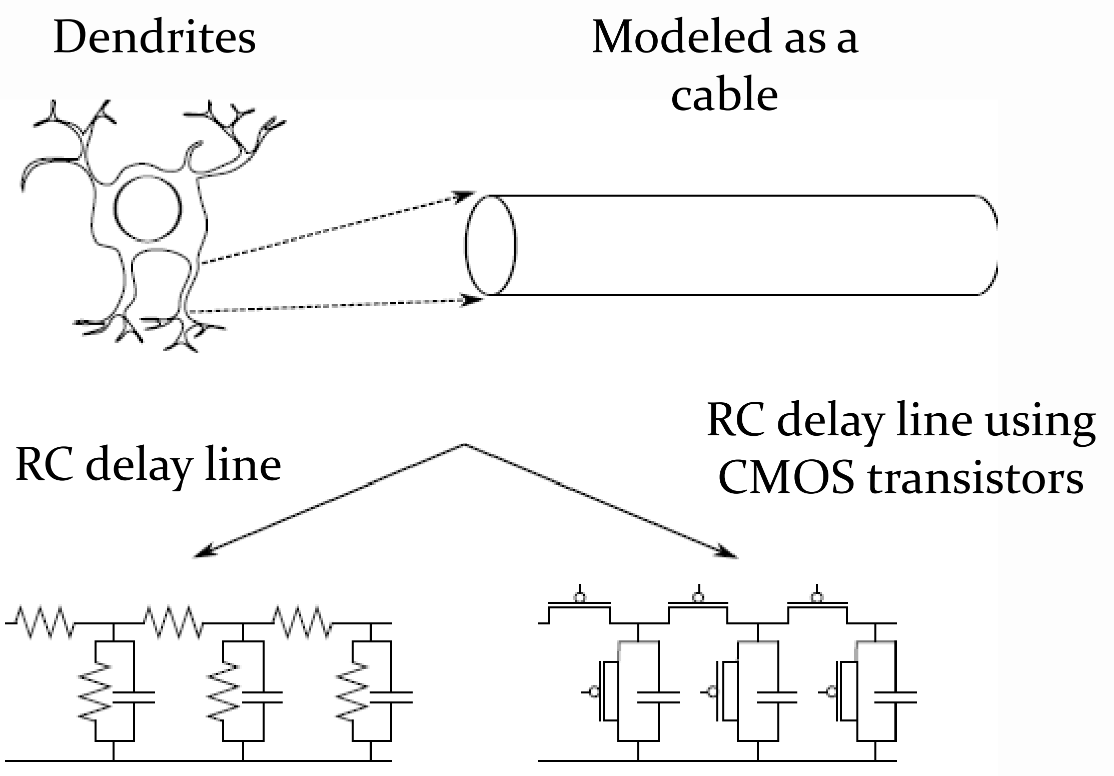

Nease et al. 2011

WHY IS THIS USEFUL?

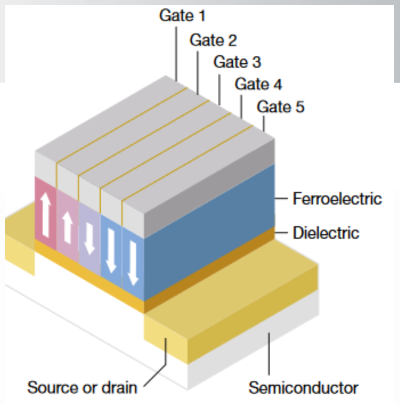

NanoDendrite
Multi-gate FeFET

- Dendrites perform non-linear computation like a multi-layer NN enabling "neuron within a neuron" capability.
- Footprint of dendritic circuits smaller compared to a neuron circuit.
  - Dendrite (1-2 transistors with sub-threshold FGs) [George 2013], HH Neuron circuit (7-8 transistors)[Farquhar and Hasler], LIF (8-10 transistors) [Indiveri 2011]
  - Proposed multi-gate FeFET for dendrites with 3D stacking for dense connectivity [Kwabena, Nature 2022], Energy estimated to program device ~29.6fJ/event [Saha et al., 2021]
- Computing in the interconnect, more energy efficient computation

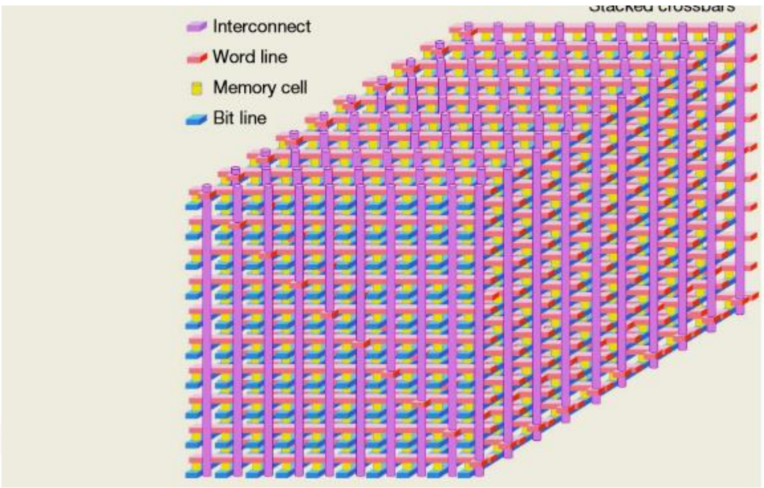

[Dendrocentric Learning- Kwabena, Nature 2022]

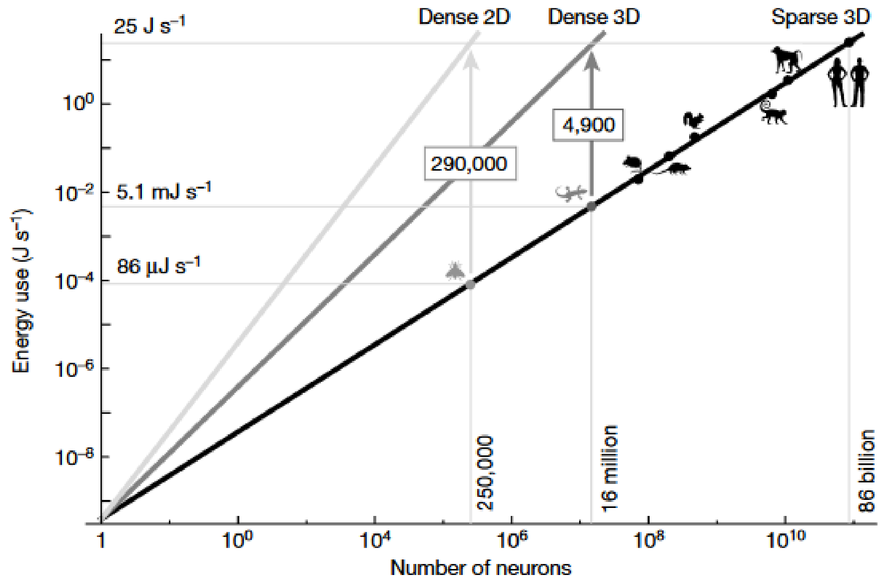

14

DRAGONFLY EXAMPLE

Algorithms *Devices & Circuits* *Physics of Computing*

Coordinate transformations from Dragonflies to Neuromorphic Hardware

Lead PI: Frances Chance, SNL

October 2021

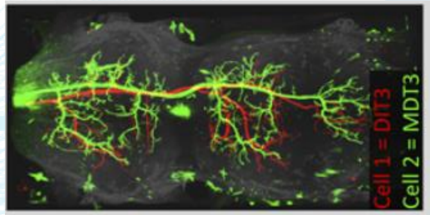

Gonzalez-Bellido, UMN

DRAGONFLY EXPERIMENTS

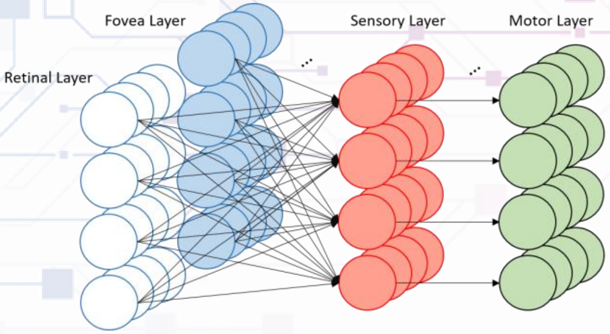

Chance 2020

COMPUTATIONAL MODEL

August 2021

GT FPAA

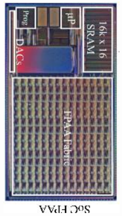

George Cardwell 2016

Intel's Loihi

Davies 2018

SNL, Baylor

NEUROMORPHIC IMPLEMENTATION

Increased collaboration between neuroscience and neuromorphic engineering will facilitate development of novel neural-inspired architectures.

DOE ASCR (FY21-24)
Department of Energy
Advanced Scientific Computing Research

Multiplication in a single neuron

Algorithms | *Devices & Circuits* | *Physics of Computing*

from fan-shaped body of *Drosophila* brain

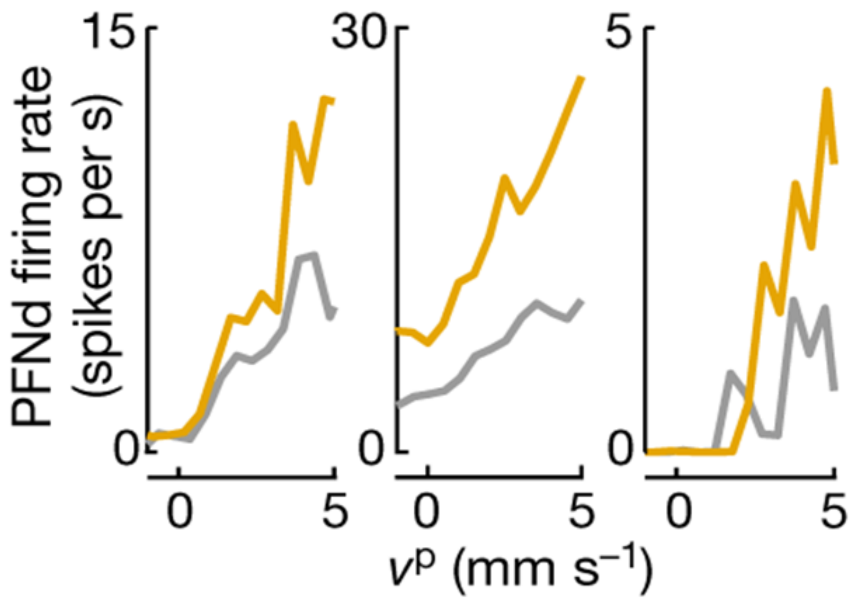

Lu et al. 2020

Chance & Cardwell, NICE 2022

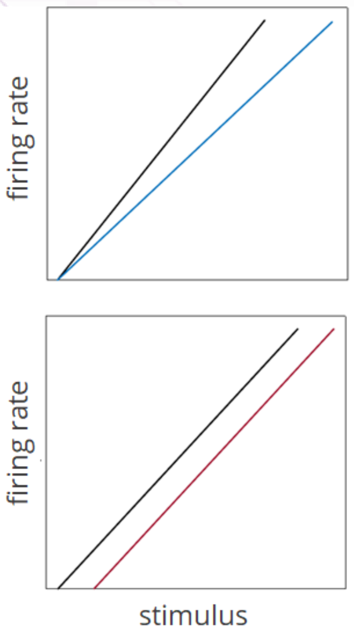

R = **A***f(x)

R = f(x) **- A**

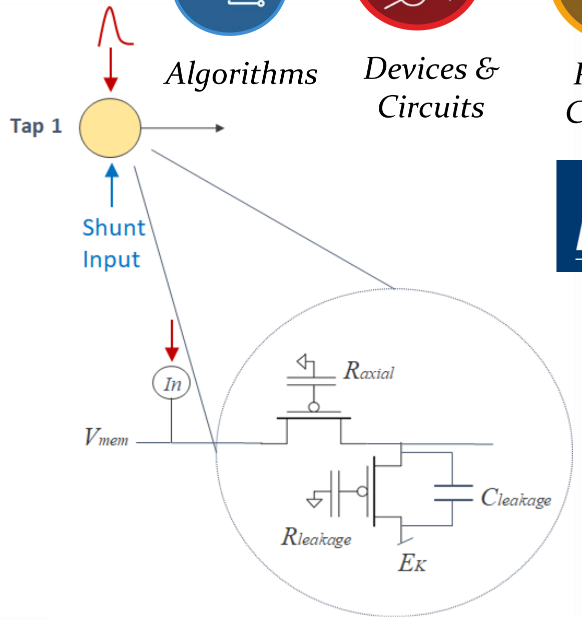

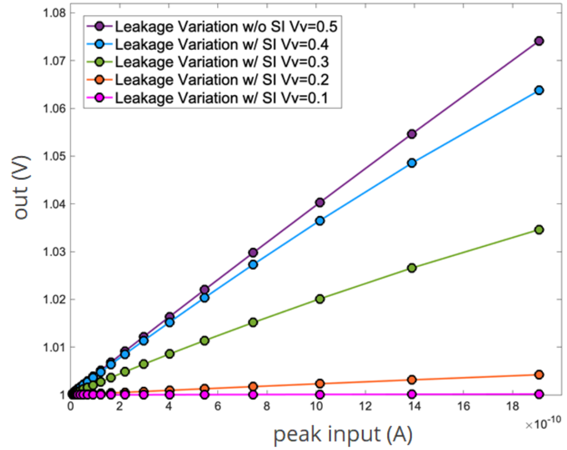

Shunting Inhibition in Neuromorphic Dendrite

Collaborators: University of Texas at Austin, Intel Neuromorphic Research Community

NEXT-GENERATION NEUROMORPHIC SYSTEMS

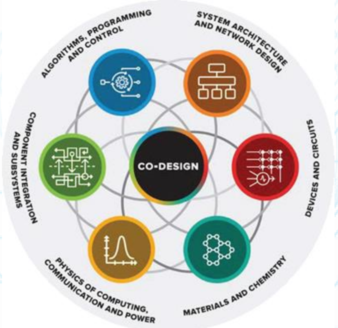

BRN Report
2018

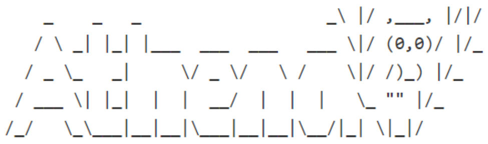

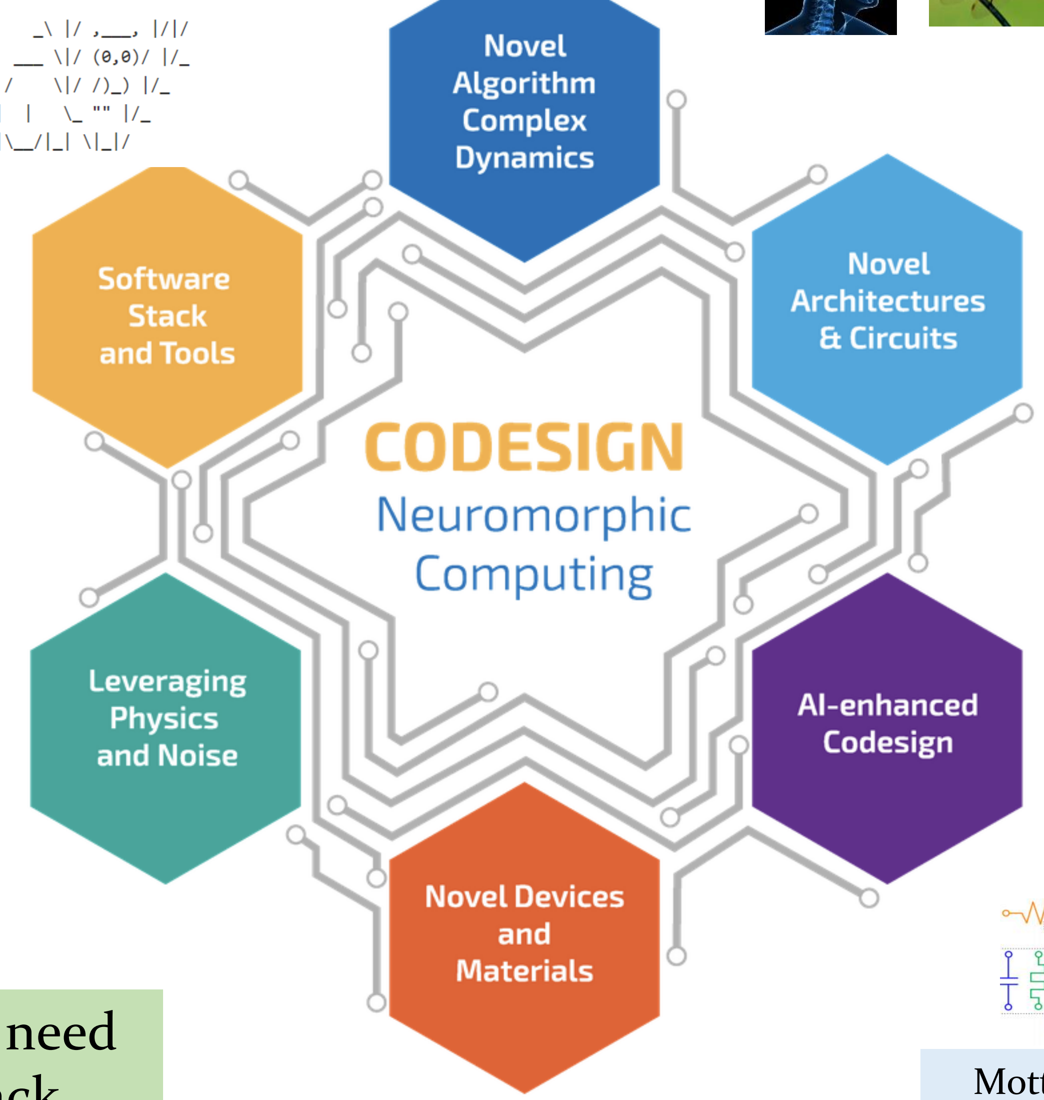

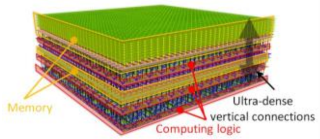

Stanford n3XT Platform
(Aly et al., 2018)

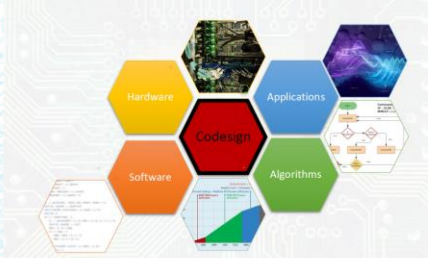

ASCR Workshop on Reimagining Codesign
2021

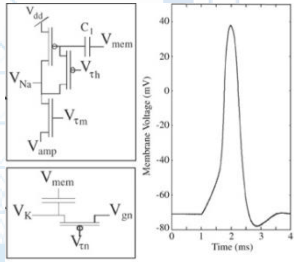

Evolutionary/
Reinforcement Learning approaches

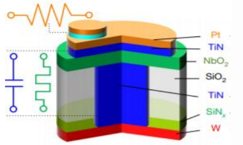

Mott- Memristor
Kumar et al., 2020

In order to innovate we need cross-talk across the stack.

CODESIGN IS CHALLENGING

Co-Design Tools for Novel Architectures

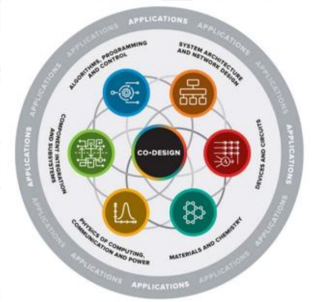

ATHENA

Analytical Tools for analog and neuromorphic ML accelerators

ASC-AML (FY20-22)
Advanced Simulation & Computing - Advanced Machine Learning

AI-enhanced Codesign

Reinforcement Learning/Evolutionary methods for Circuit and System design

AS&T LDRDs (FY21-23)
Advanced Science & Technology Laboratory Directed Research and Development

Next-generation Neuromorphic Architectures

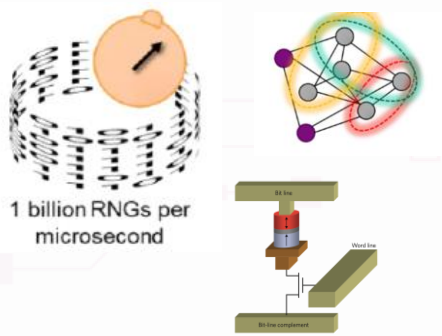

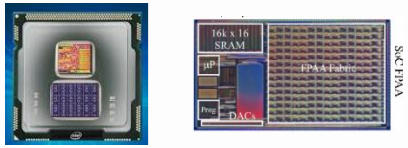

COINFLIPS

Probabilistic Neural Computing, Leverage stochasticity in beyond-CMOS devices

DOE ASCR/BES (FY21-24)
Department of Energy Advanced Scientific Computing Research Basic Energy Sciences

DRAGONFLY

Dendritic processing, Coordinate transformation from Dragonflies to Neuromorphic hardware, Analog and digital

DOE ASCR (FY21-24)
Department of Energy Advanced Scientific Computing Research

External Collaborators: UT Austin, Intel, Infineon Memory Solutions, Georgia Tech, UMN, Baylor University, UT Knoxville, Temple University, NYU, ORNL

AI-ENHANCED CODESIGN

Data and Models

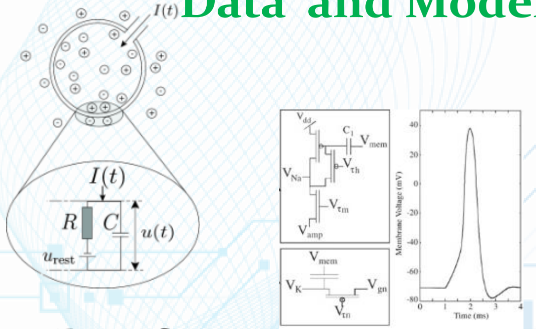

- Data Sweeps
- Device Models
- ASIC behavior models

Neural Circuits & Architectures

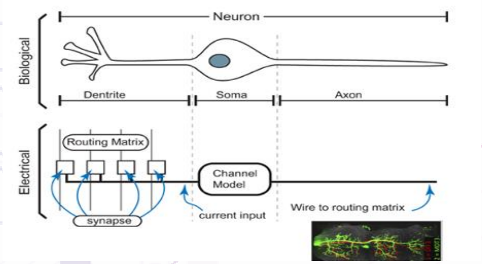

Hyper Parameters

- Learning Rate
- # of Epochs
- Hardware based constraints in architecture search

Works based on Reinforcement learning (Goldie & Mirhoseini, 2020) and Game theory augmented techniques (Cooksey & Mavris 2011, Ganesan et al., 2015)

Machine Learning

Topological Analysis

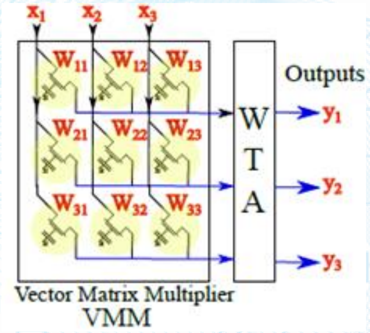

- Size constraints
- Discover novel circuit topologies

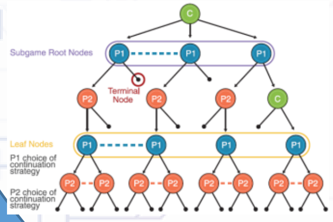

Game Theory

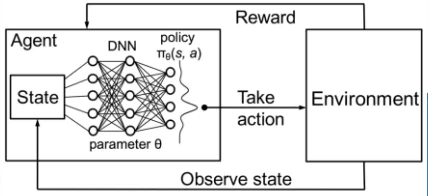

Reinforcement Learning

Device and Architectural Constraints

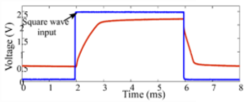

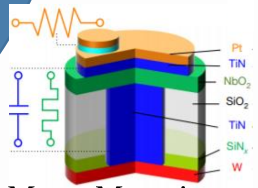

Mott- Memristor
Kumar et al., 2020

- Charge time
- Energy efficiency
- SWaP
- Connectivity
- Extreme Temperature environments

Our AI-enhanced framework would need inputs from algorithms, devices, architectures and ML-based hyper-parameters. The framework will enable new capabilities.

AI-ENHANCED CODESIGN: NEURAL CIRCUITS

We developed an RL algorithm approach which is capable of building very simple circuits.

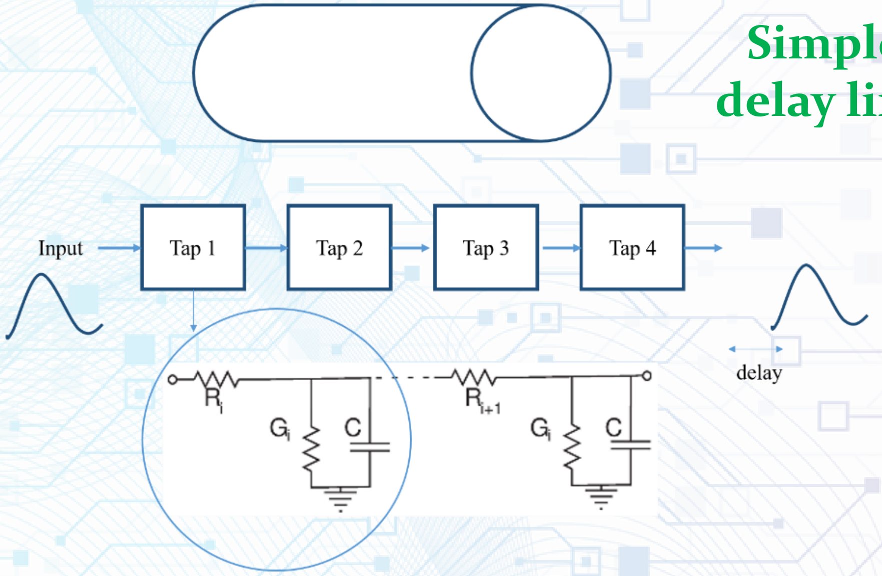

Simple delay line

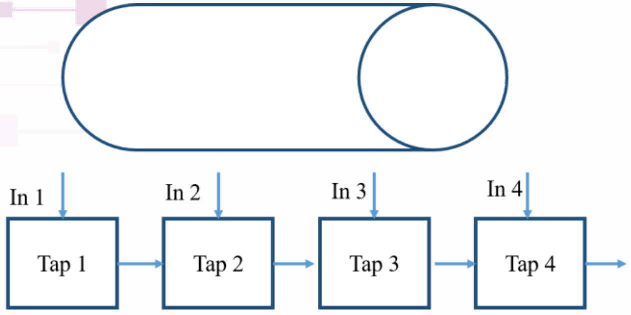

Pattern Detection

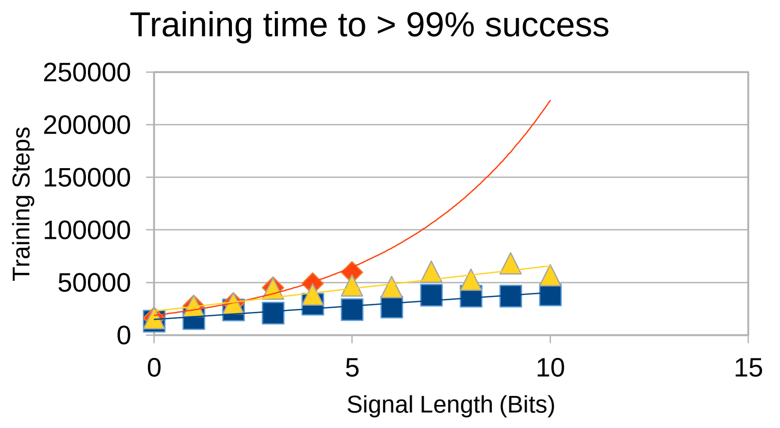

Crowder et al. 2023 (In review)

Algorithms

Physics of Computing

Devices & Circuits

AI-ENHANCED CODESIGN: COINFLIPS

Lead PI: Brad Aimone

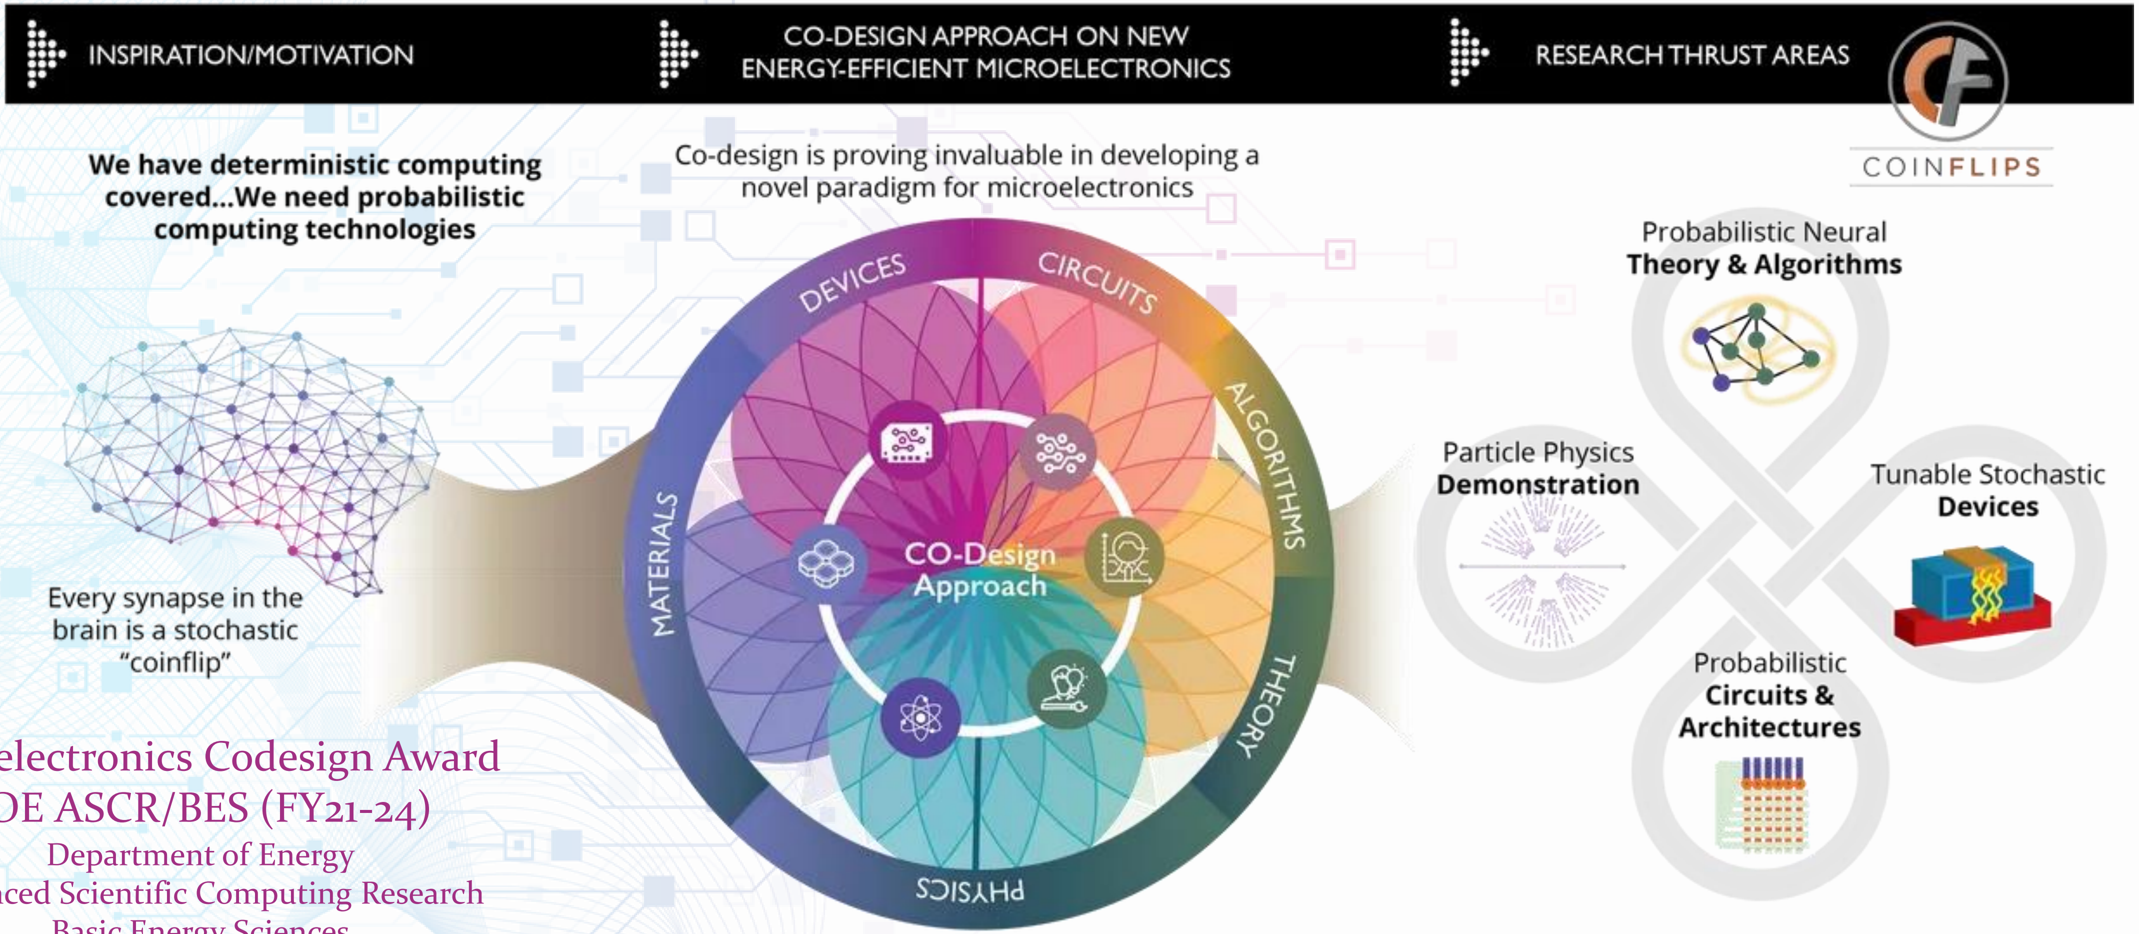

Microelectronics Codesign Award
DOE ASCR/BES (FY21-24)
Department of Energy
Advanced Scientific Computing Research
Basic Energy Sciences

Collaborators: NYU, ORNL, Temple University, UT-Austin and UT-Knoxville

https://coinflipscomputing.org/

AI-GUIDED CODESIGN OF PROBABILSITIC CIRCUITS

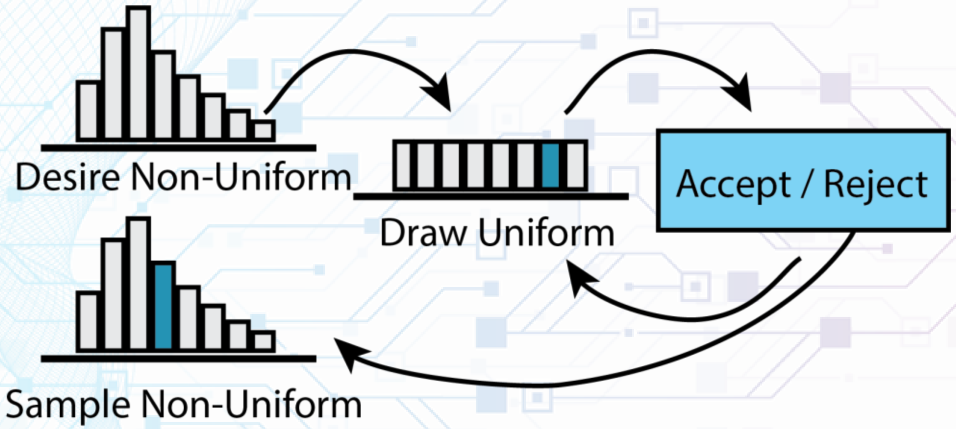

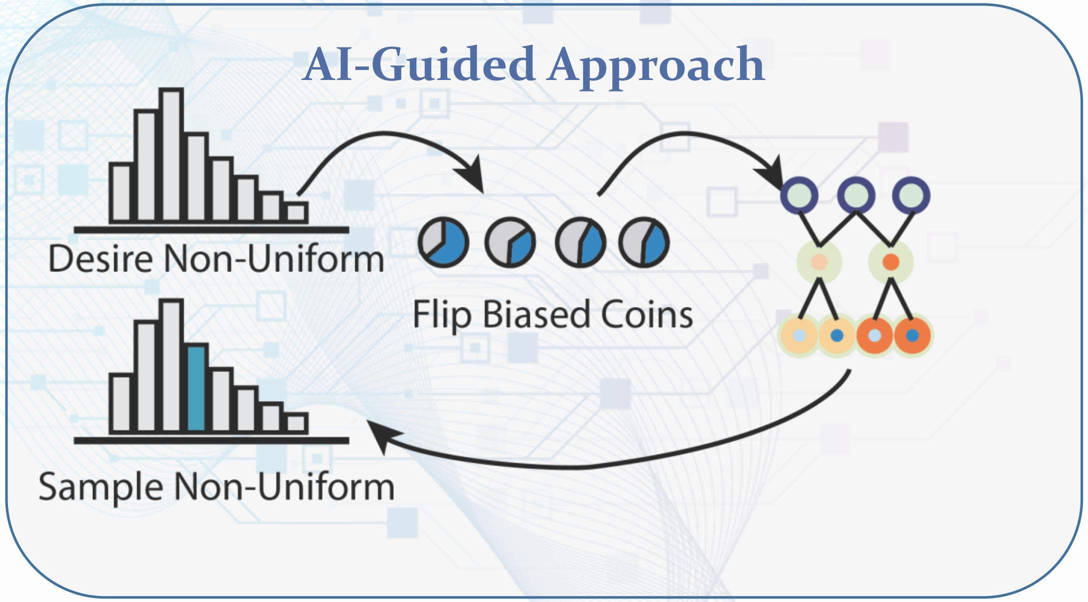

- Unfair coins can be combined with AI-designed neural circuits to allow sampling of application desired probability distributions, avoiding accept/reject steps.
- We leveraged evolutionary algorithms for circuit design and optimization
  - **LEAP** (Library of Evolutionary Algorithms in Python)
  - **EONS** (Evolutionary Optimization for Neuromorphic Systems)- Schuman et al. , 2020
- We used abstracted device models for TD and MTJ to capture functionality and energy usage.

Cardwell et al., International Conference on Rebooting Computing (ICRC) 2022

AI-GUIDED CODESIGN OF PROBABILSITIC CIRCUITS

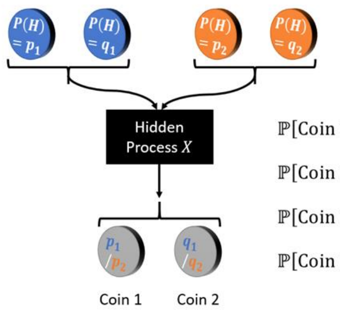

$\mathbb{P}[\text{Coin } 1 = H \text{ and Coin } 2 = H] = \frac{1}{2}$

$\mathbb{P}[\text{Coin } 1 = H \text{ and Coin } 2 = T] = \frac{1}{6}$

$\mathbb{P}[\text{Coin } 1 = T \text{ and Coin } 2 = H] = \frac{1}{6}$

$\mathbb{P}[\text{Coin } 1 = T \text{ and Coin } 2 = H] = \frac{1}{6}$

Probabilistic Mixing Algorithm

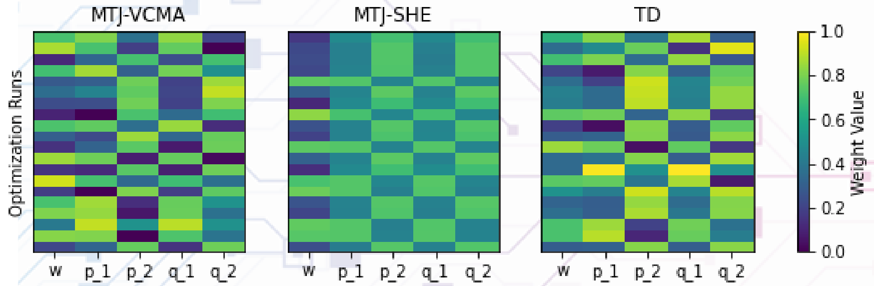

Optimized weight values for each device

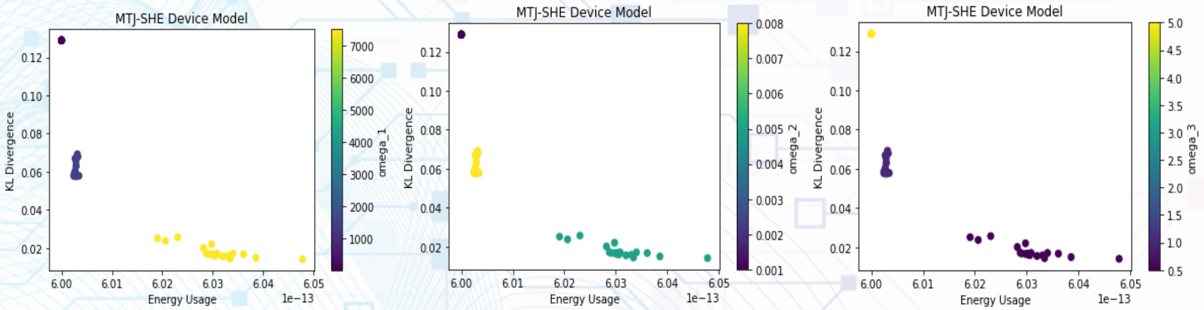

Multi-objective optimization of weights of fitness function for optimal KL divergence, biased weight and energy usage.

Cardwell et al., International Conference on Rebooting Computing (ICRC) 2022

- Weights are customized for the device's behavior to target the best performance in terms of KL divergence and energy usage.
- One of the challenges in optimizing for both algorithms and devices was appropriately abstracting the device models and algorithmic constraints.
- The functional models developed will also be evolved in time as new device data and research emerges.
- Our framework can accommodate any emerging device type.

24

COINFLIPS APPLICATION: NUCLEAR PHYSICS SIMULATIONS

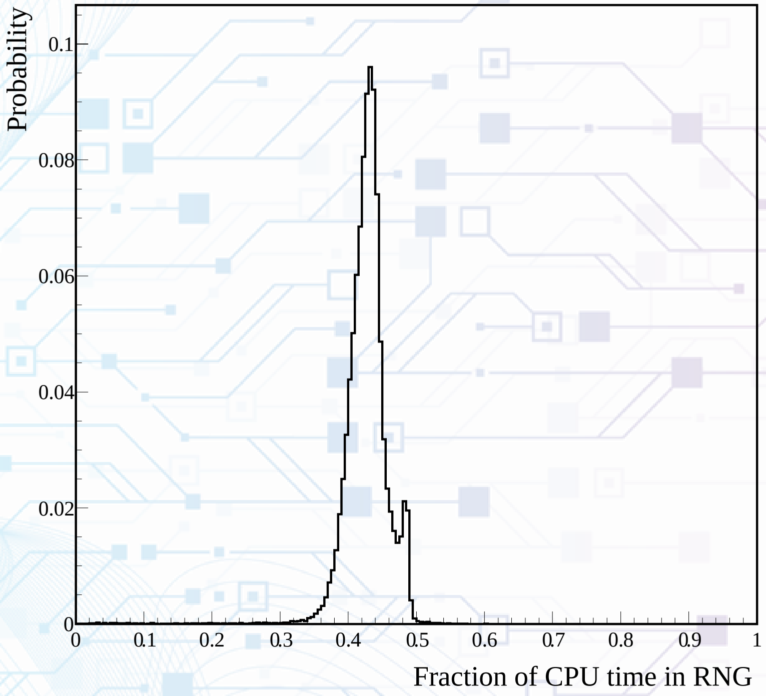

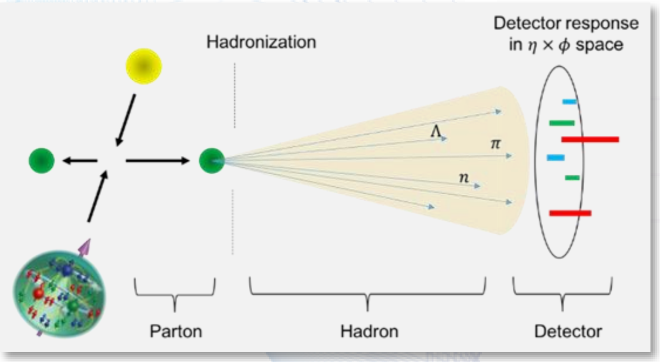

- For a particular collider physics simulation [Pierog et al., Phy Rev. 2022], ~ 270K pseudo- random numbers needed for a single event, with billions of events needing to be simulated.
- CPU time is ~ 30-50% of the total compute time
- Direct random number generation leveraging stochastic devices can promise significant energy savings for such applications

Misra et al., *Advanced Materials 2022*

Random numbers are a limiting computational cost for some nuclear physics applications

25

COINFLIPS APPLICATION: NEUROMORPHIC MAXCUT

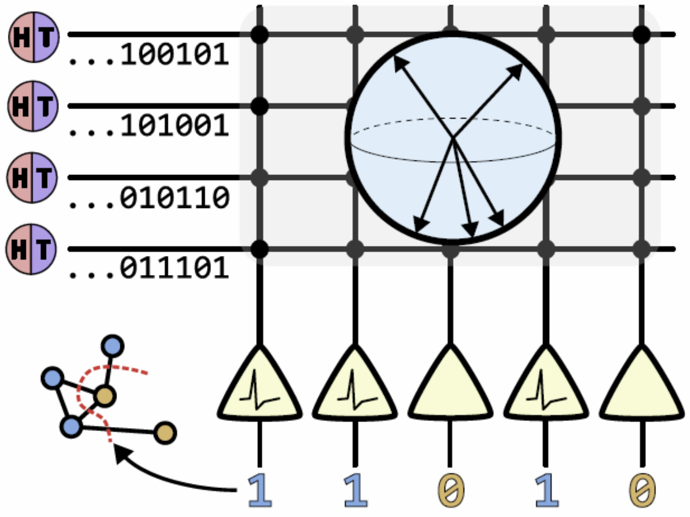

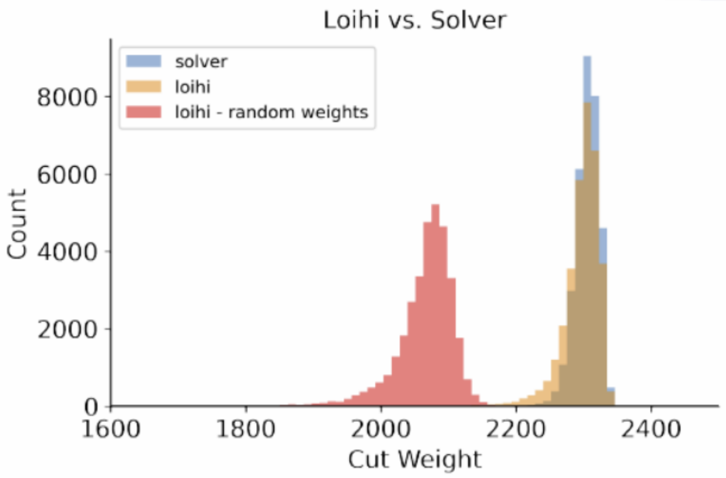

Loihi generated graph cuts of neuromorphic GW algorithm

- Loihi generated graph cuts of neuromorphic GW algorithm match conventional MAXCUT solver generated cuts.
- Results demonstrate effective implementation on current neuromorphic platform with minimal loss and potential to take advantage of future accelerated neuromorphic platforms.
- MAXCUT has broad real-world applications ranging from circuit design to power grid resilience, and these applications are well positioned to take advantage of dramatically accelerated neuromorphic implementations.

Theilman & Aimone, *Neuro-Inspired Computational Elements Conference* (NICE 2023)

AI-ENHANCED CODESIGN ACROSS SCALES

Device Design	Circuit Design	System Design	Architecture Design	Algorithm Design

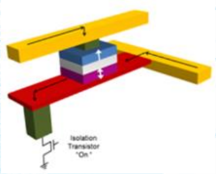

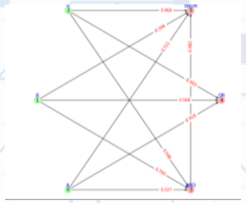

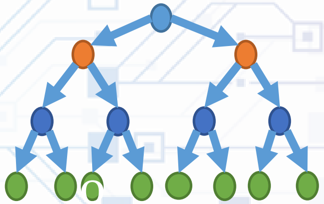

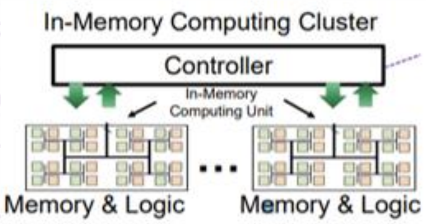

Fan (UCF), 2018

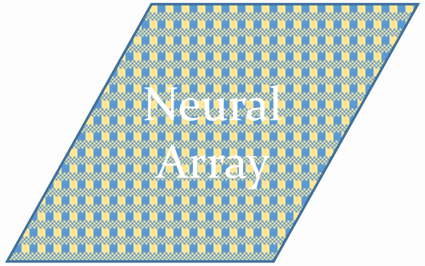

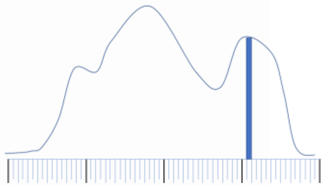

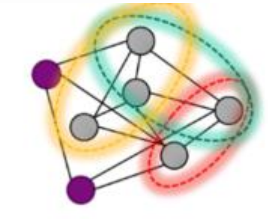

Approach

Can we leverage AI to generate specifications for novel devices?

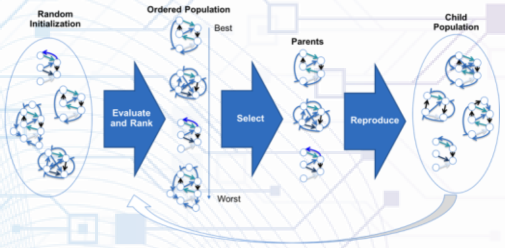

Evolutionary/RL approaches

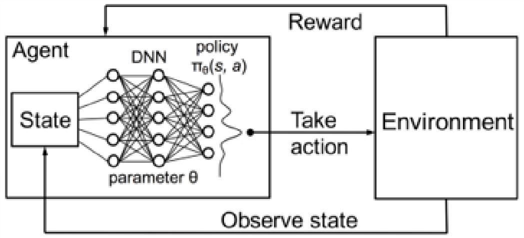

RL approaches

Analytical and cycle-accurate tools, network simulation tools

RL approaches

AI-ENHANCED CODESIGN ACROSS SCALES

Device Design	Circuit Design	System Design	Architecture Design	Algorithm Design

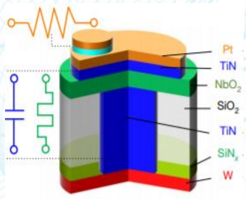

Mott- Memristor
Kumar et al., 2020

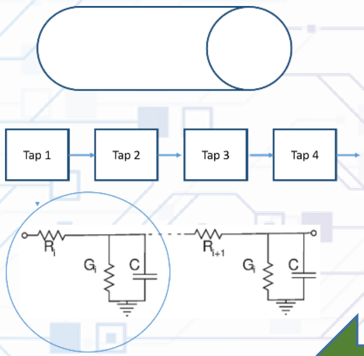

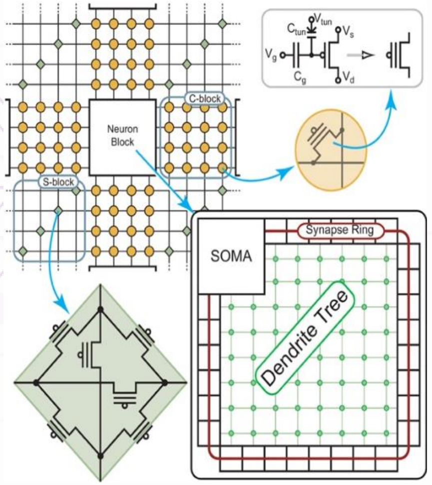

Ramakrishnan, 2013

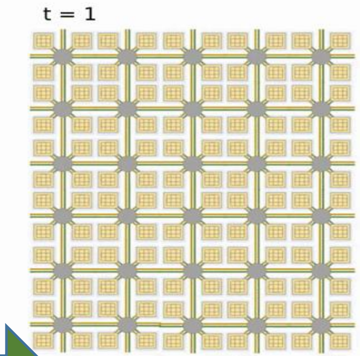

Davies, 2018

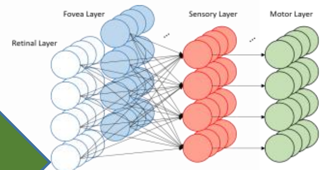

Chance 2020

Approach

Can we leverage AI to generate specifications for novel devices?

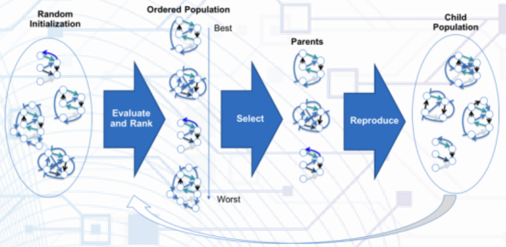

Evolutionary/RL approaches

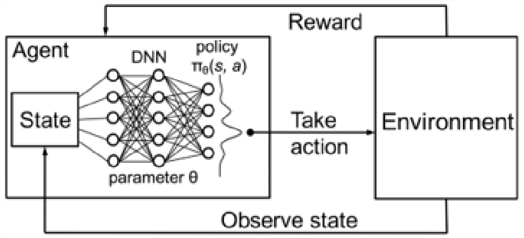

RL approaches

Analytical and cycle-accurate tools, network simulation tools

RL approaches

CHALLENGE: SOFTWARE TOOLS

- ATHENA (Analytical Tool to evaluate Neuromorphic Architectures) will be leveraged to do design space exploration of novel architectures that leverage neuromorphic and emerging devices.

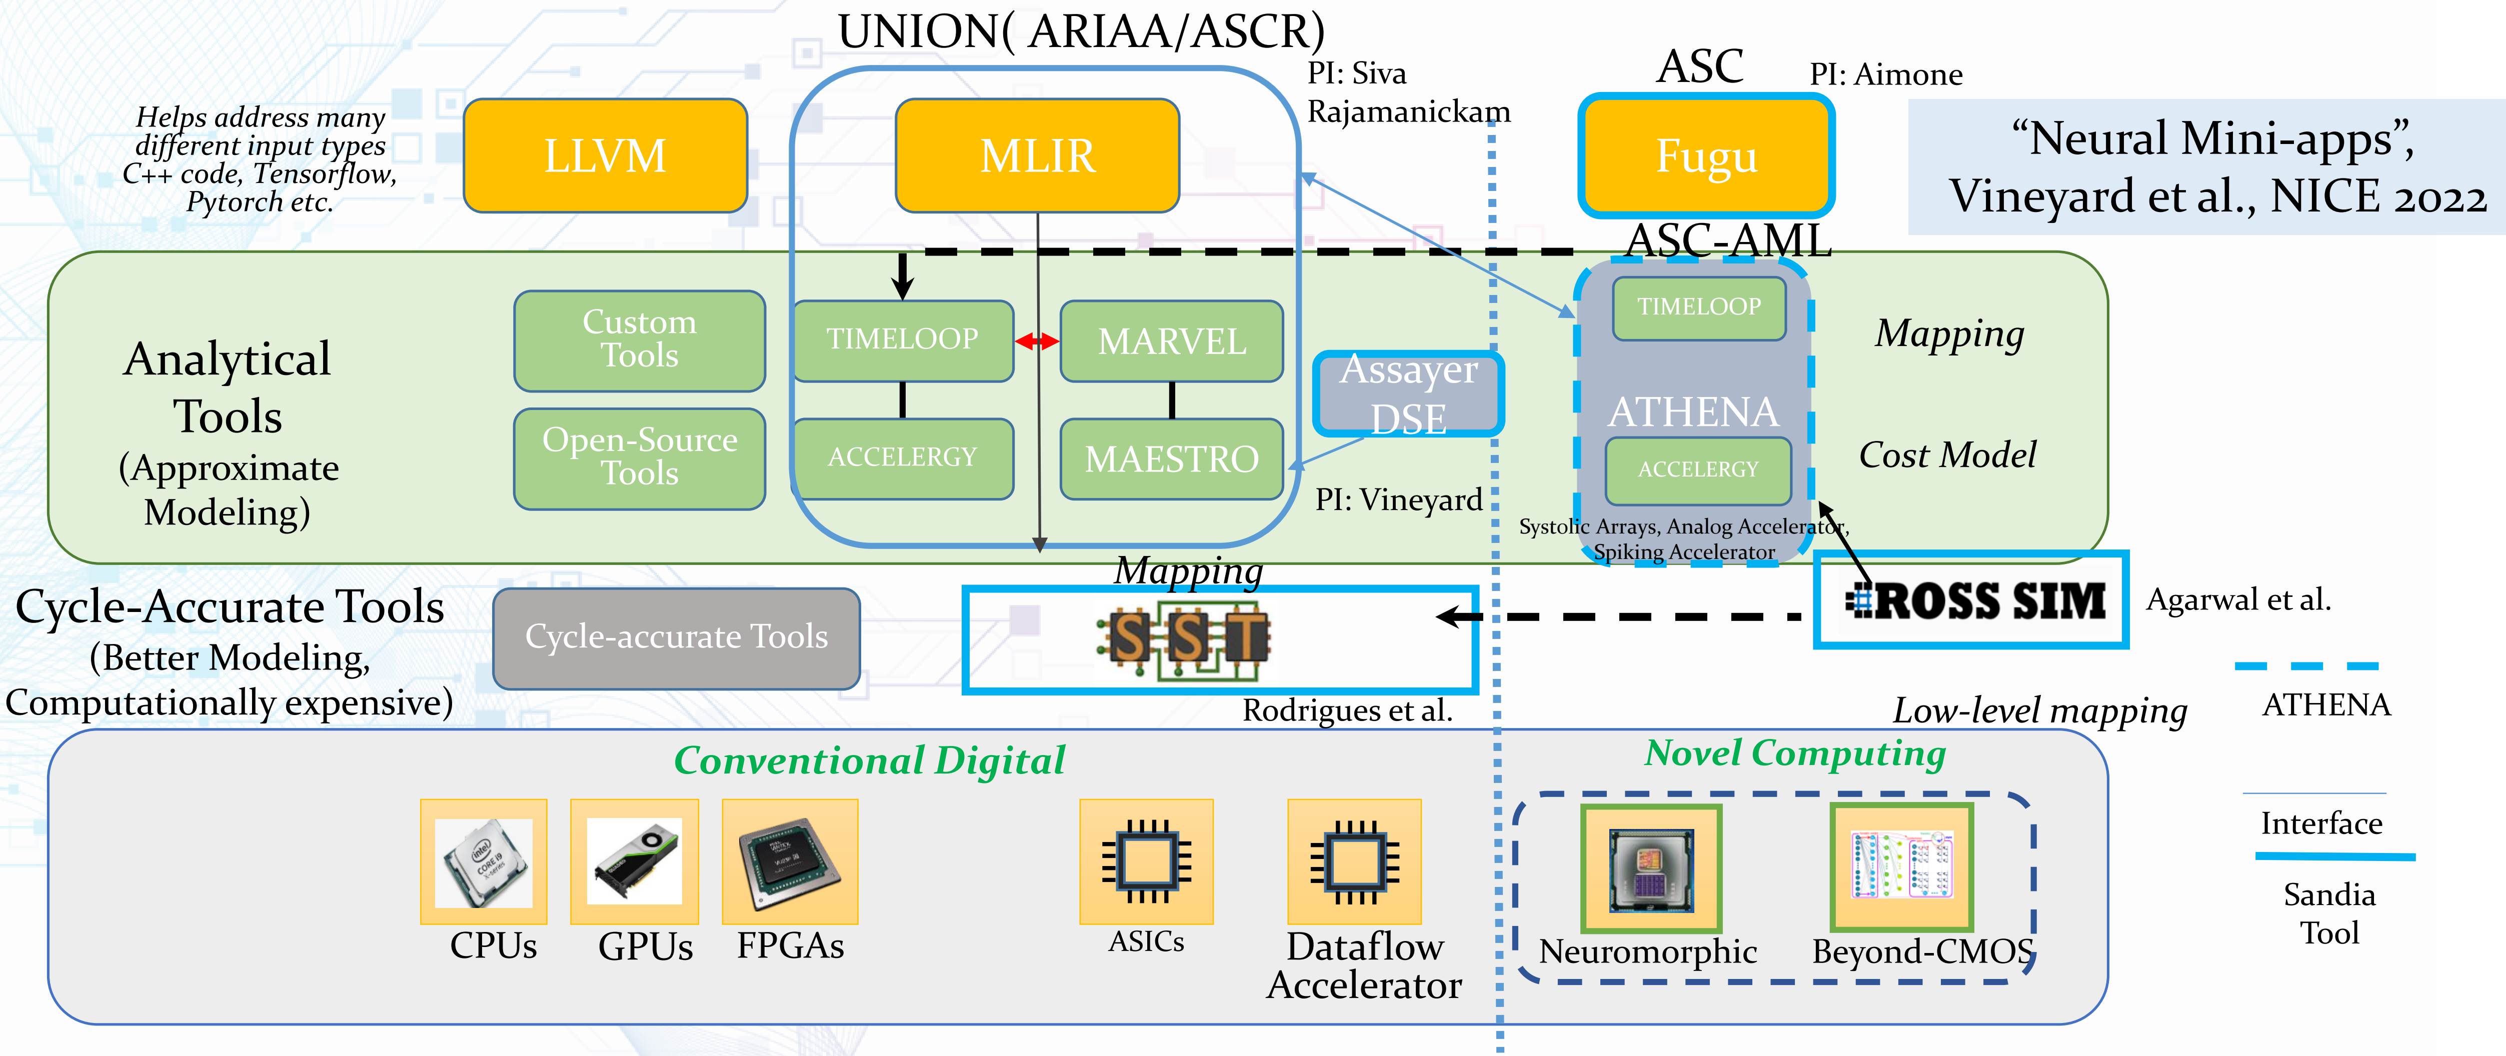

ATHENA

(Analytical Tool to Evaluate Heterogeneous Neuromorphic Architectures)

- ATHENA will quickly evaluate performance metrics of analog architectures

- Developed as part of a larger ecosystem
  - Tools to enable next-generation hardware design prototyping

Plagge et al., International Conference on Rebooting Computing (ICRC) 2022

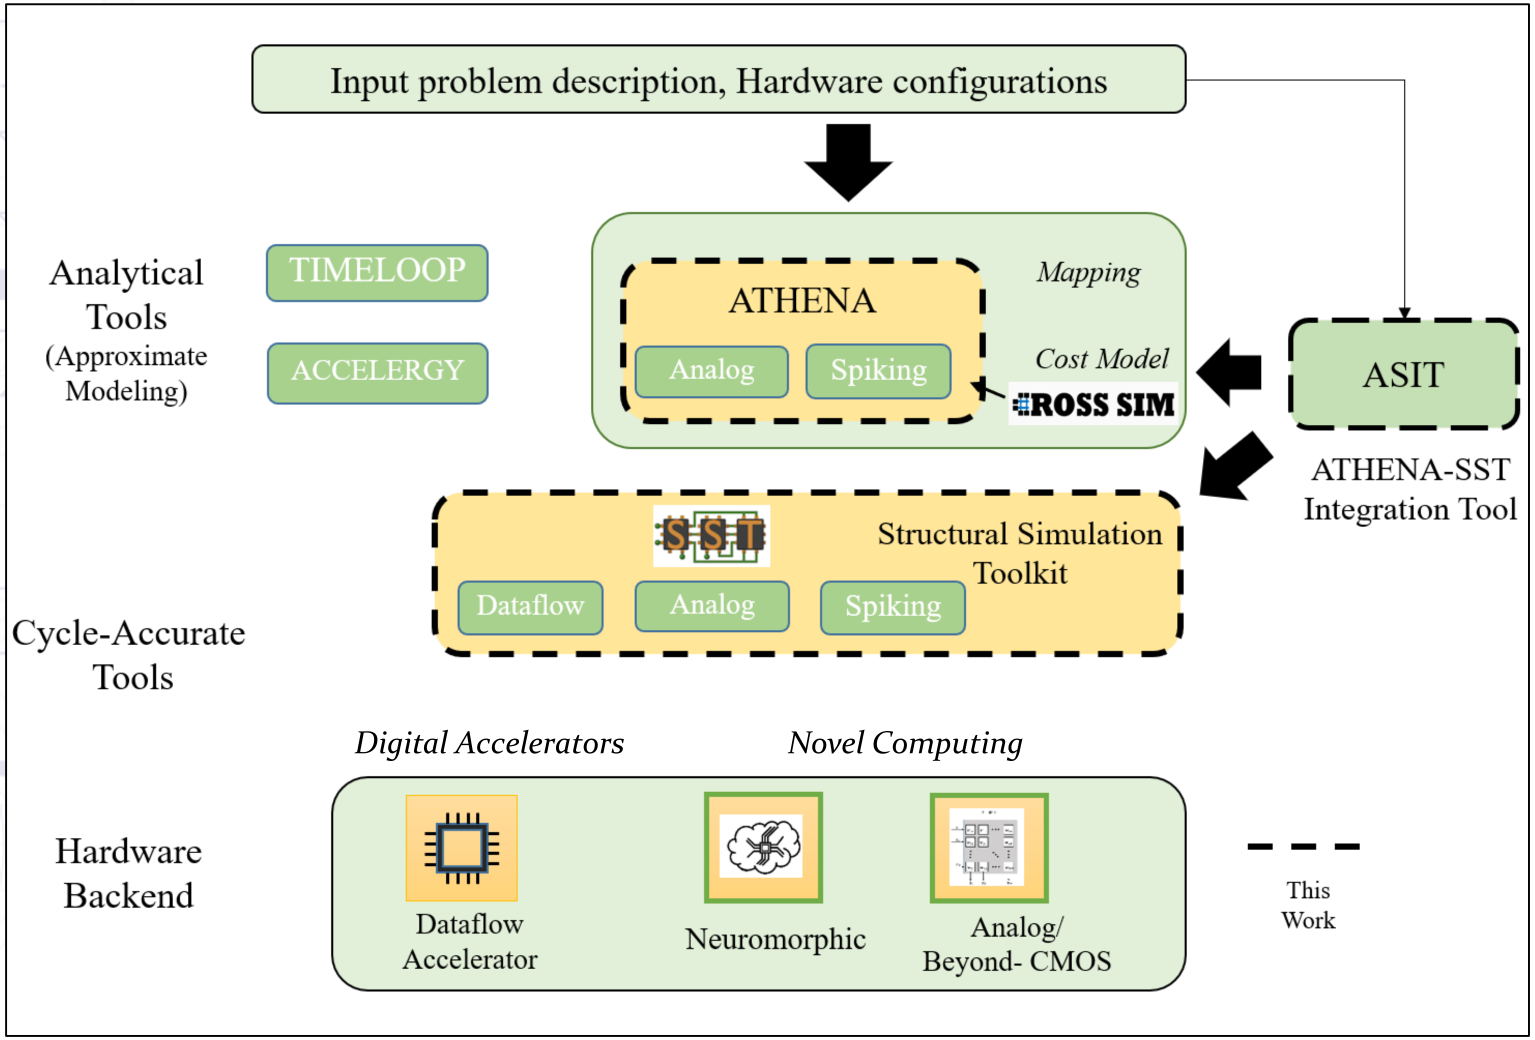

30

ATHENA – HARDWARE PERFORMANCE

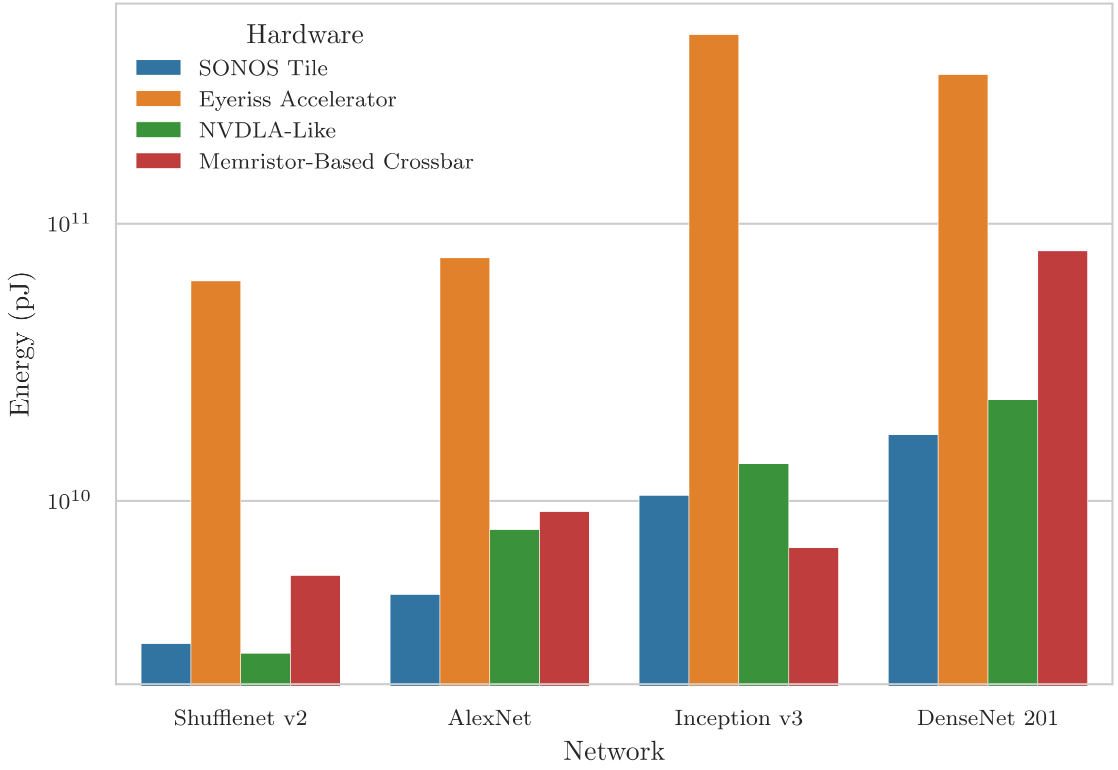

Plagge et al., International Conference on Rebooting Computing (ICRC) 2022

- ATHENA was used to compare the performance of multiple hardware devices against various deep learning networks
- The SONOS tile-based architecture performed well across networks, with one notable exception: the Inception v3 network
- This performance difference could be explored – showing ATHENA's potential for codesign work

SANA-FE: NEUROMORPHIC SYSTEM MODELING & CODESIGN

- Tools are needed to rapidly estimate performance of neuromorphic architectures for design-space exploration
- General & extensible spiking H/W simulator
- Model functional behavior & track performance
- Schedule messages & intra-core interactions
- Calibrate simulator to real-world systems

- Accurately predicts latency & energy of gesture categorization spiking neural network (SNN)
- Faster than existing simulator (NeMo)

SANA-FE: Simulating Advanced Neuromorphic Architectures for Fast Exploration

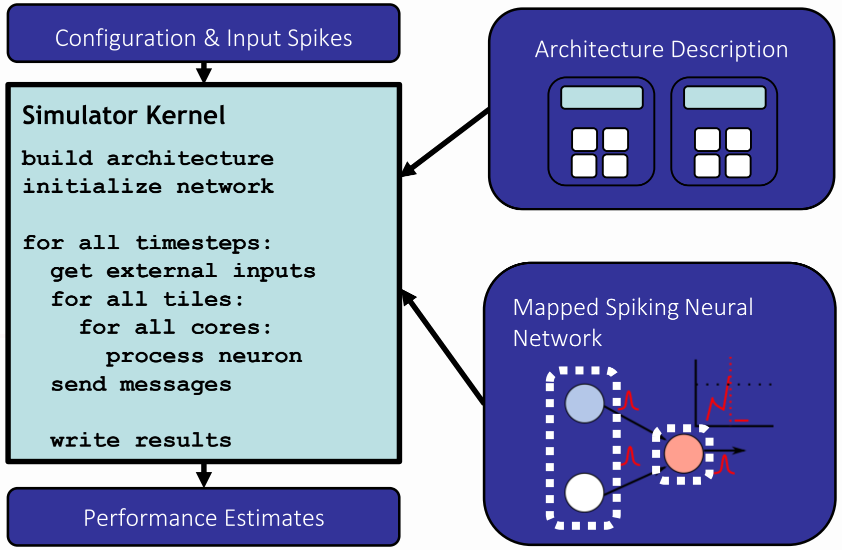

Latency Estimation for Hand Gesture Categorization

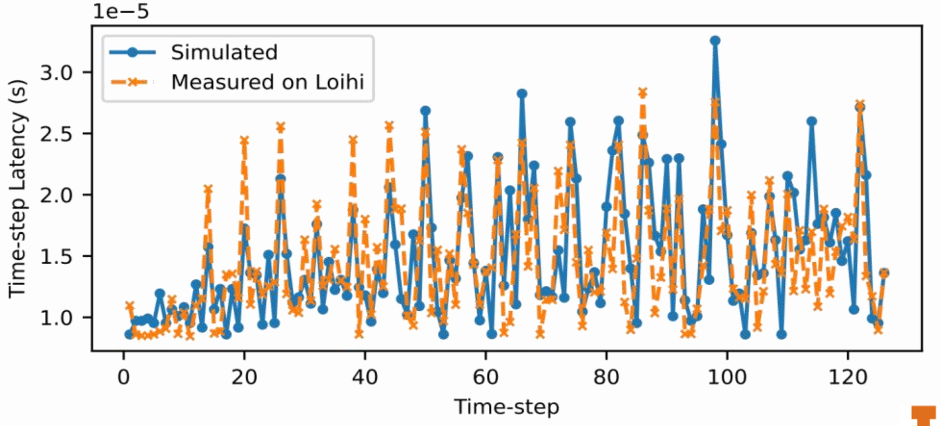

NeMo, Plagge et al. SIGSIM-PADS 2016

J. Boyle et al. 2023 (In Review)

SANA-FE: NEUROMORPHIC SYSTEM MODELING & CODESIGN

SANA-FE: Simulating Advanced Neuromorphic Architectures for Fast Exploration

- **Neuromorphic system architectures**
  - System design space exploration
  - Spiking Neural Network-based H/W
  - Novel features e.g., analog compute
  - Modeling & benchmarking for codesign

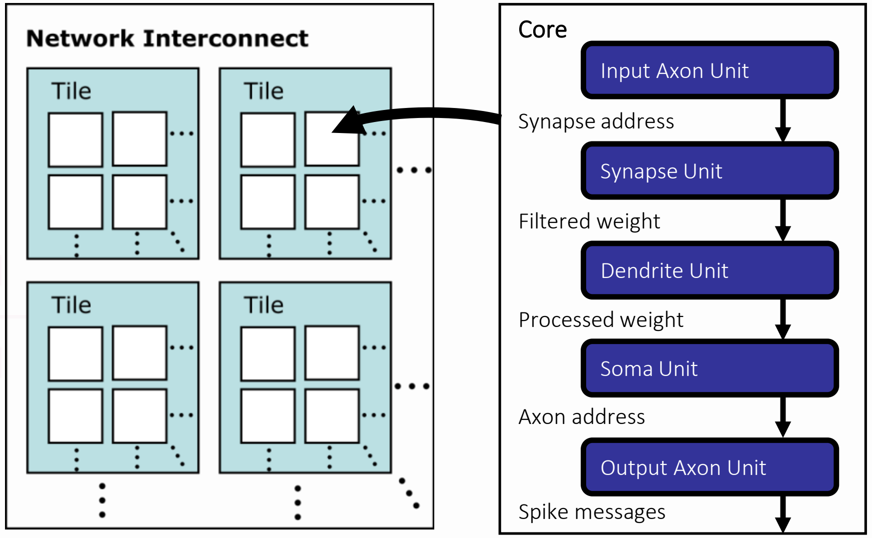

- **Fast spiking H/W simulation**
  - General & extensible framework
  - Functional model & performance
  - Calibrate simulator to real-world H/W
  - Accurately estimate latency & energy

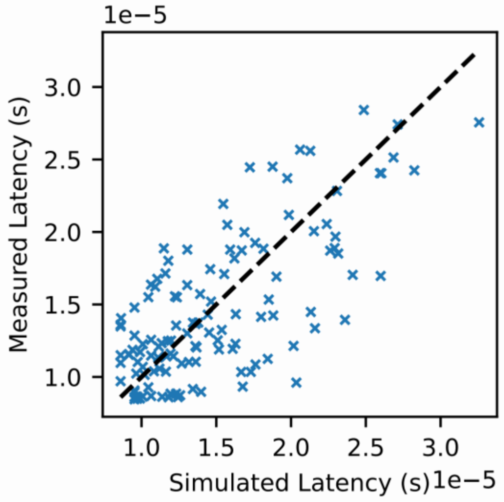

J. Boyle et al. 2023 (In Review)

NEUROMORPHIC COMPUTING CHALLENGE

SCALABILITY VS. COMPLEXITY

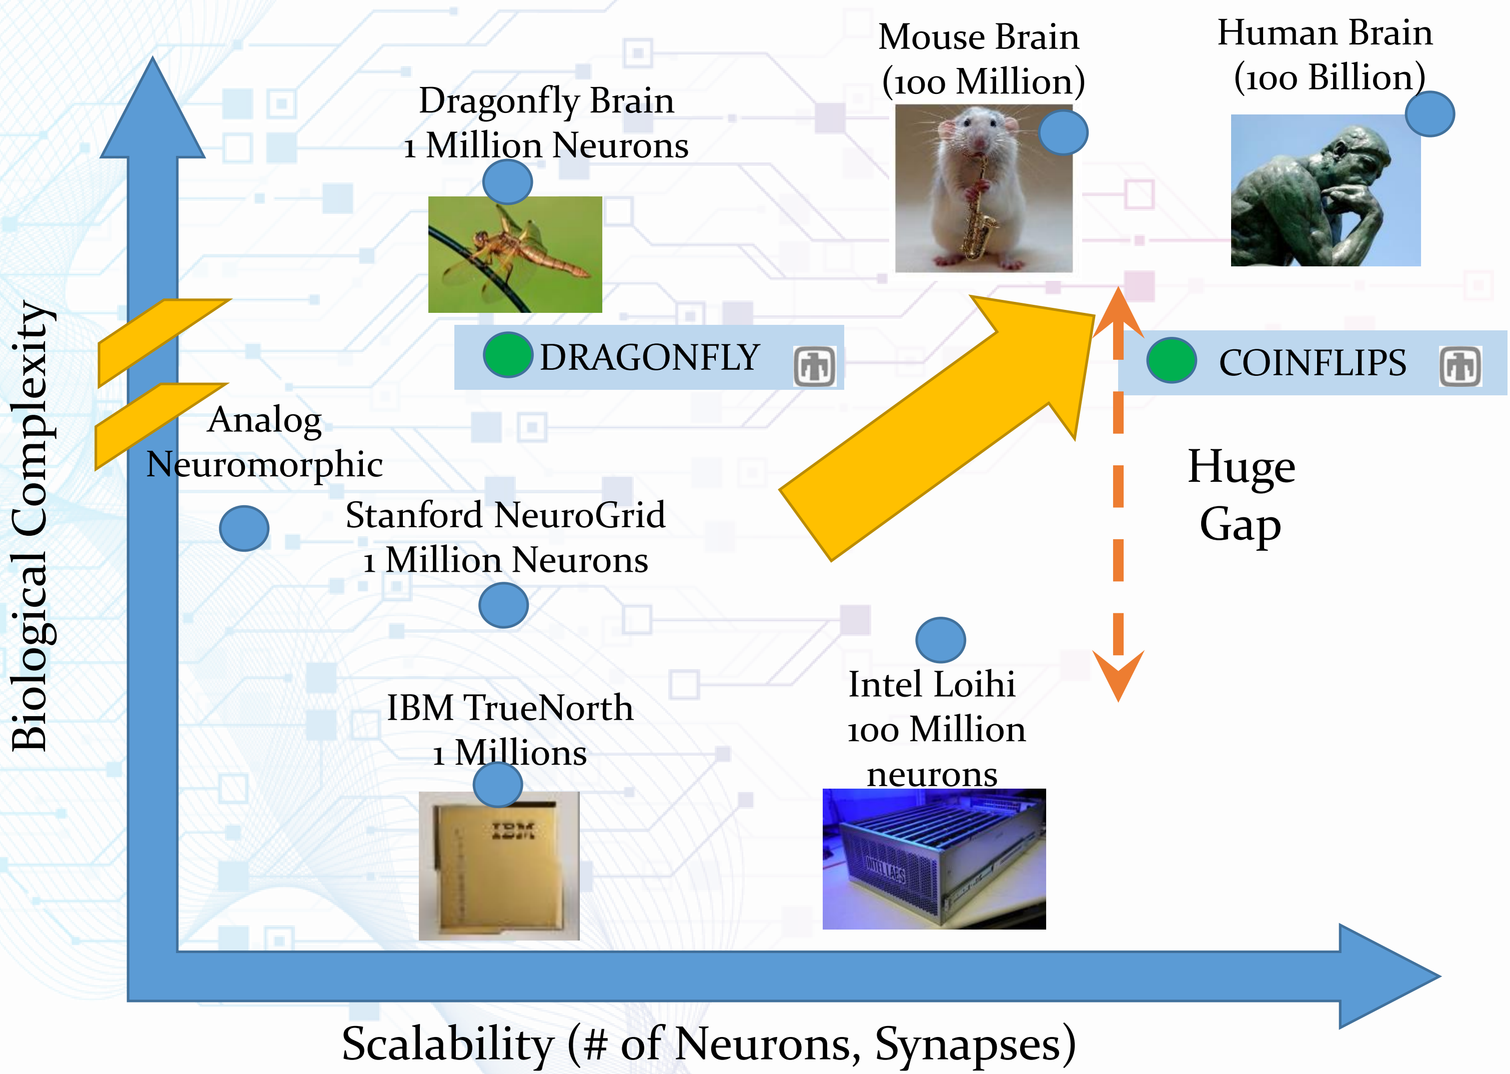

Hardware aware algorithms are critical for AI

- Novel neuroscience information translated to algorithms and models will influence next-generation devices, architectures
- Novel Algorithms
- Novel Devices
- Increased connectivity and communication (3D, wafer-scale, photonics)?

NEUROMORPHIC APPLICATIONS

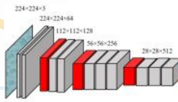

AI/ML Applications
ANNs
SNNs

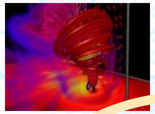

Scientific Computing
- Random Walks

High-fidelity Physics Simulations

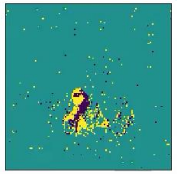

Edge Computing
- Event sensors
- Spatio-temporal processing

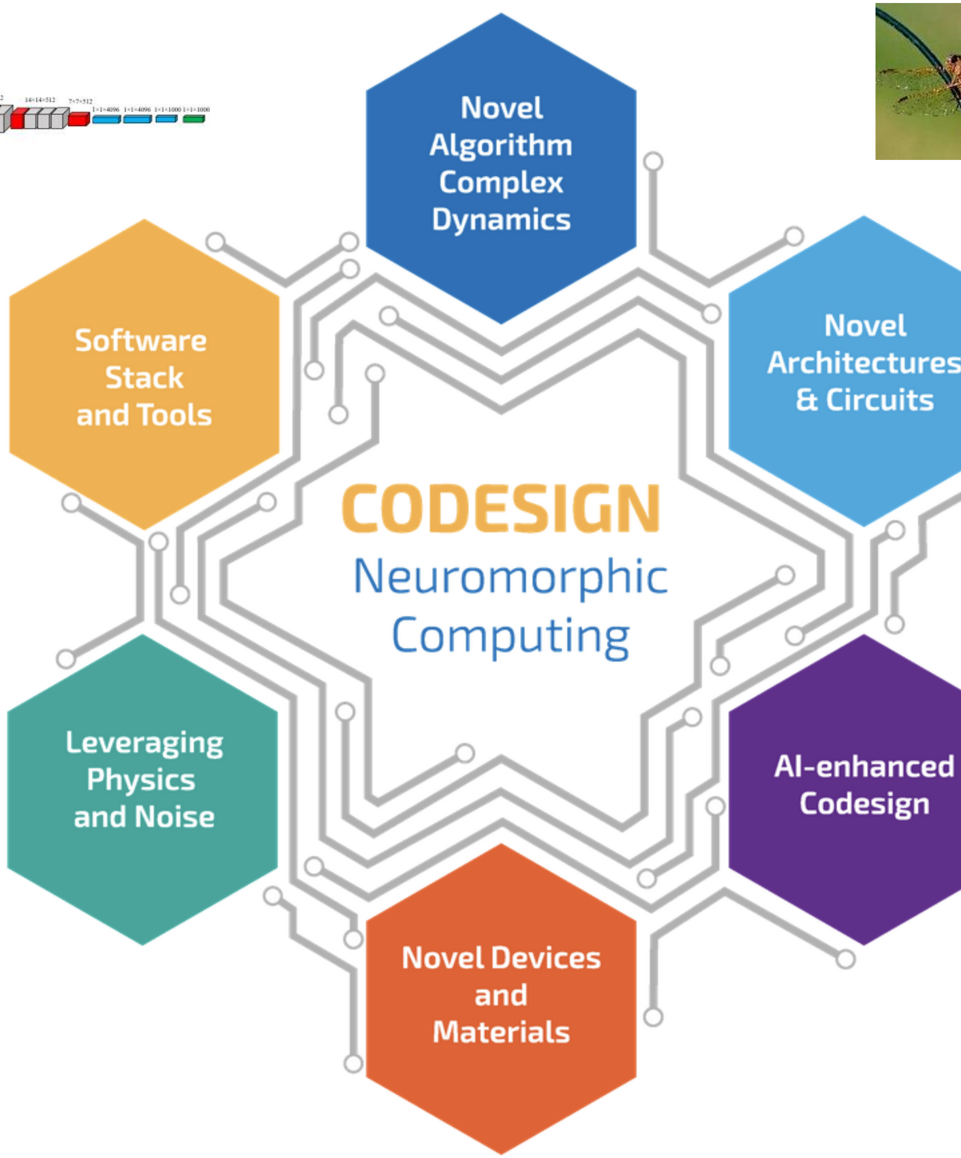

Brain-inspired Algorithms
Dragonfly
Dendritic Processing

Probabilistic Computing
COINFLIPS

Heterogeneous Computing Applications

GOING FORWARD

price per computation

Finally, policymakers should take proactive steps to ensure that researchers with small or moderate budgets can effectively contribute to the AI research field.

Limited direct neuromorphic investment even though it has been a more established field

Concentrating state-of-the-art technologies among the small number of research centers possessing extremely large compute budgets risks creating oligopolistic markets and shrinking the talent pool and opportunities for researchers.

Enable resources for research. Not just compute but training

The price of computations in gigaFLOPS has not decreased since 2017.15

Major overhauls of the computing paradigm like quantum computing or neuromorphic chips might one day allow for vast amounts of plentiful new compute.

LONG-TERM GOALS FOR NEXT-GENERATION OF NEUROMORPHIC SYSTEMS

- Heterogeneous architectures
- CoDesign to optimize communication and memory bottlenecks
- 3D architectures, Photonics

- Algorithms are cognizant of architecture and device constraints.
- Leverage the complex dynamics of devices.
- Bio-inspired techniques, adoption in computing

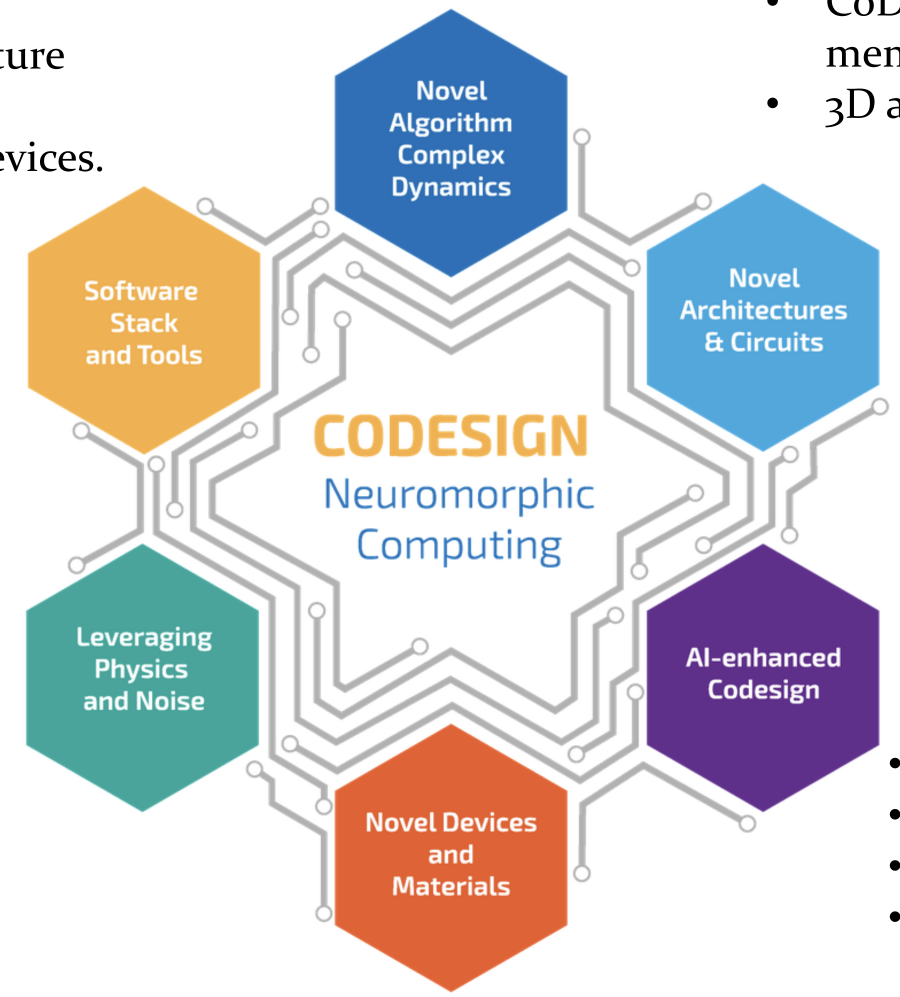

- How can AI-enhanced techniques accelerate scientific discovery?
- Different AI techniques at the device, circuit, system design and architecture level.
- Enable encoding of domain knowledge
- Enable concurrent contribution from researchers

- Software tools to support design and development
- Integration with AI-enhanced techniques?

- Leverage the physics of devices to do computation (analog)
- Embrace stochasticity of devices
- Analog devices are noisy. How can we incorporate this into algorithms?

- Novel devices with complex dynamics
- Radiation-hardened devices
- Reconfigurable devices
- Computational efficiency and computational density

THANK YOU!

References

- Mehonic, A., & Kenyon, A. J. (2022). Brain-inspired computing needs a master plan. *Nature*, *604*(7905), 255-260.
- Hasler, Jennifer. "Opportunities in physical computing driven by analog realization." *2016 IEEE international conference on rebooting computing (ICRC)*. IEEE, 2016.
- Cardwell, Suma George, Craig Vineyard, Willam Severa, Frances S. Chance, Frederick Rothganger, Felix Wang, Srideep Musuvathy, Corinne Teeter, and James B. Aimone. "Truly heterogeneous HPC: Co-design to achieve what science needs from HPC." In *Smoky Mountains Computational Sciences and Engineering Conference*, pp. 349-365. Springer, Cham, 2020.
- Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., & others. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. Ieee Micro, 38(1), 82–99.
- Furber, Steve B., Francesco Galluppi, Steve Temple, and Luis A. Plana. "The spinnaker project." *Proceedings of the IEEE* 102, no. 5 (2014): 652-665.
- McHarg, Matthew G., et al. "Falcon Neuro: an event-based sensor on the International Space Station." *Optical Engineering* 61.8 (2022): 085105.
- Gonzalez-Bellido, Paloma T., et al. "Eight pairs of descending visual neurons in the dragonfly give wing motor centers accurate population vector of prey direction." *Proceedings of the National Academy of Sciences* 110.2 (2013): 696-701.
- Polsky, A., Mel, B. W., & Schiller, J. (2004). Computational subunits in thin dendrites of pyramidal cells. Nature Neuroscience, 7(6), 621–627
- Posch, Christoph, Teresa Serrano-Gotarredona, Bernabe Linares-Barranco, and Tobi Delbruck. "Retinomorphic event-based vision sensors: bioinspired cameras with spiking output." *Proceedings of the IEEE* 102, no. 10 (2014): 1470-1484.
- Häusser, M., Spruston, N., & Stuart, G. J. (2000). Diversity and dynamics of dendritic signaling. Science, 290(5492), 739–744.
- London, M., & Häusser, M. (2005). Dendritic computation. Annu. Rev. Neurosci., 28, 503–532.
- Kumar, Suhas, R. Stanley Williams, and Ziwen Wang. "Third-order nanocircuit elements for neuromorphic engineering." *Nature* 585.7826 (2020): 518-523.
- George-Cardwell, S., et al. "Low power dendritic computation for wordspotting." *Journal of Low Power Electronics and Applications* 3.2 (2013): 73-98.
- Ramakrishnan, S., Wunderlich, R., Hasler, J., & George-Cardwell, S. (2013). Neuron array with plastic synapses and programmable dendrites. IEEE Transactions on Biomedical Circuits and Systems, 7(5), 631–642.
- Aly, Mohamed M. Sabry, et al. "The N3XT approach to energy-efficient abundant-data computing." *Proceedings of the IEEE* 107.1 (2018): 19-48.
- George-Cardwell, S., Hasler, J., Koziol, S., Nease, S., & Ramakrishnan, S. (2013). Low power dendritic computation for wordspotting. Journal of Low Power Electronics and Applications, 3(2), 73–98.
- Häusser, M., Spruston, N., & Stuart, G. J. (2000). Diversity and dynamics of dendritic signaling. Science, 290(5492), 739–744.
- Rall, W. (1995). 4.2 Theoretical Significance of Dendritic Trees for Neuronal Input-Output Relations (1964), in Neural Theory and Modeling, ed. RF Reiss, Palo Alto: Stanford University Press. The Theoretical Foundation of Dendritic Function: Selected Papers of Wilfrid Rall with Commentaries, 122.
- Ganesan, Timothy, I. Elamvazuthi, and Pandian Vasant. "Multiobjective design optimization of a nano-CMOS voltage-controlled oscillator using game theoretic-differential evolution." *Applied Soft Computing* 32 (2015): 293-299.
- Goldie, Anna, and Azalia Mirhoseini. "Placement optimization with deep reinforcement learning." *Proceedings of the 2020 International Symposium on Physical Design*. 2020.
- Cooksey, K. Daniel, and Dimitri Mavris. "Game theory as a means of modeling system of systems viability and feasibility." *2011 Aerospace Conference*. IEEE, 2011.

References

- George, Suma, Sihwan Kim, Sahil Shah, Jennifer Hasler, Michelle Collins, Farhan Adil, Richard Wunderlich, Stephen Nease, and Shubha Ramakrishnan. "A programmable and configurable mixed-mode FPAA SoC." *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 24, no. 6 (2016): 2253-2261.
- Lu, Jenny, et al. "Transforming representations of movement from body-to world-centric space." *Nature* 601.7891 (2022): 98-104.
- Boahen, Kwabena. "Dendrocentric learning for synthetic intelligence." *Nature* 612, no. 7938 (2022): 43-50.
- Farquhar, Ethan, and Paul Hasler. "A bio-physically inspired silicon neuron." *IEEE Transactions on Circuits and Systems I: Regular Papers* 52, no. 3 (2005): 477-488.
- Indiveri, Giacomo, et al. "Neuromorphic silicon neuron circuits." *Frontiers in neuroscience* 5 (2011): 73.
- Saha, Arnob, et al. "Intrinsic synaptic plasticity of ferroelectric field effect transistors for online learning." *Applied Physics Letters* 119.13 (2021): 133701.
- Parameshwara, Chethan M., Simin Li, Cornelia Fermüller, Nitin J. Sanket, Matthew S. Evanusa, and Yiannis Aloimonos. "SpikeMS: Deep spiking neural network for motion segmentation." In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3414-3420. IEEE, 2021.
- Nease, Stephen, Suma George, Paul Hasler, Scott Koziol, and Stephen Brink. "Modeling and implementation of voltage-mode CMOS dendrites on a reconfigurable analog platform." *IEEE transactions on biomedical circuits and systems* 6, no. 1 (2011): 76-84.
- Chance, Frances S., and Suma G. Cardwell. "Shunting Inhibition as a Neural-Inspired Mechanism for Multiplication in Neuromorphic Architectures." In *Neuro-Inspired Computational Elements Conference*, pp. 41-46. 2023.
- Plagge, Mark, Ben Feinberg, John McFarland, Fred Rothganger, Sapan Agarwal, Amro Awad, Clayton Hughes, and Suma G. Cardwell. "ATHENA: Enabling Codesign for Next-Generation AI/ML Architectures." In *2022 IEEE International Conference on Rebooting Computing (ICRC)*, pp. 13-23. IEEE, 2022.
- Misra, Shashank, Leslie C. Bland, Suma G. Cardwell, Jean Anne C. Incorvia, Conrad D. James, Andrew D. Kent, Catherine D. Schuman, J. Darby Smith, and James B. Aimone. "Probabilistic Neural Computing with Stochastic Devices." *Advanced Materials* (2022): 2204569
- Cardwell, Suma G., Catherine D. Schuman, J. Darby Smith, Karan Patel, Jaesuk Kwon, Samuel Liu, Christopher Allemang, Shashank Misra, Jean Anne Incorvia, and James B. Aimone. "Probabilistic Neural Circuits leveraging AI-Enhanced Codesign for Random Number Generation." *arXiv preprint arXiv:2212.00625* (2022).
- Vineyard, Craig, Suma Cardwell, Frances Chance, Srideep Musuvathy, Fred Rothganger, William Severa, John Smith, Corinne Teeter, Felix Wang, and James Aimone. "Neural Mini-Apps as a Tool for Neuromorphic Computing Insight." In *Neuro-Inspired Computational Elements Conference*, pp. 40-49. 2022.
- Theilman, Bradley H., and James B. Aimone. "Goemans-Williamson MAXCUT approximation algorithm on Loihi." *Neuro-Inspired Computational Elements Conference*. 2023.
- Theilman, Bradley H., and James B. Aimone. "Goemans-Williamson MAXCUT approximation algorithm on Loihi." In *Neuro-Inspired Computational Elements Conference*, pp. 1-5. 2023.
- Aimone, James B., William Severa, and Craig M. Vineyard. "Composing neural algorithms with Fugu." In *Proceedings of the International Conference on Neuromorphic Systems*, pp. 1-8. 2019.
- Jeong, Geonhwa, Gokcen Kestor, Prasanth Chatarasi, Angshuman Parashar, Po-An Tsai, Sivasankaran Rajamanickam, Roberto Gioiosa, and Tushar Krishna. "Union: A unified HW-SW co-design ecosystem in MLIR for evaluating tensor operations on spatial accelerators." In *2021 30th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pp. 30-44. IEEE, 2021.
- Rodrigues, Arun F., K. Scott Hemmert, Brian W. Barrett, Chad Kersey, Ron Oldfield, Marlo Weston, Rolf Risen et al. "The structural simulation toolkit." *ACM SIGMETRICS Performance Evaluation Review* 38, no. 4 (2011): 37-42.
- Xiao, T., Bennett, C., Feinberg, B., Marinella, M., & Agarwal, S. (2022). *CrossSim Inference Manual v2. 0* (No. SAND2022-7074R). Sandia National Lab.(SNL-NM), Albuquerque, NM (United States); Sandia National Laboratories,, Livermore, CA.