

A survey of Emerging Beyond CMOS Devices and Architectures



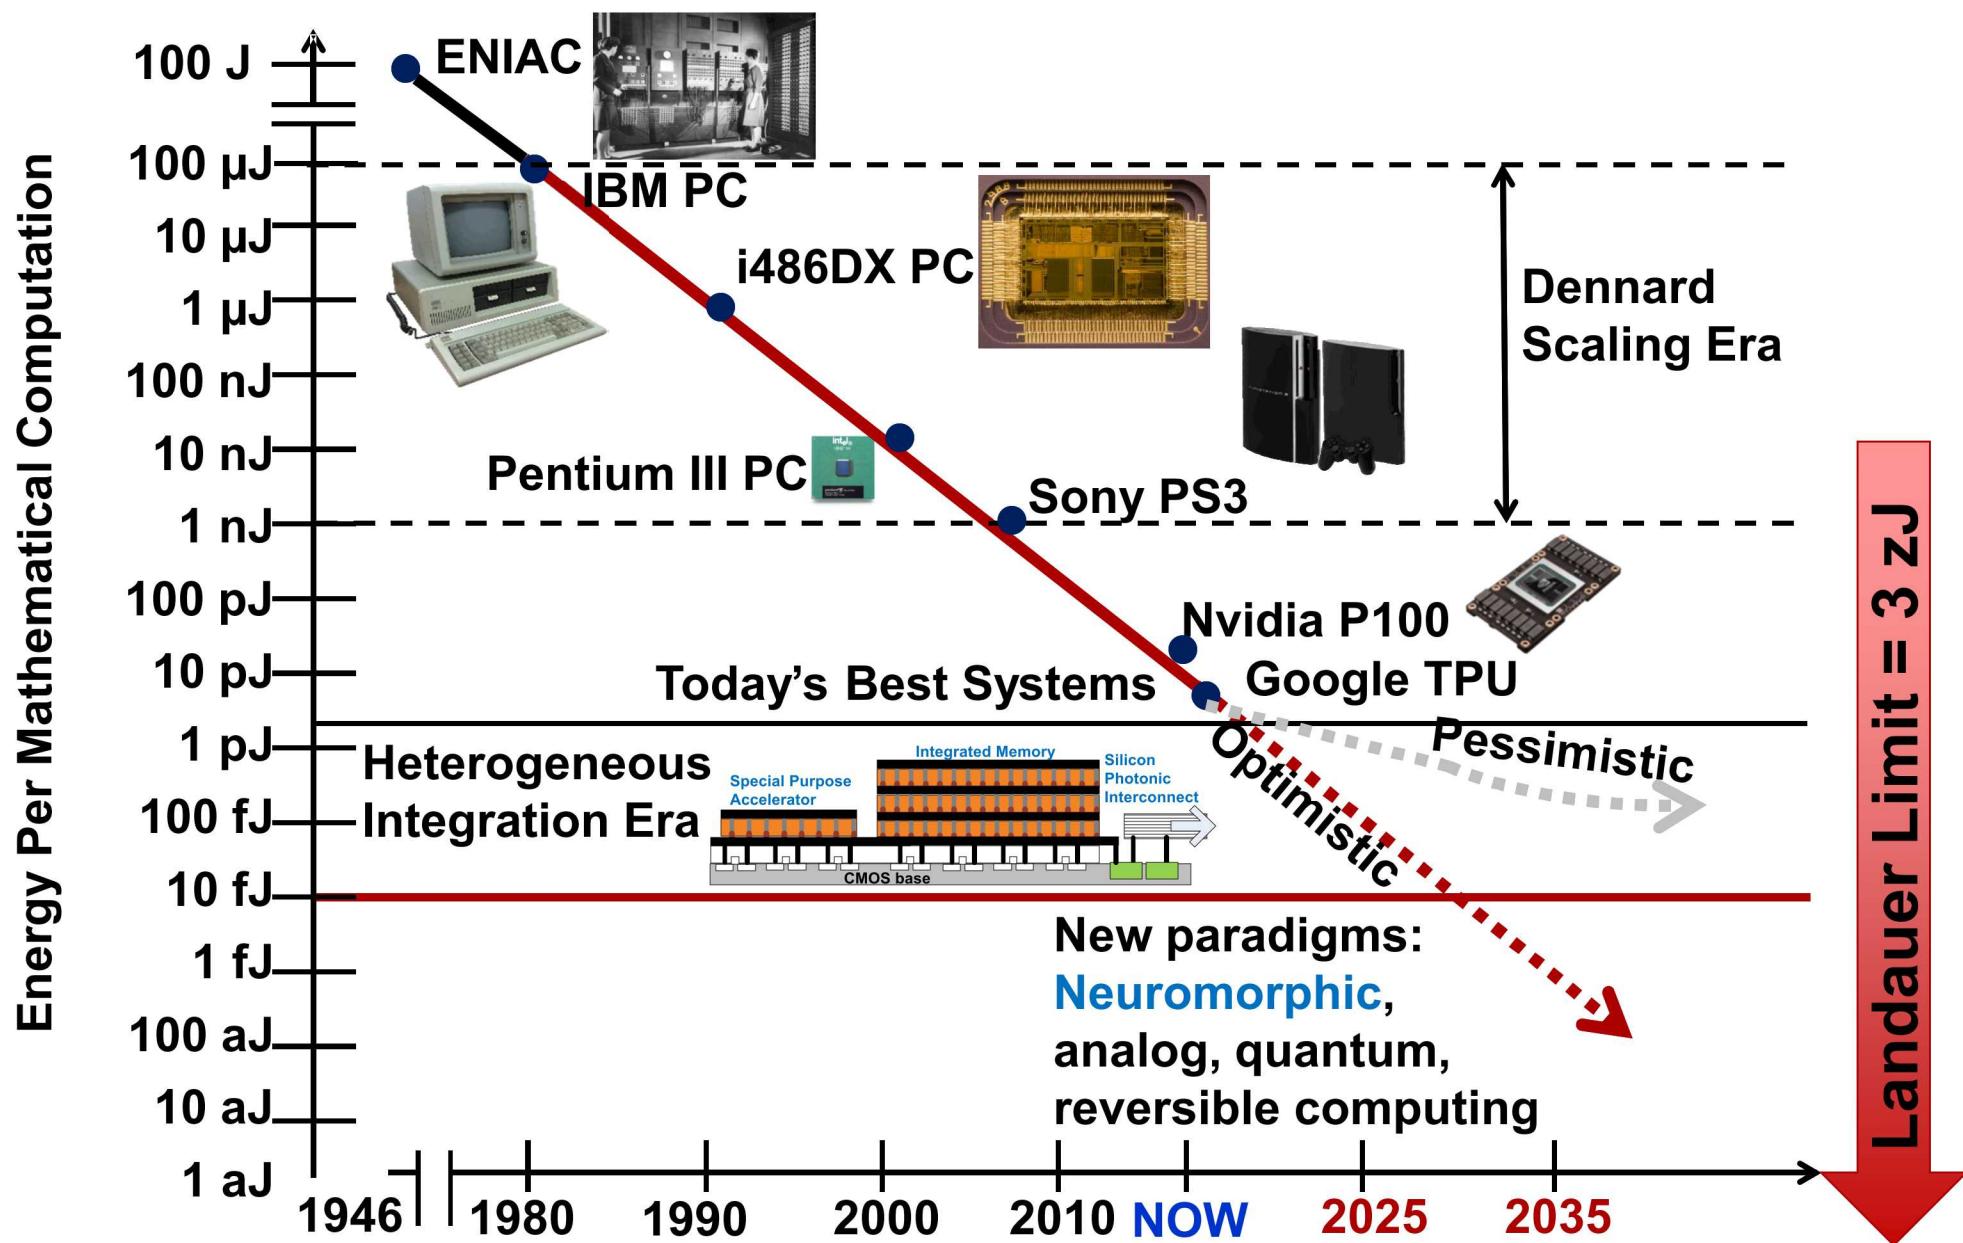
Sapan Agarwal

Sandia National Laboratories



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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.

Evolution of Computing Machinery

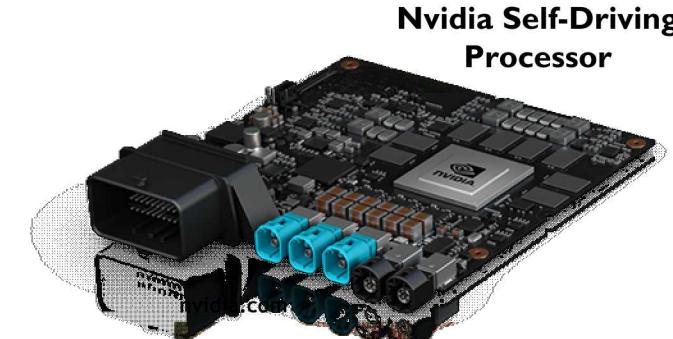


Computing Across Power Envelopes

IoT, Edge, and Mobile Computing



Self Driving Cars, Unmanned Arial Vehicles, and Satellite Computing



Datacenters, HPC



ASCI Red Supercomputer

1W

10W

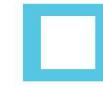
10^2 W

10^3 W

10^4 W

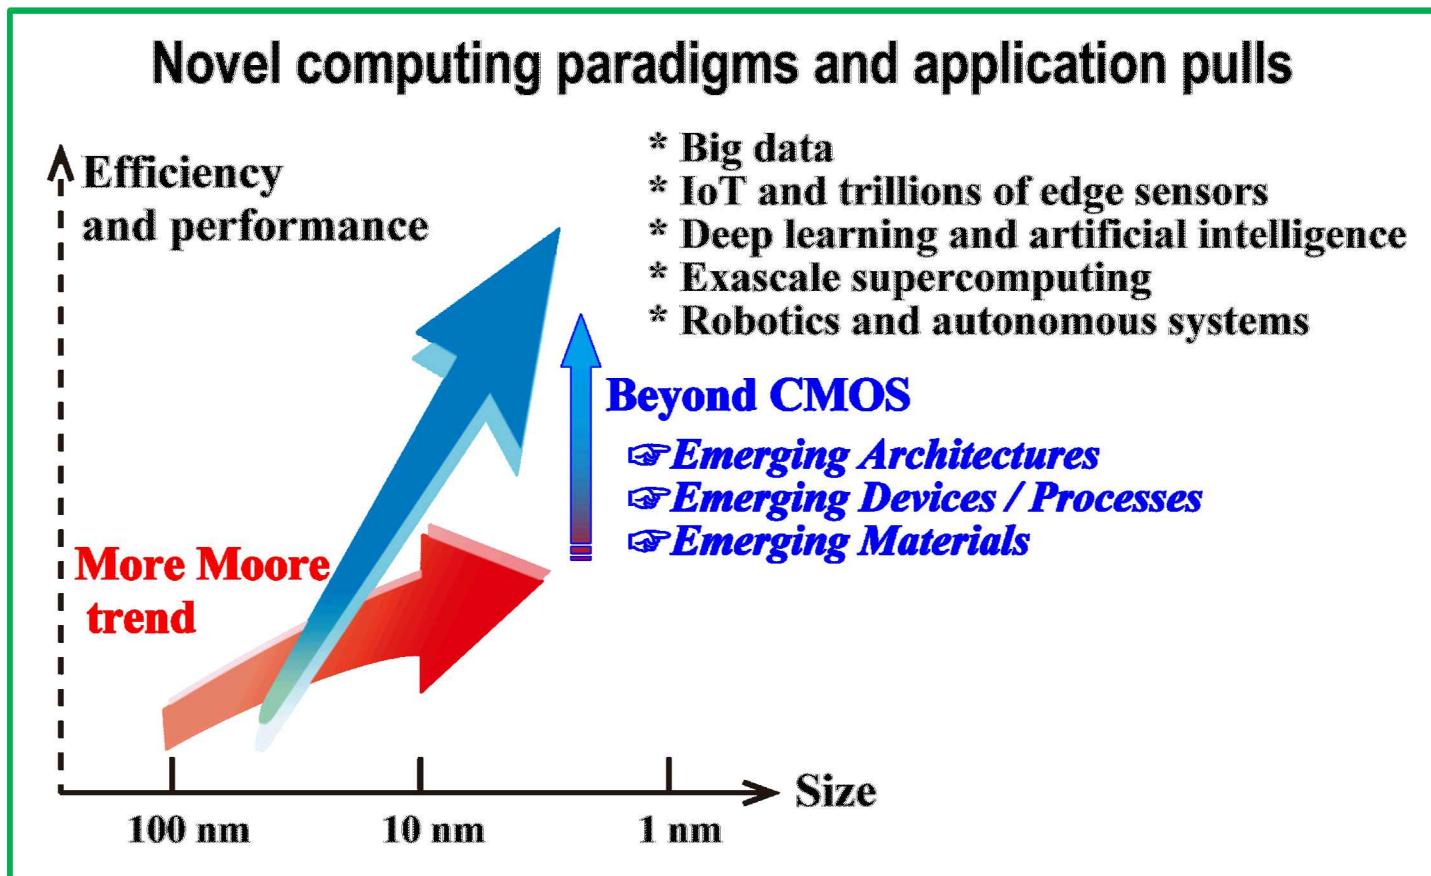
10^5 W

10^6 W

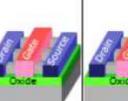
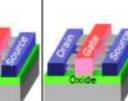
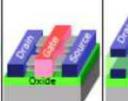
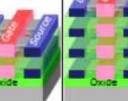


Successor to ITRS (International Technology Roadmap for Semiconductors)

Objective of the Beyond CMOS (BC) Chapter



Still road mapping near term semiconductors (Moore Moore)

YEAR OF PRODUCTION	2020	2022	2025	2028	2031	2034
Logic industry "Node Range" Labeling (nm)	G48M36	G45M24	G42M20	G40M16	G38M16T2	G38M16T4
"5"	"5"	"3"	"2.1"	"1.5"	"1.0 eq"	"0.7 eq"
IDM-Foundry node labeling	i7-f5	i5-f3	i3-f2.1	i2.1-f1.5	i1.5e-f1.0e	i1.0e-f0.7e
Logic device structure options	FinFET	finFET LGAA	LGAA	LGAA	LGAA-3D	LGAA-3D
Mainstream device for logic	finFET	finFET	LGAA	LGAA	LGAA-3D	LGAA-3D
						
LOGIC DEVICE GROUND RULES						
M _x pitch (nm)	36	32	24	20	16	16
M ₁ pitch (nm)	32	30	24	20	19	19
M ₀ pitch (nm)	30	24	20	16	16	16
Gate pitch (nm)	48	45	42	40	38	38
L _g - Gate Length - HP (nm)	18	16	14	12	12	12
L _g - Gate Length - HD (nm)	20	18	14	12	12	12
Channel overlap ratio - two-sided	0.20	0.20	0.20	0.20	0.20	0.20
Spacer width (nm)	7	6	5	4	4	4
Contact CD (nm) - finFET, LGAA	16	17	18	20	18	18
Contact CD (nm) - VGAA						
Device architecture key ground rules						
FinFET pitch (nm)	28.0	24.0				
FinFET Fin width (nm)	7.0	6.0				
FinFET Fin height (nm)	50	60				
Footprint drive efficiency - finFET	3.82	5.25				
Lateral GAA lateral pitch (nm)			22.0	20.0	20.0	20.0
Lateral GAA vertical pitch (nm)			18.0	16.0	14.0	14.0
Lateral GAA (nanosheet) thickness (nm)			7.0	6.0	5.0	5.0
Number of vertically stacked nanosheets			3	3	4	4
LGAA width (nm) - HP			30	20	15	10
LGAA width (nm) - HD			20	11	6	6
LGAA width (nm) - SRAM			7	6	6	6
LGAA total height (nm)			53	48	57	57
Footprint drive efficiency - lateral GAA - HP			4.80	4.59	5.52	5.00
Device effective width (nm) - HP	107.0	126.0	192.0	156.0	160.0	120.0
Device effective width (nm) - HD	107.0	126.0	132.0	102.0	88.0	88.0
Device lateral pitch (nm)	28	24	22	20	20	20
Device height (nm)	50.0	60.0	53.0	48.0	57.0	57.0
Device width (nm) - HP	7	6	25	20	15	10
Device width (nm) - HD	7	6	15	11	6	6
Device width (nm) - SRAM	7	6	7	6	6	6

2019 Beyond CMOS Team Members and Contributors

Sapan Agarwal

Brad Aimone
Hiro Akinaga
Otitaleke Akinola
Mustafa Badaroglu
Gennadi Bersuker
Christian Binek
Geoffrey Burr
Leonid Butov
Kerem Camsari
Gert Cauwenberghs
An Chen
Winston Chern
Supriyo Datta
John Dallesasse
Shamik Das
Erik DeBenedictis
Peter Dowben
Tetsuo Endoh
Ben Feinberg
Thomas Ferreira de Lima
Akira Fujiwara
Elliot Fuller
Michael Frank
Paul Franzon

Michael Fuhrer
Mike Garner
Chakku Gopalan
Bogdan Govoreanu
Cat Graves
Kohei Hamaya
Masami Hane
Jennifer Hasler
Yoshihiro Hayashi
Toshiro Hiramoto
D. Scott Holmes
Sharon Hu
Francesca Iacopi
Danielle Ilmeni
Jean Anne Incorvia
Engin Ipek
Satoshi Kamiyama
Kiyoshi Kawabata
Asif Khan
Hajime Kobayashi
Suhas Kumar
Ilya Krivorotov
Xiuling Li
Xiang (Shaun) Li
Shy-Jay Lin

Tsu-Jae King Liu
Matthew Marinella
Bicky A. Marque
Rivu Midya
Yoshiyuki Miyamoto
Johannes Muller
Azad Naeemi
Mitchell A. Nahmias
Emre Neftci
Mike Niemier
Dmitri Nikonov
Yutaka Ohno
Chenyun Pan
Ferdinand Peper
Shriram Ramanathan
Mingyi Rao
Shashi Paul
Paul R. Prucnal
Titash Rakshit
Arijit Raychowdhury
Sayeef Salahuddin
Shintaro Sato
Michael Schneider
Bhavin J. Shastri
Xia Sheng

Takahiro Shinada
Urmita Sikder
Greg Snider
John-Paul Strachan
Dimitri Strukov
Naoyuki Sugiyama
Tarek Taha
Alexander N. Tait
Shinichi Takagi
Norikatsu Takaura
Tsutomu Teduka
Yasuhide Tomioka
Wilman Tsai
Tohru Tsuruoka
Zhongrui Wang
R. Stanley Williams
Justin Wong
Dirk Wouters
Patrick Xiao
Kojiro Yagami
J. Joshua Yang
Noboyuki Yoshikawa
Victor Zhirnov

What's the Minimum Energy to Operate a Transistor?

Consider a signal Energy E_{signal}

The probability of an error due to thermal noise is:

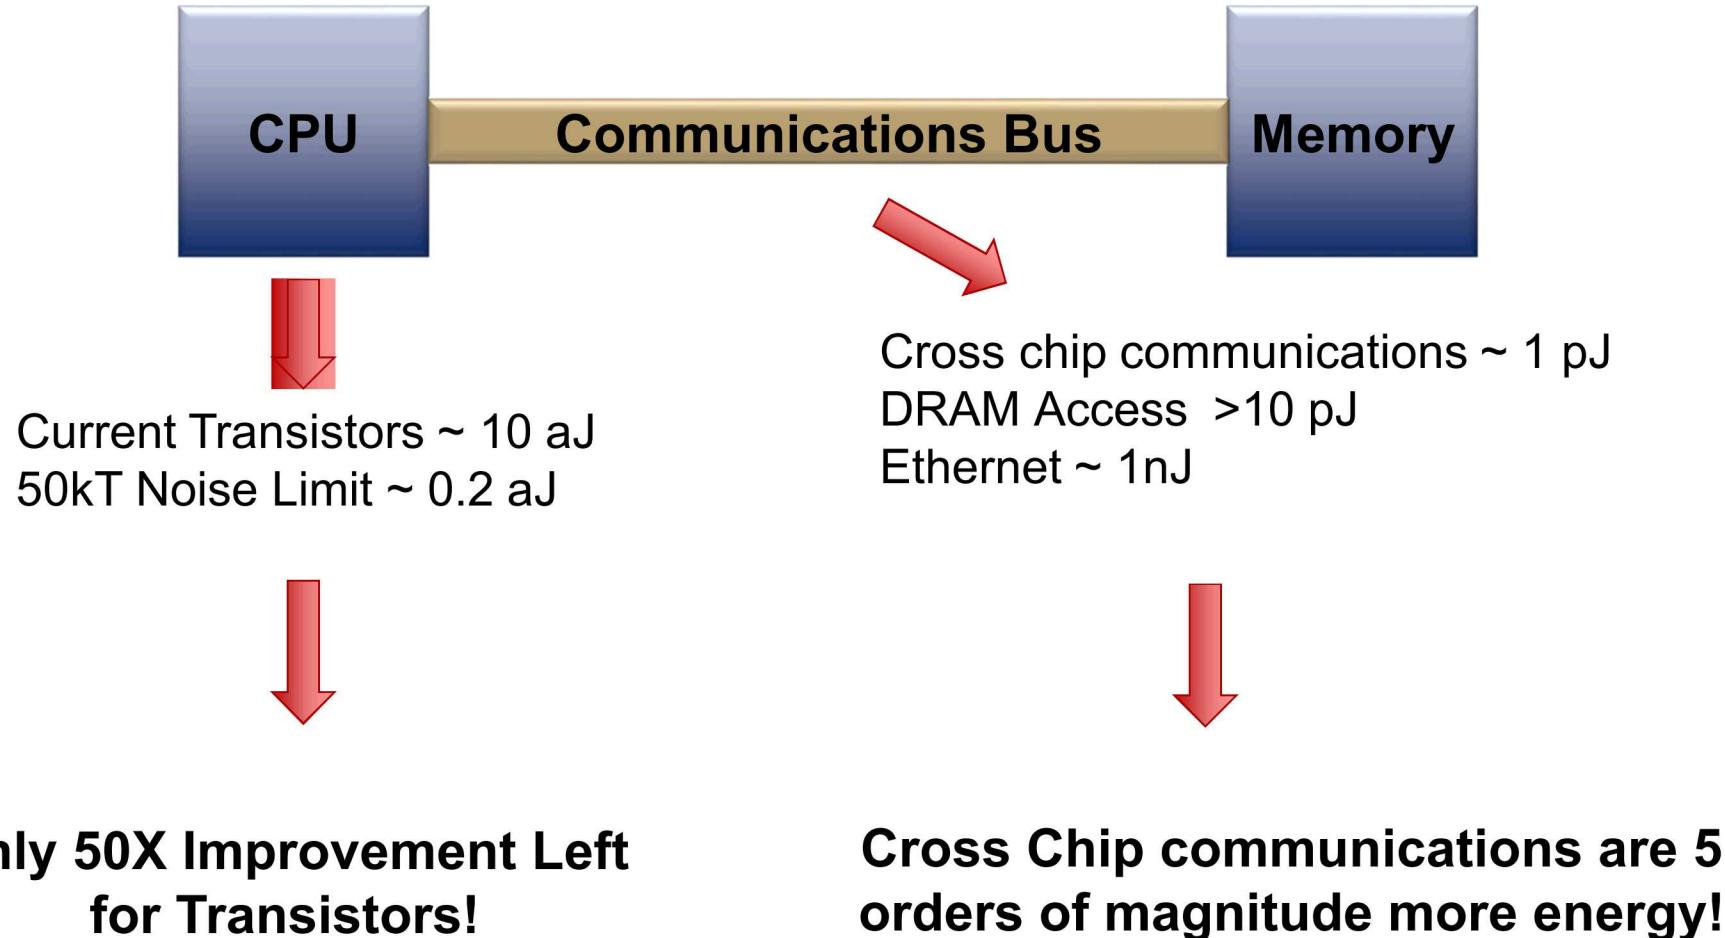
$$P(\text{Error}) = e^{-E_{signal}/kT}$$

In order to ensure a full system with billions of transistors is reliable, we need:

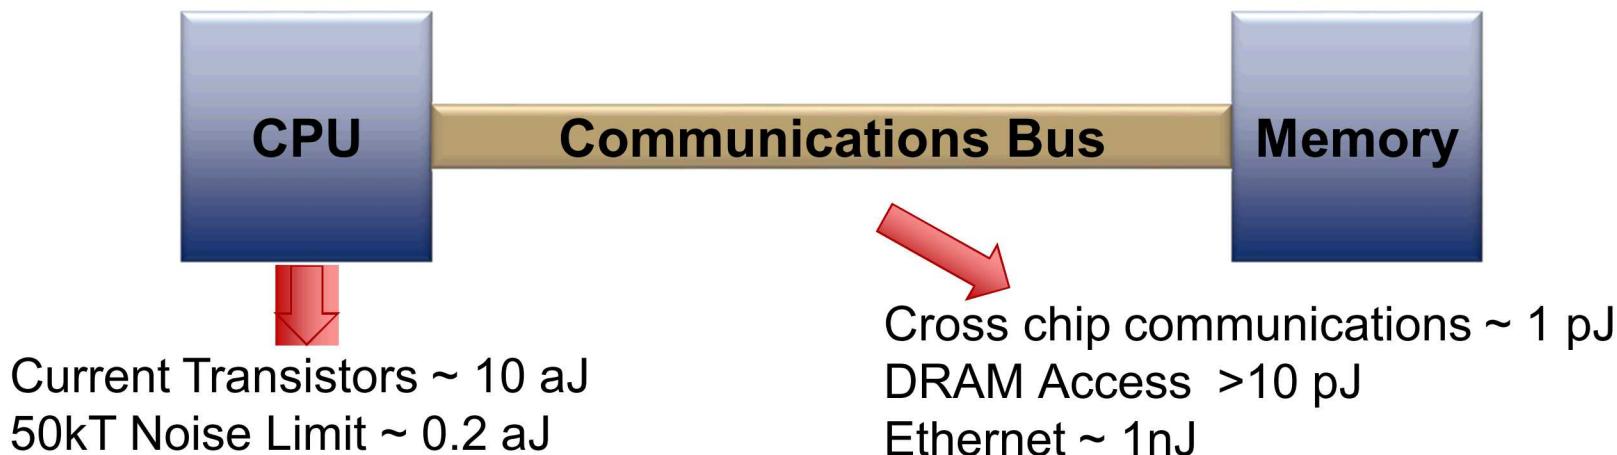
$$E_{signal} \sim 50 \text{ kT}$$

Landauer – Shannon Limit

How Efficient are Current Systems?



Beyond Moore Technologies



Extending Von Neumann

- Low Voltage or Novel Transistors
- Optical Communications
- Reduced Data Movement
 - New On-Chip Memory
 - Processing near Memory

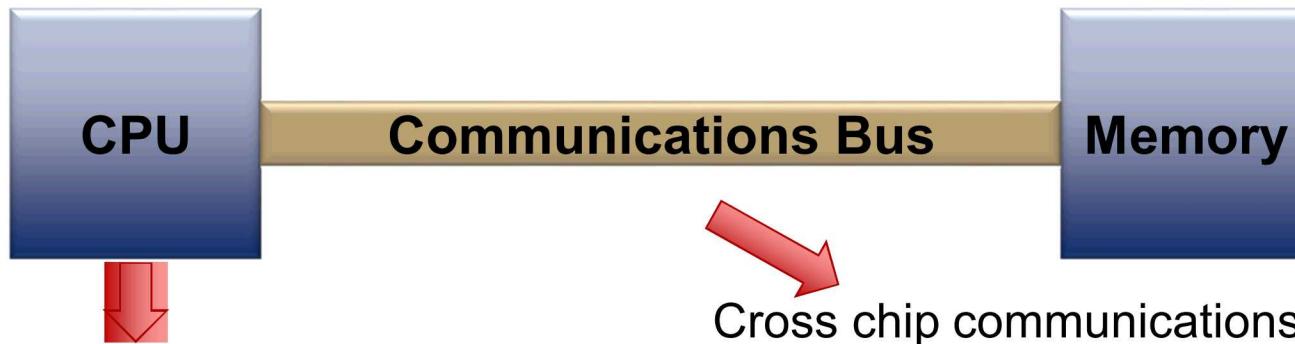
Alternate Computing Paradigms

- Neuromorphic
- Analog
- Computing with memory devices
- Quantum
- Stochastic
- Approximate

Going Below 50 kT

- Error Correction
- Reversible Computing
 - Adiabatic Computing / Energy Recycling
- Superconducting

Beyond Moore Technologies



Current Transistors ~ 10 aJ
50kT Noise Limit ~ 0.2 aJ

Cross chip communications ~ 1 pJ
DRAM Access >10 pJ
Ethernet ~ 1nJ

Extending Von Neumann

- Low Voltage or Novel Transistors
- Optical Communications
- Reduced Data Movement
 - New On-Chip Memory
 - Processing near Memory

Alternate Computing Paradigms

- Neuromorphic
- Analog
- Computing with memory devices
- Quantum
- Stochastic
- Approximate

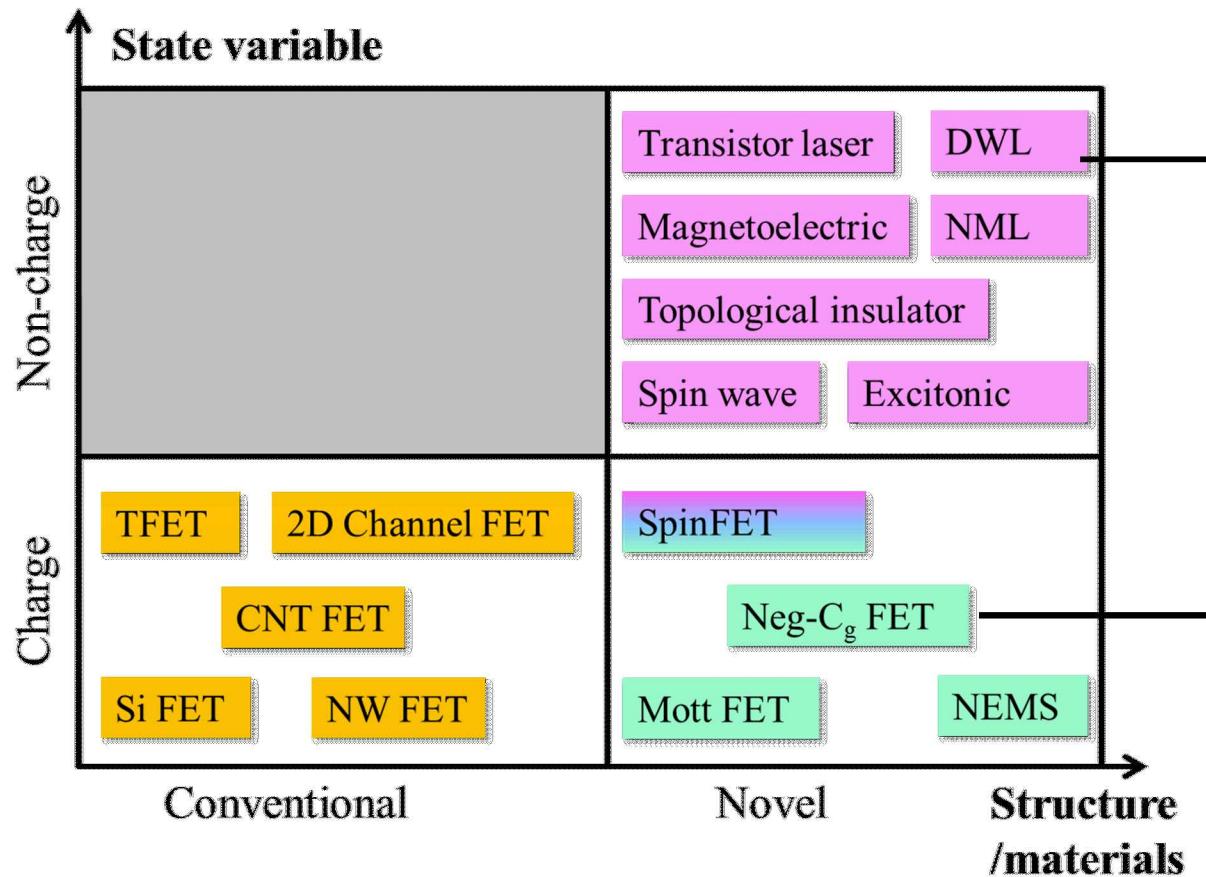
Going Below 50 kT

- Error Correction
- Reversible Computing
 - Adiabatic Computing / Energy Recycling
- Superconducting

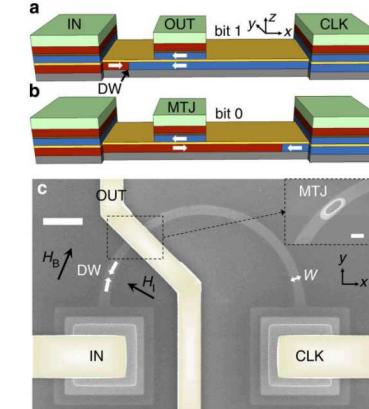
Extending Von Neumann – New Transistors

Lowering voltage lowers CV^2 energy of communications

Logic and Information Processing Devices

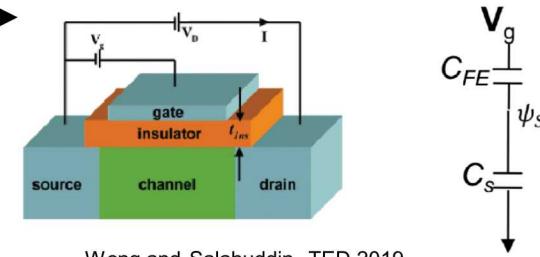


Domain Wall Logic



J.A. Incorvia et al, Nature Comm 7, 2016

Negative Capacitance FET

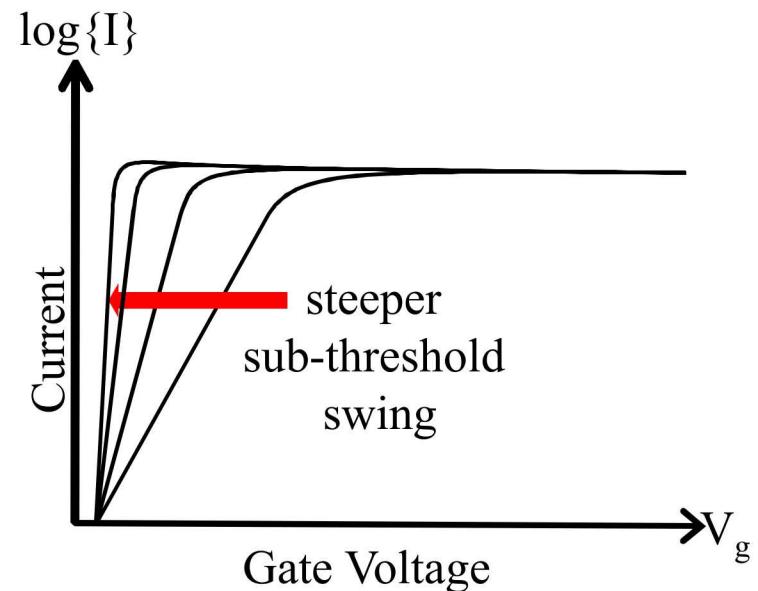


Wong and Salahuddin, TED 2019

A New Switch has to Satisfy Three Specifications

Low Active Power (CV² energy)

- Steepness (or sensitivity)
 - switches with only a few milli-volts
 - 60mV/decade \Rightarrow **1mV/decade**

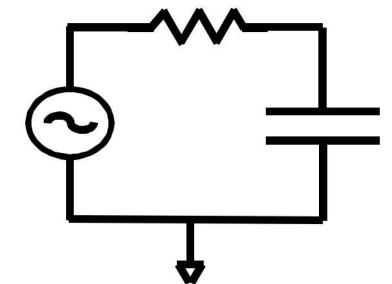


Low Leakage Power

- On/Off ratio: **10⁵ : 1**

High Speed (RC delay)

- High Conductance Density
 - **1 milli-Siemen/micron**



Need Restoring Logic for Novel State Variables

- Need to be able to drive multiple output transistors across multiple stages of logic

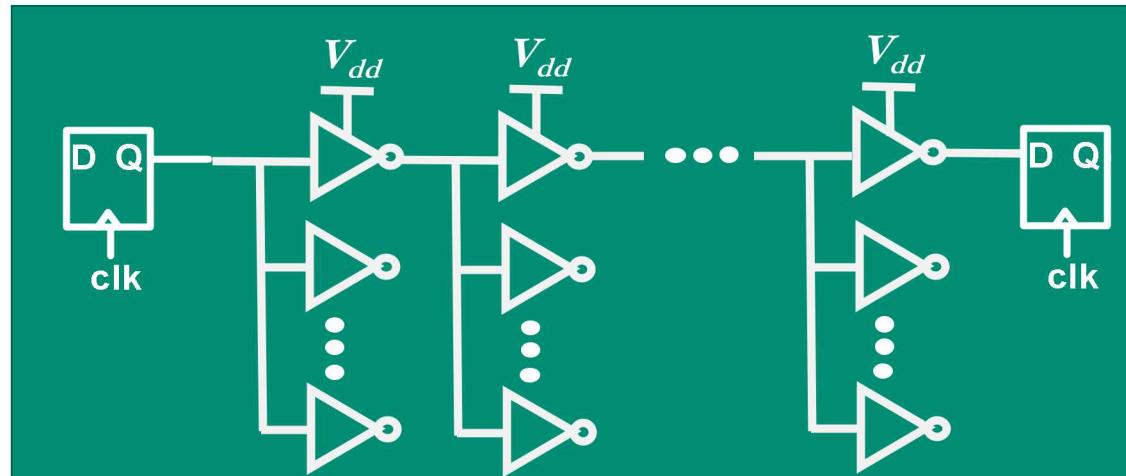
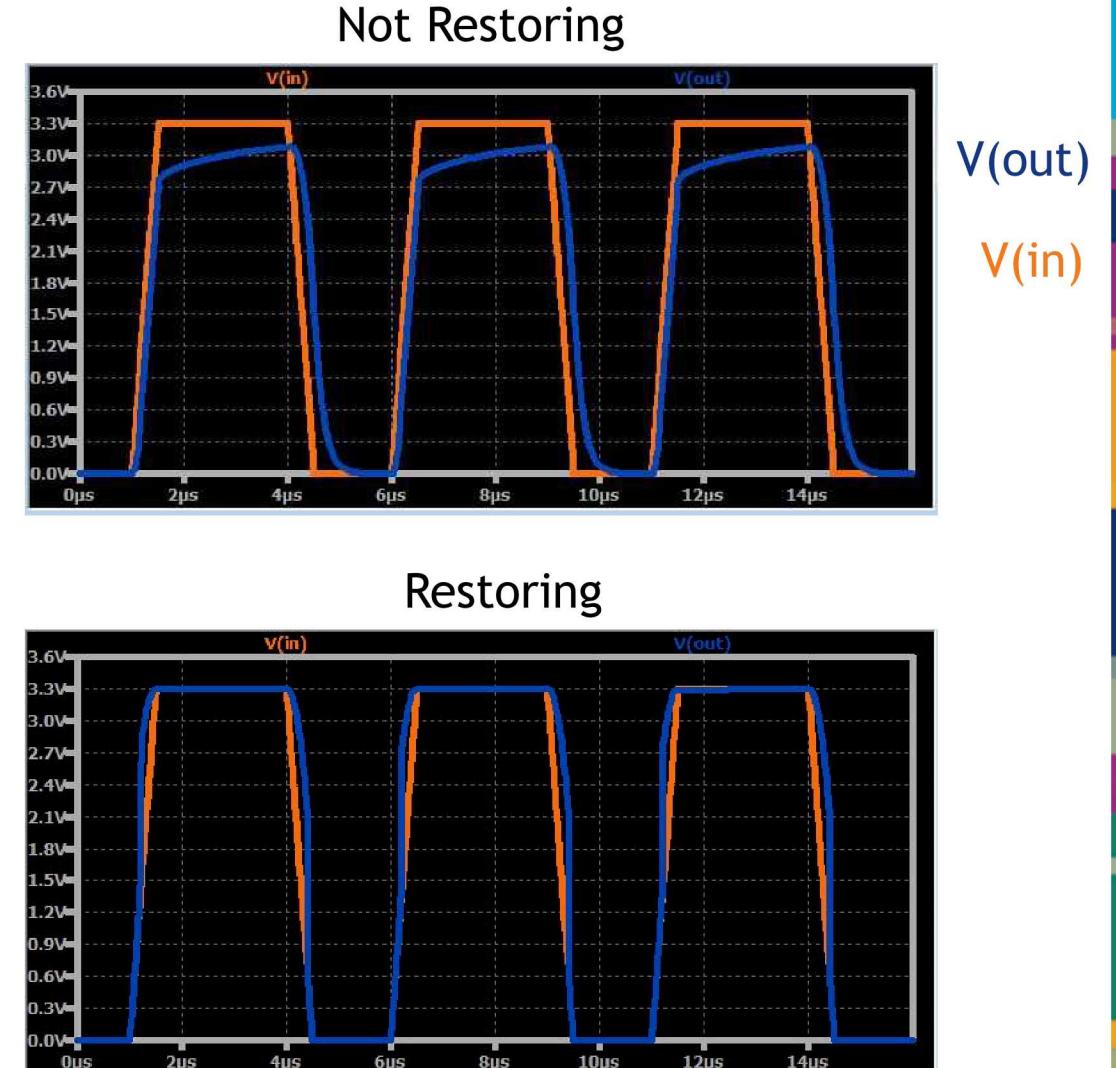


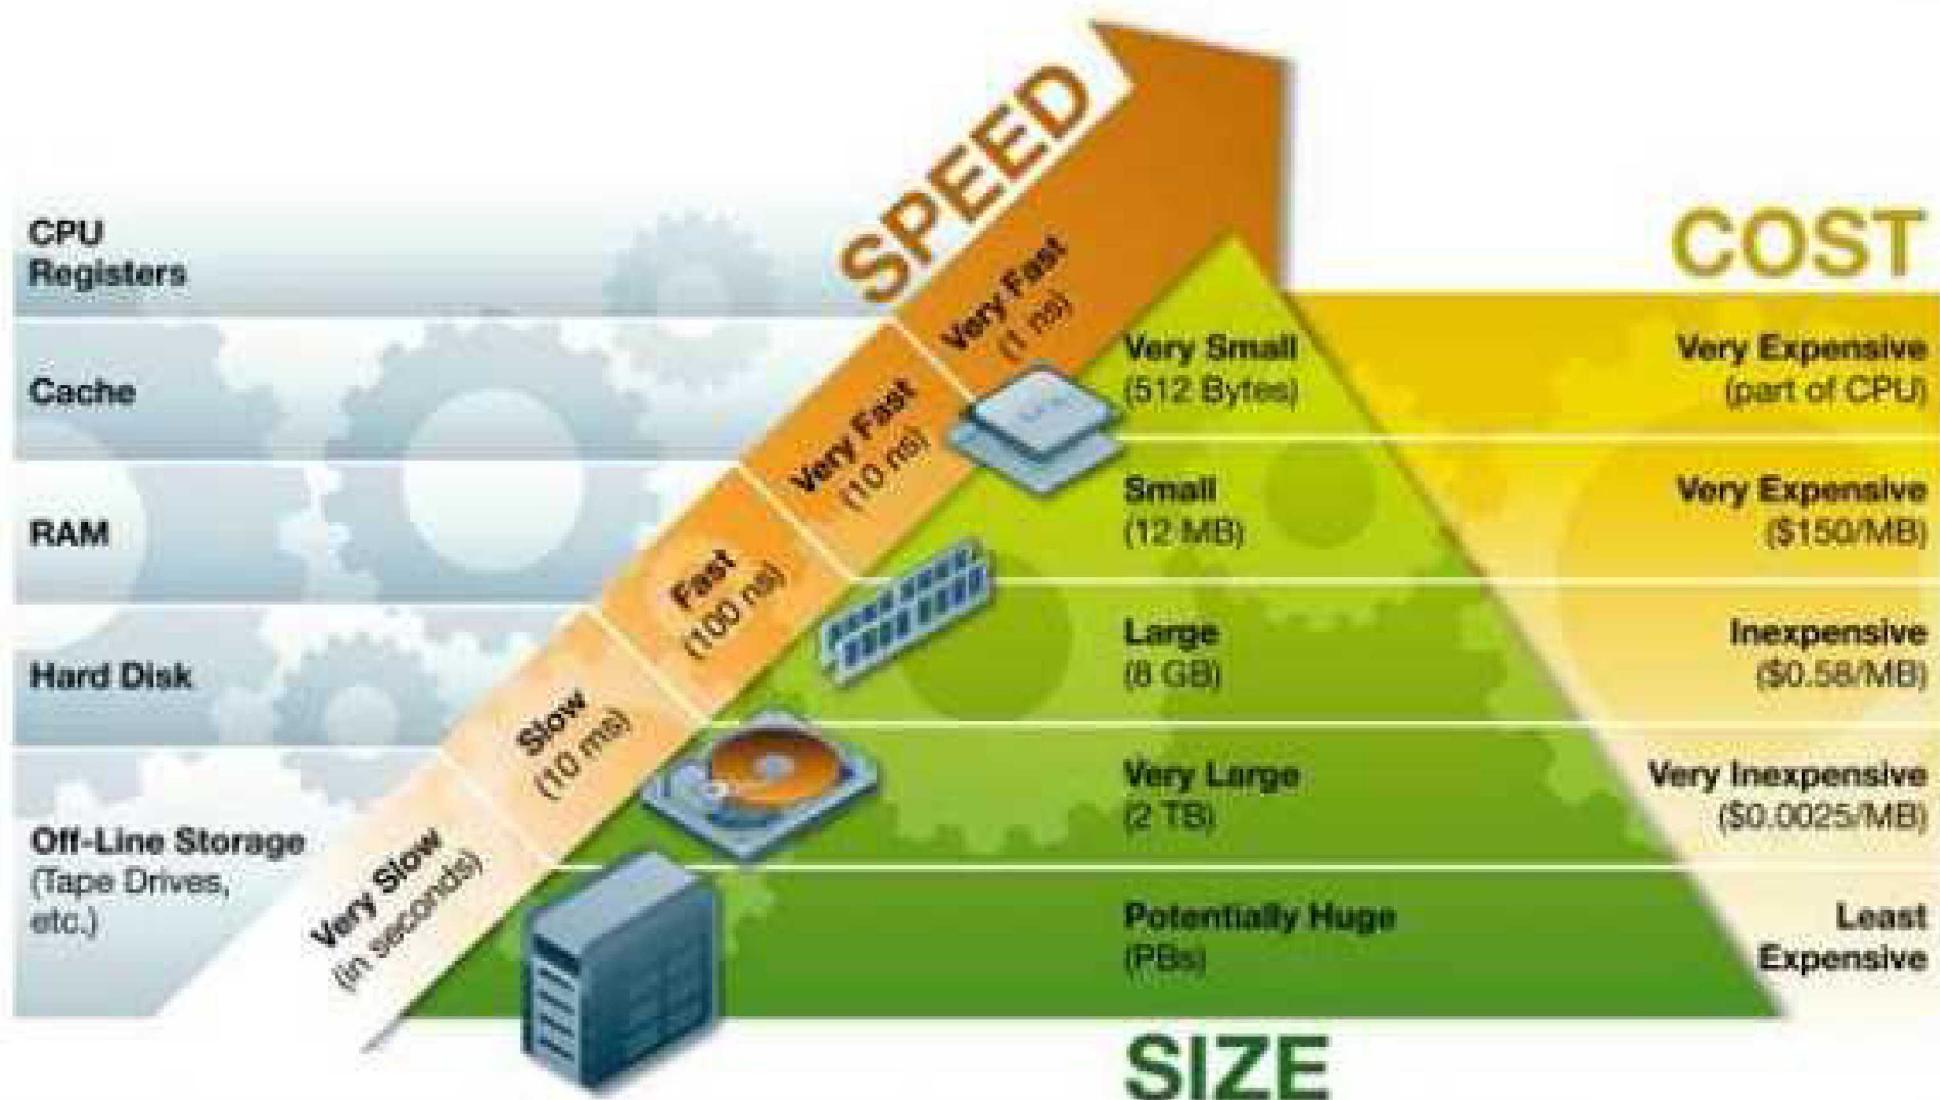
Figure from Elad Elon, 2010 E3S Retreat

Inputs slightly below V_{dd} restored to V_{dd} and inputs slightly above 0 are driven to zero



The Memory Hierarchy

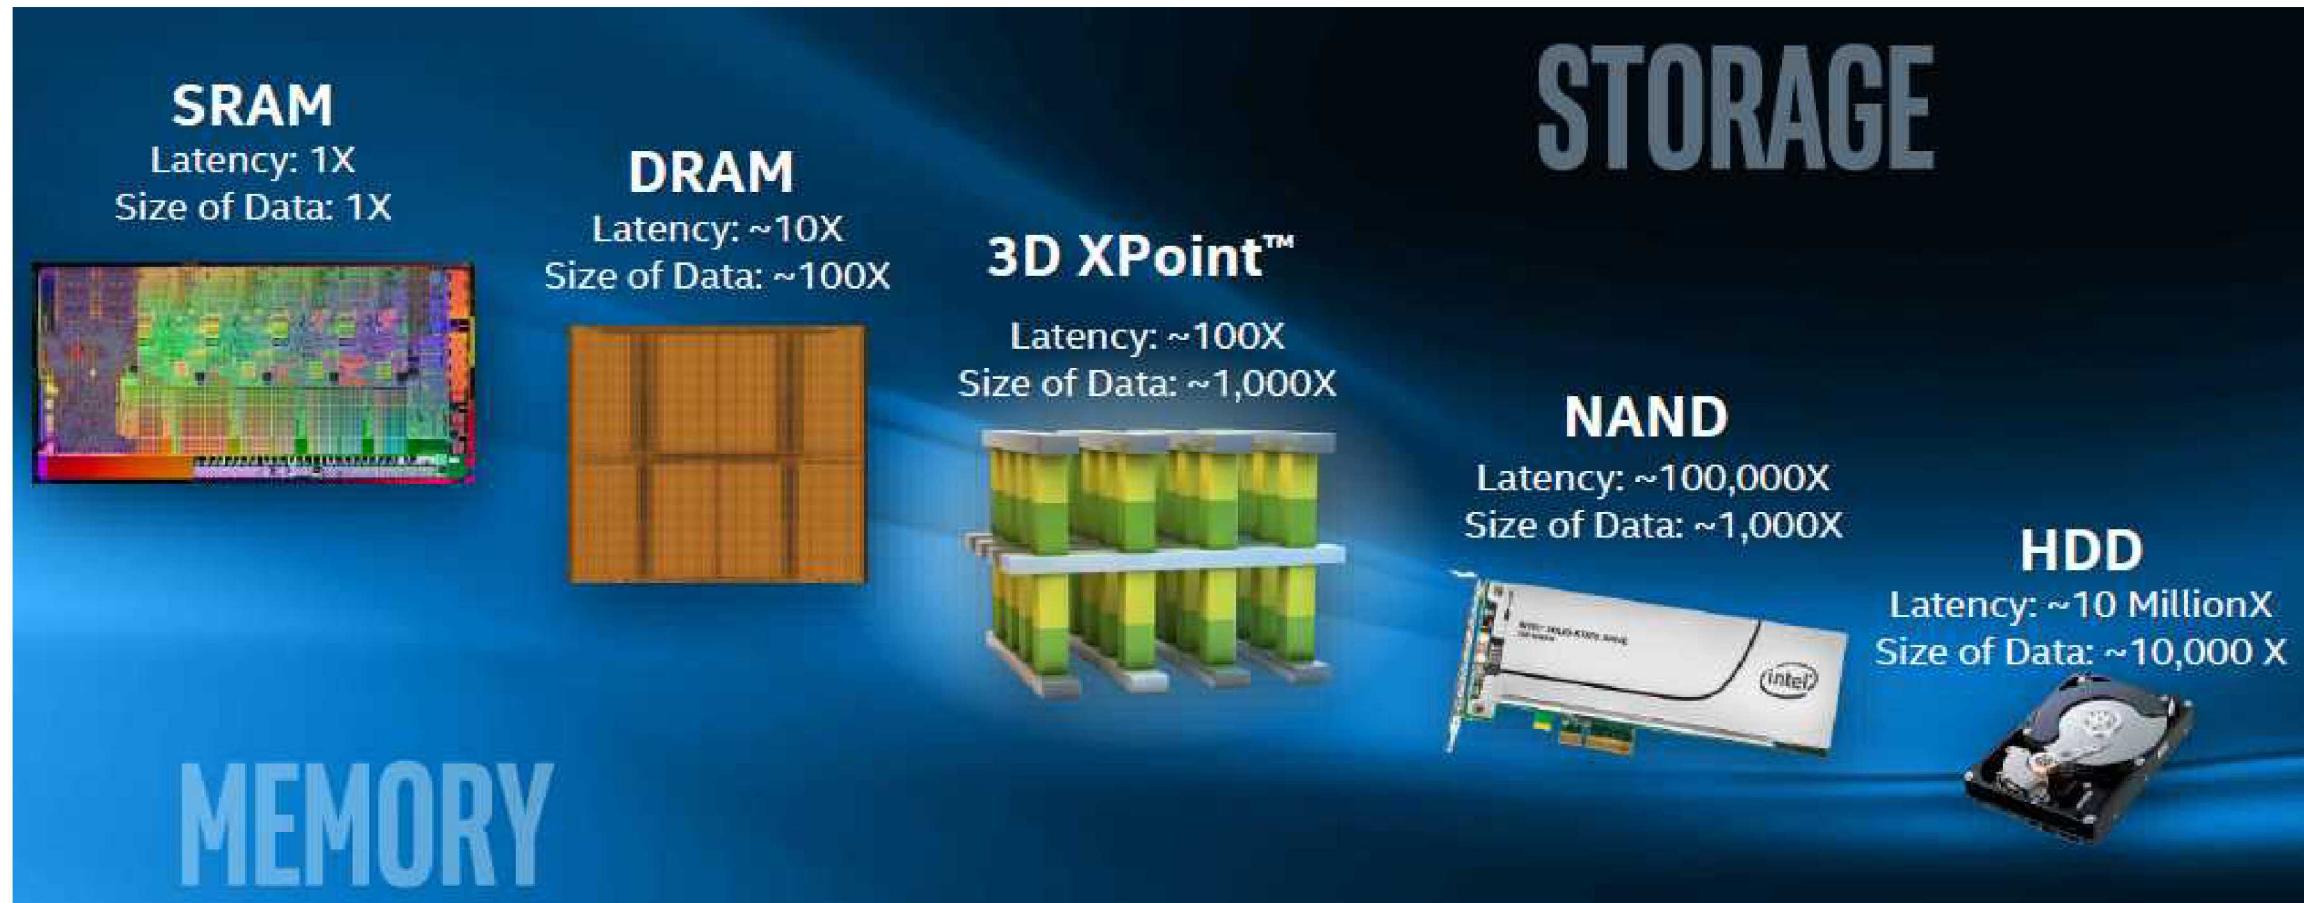
10⁷ difference in speed from HDD to Registers



Source: http://www.ts.avnet.com/uk/products_and_solutions/storage/hierarchy.html

10¹⁰ difference in density from HDD to Registers

Storage Class Memory - Intel/Micron 3D XPoint

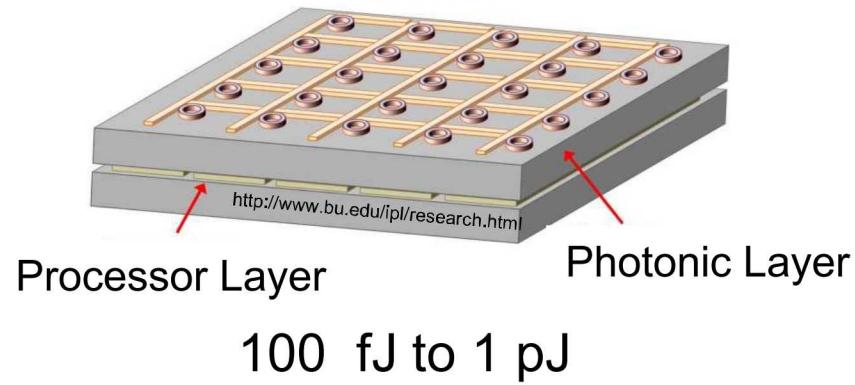


Also Samsung Z-NAND:

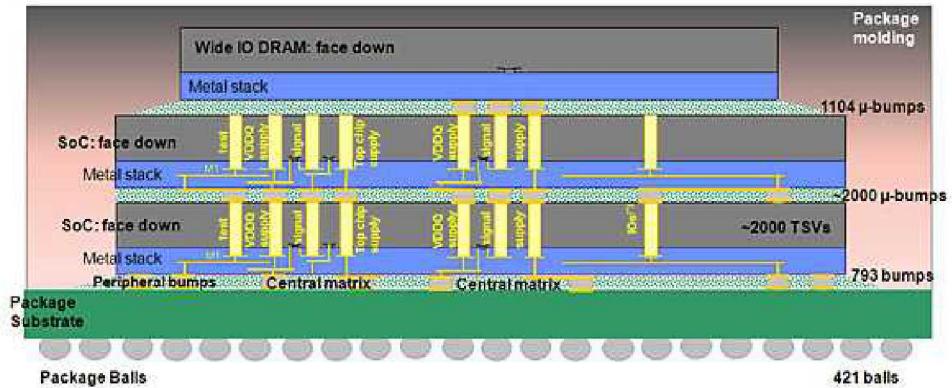
- Re-optimize flash for speed rather than density (single level per cell)
- Comparable to Intel Optane (3D XPoint) products

Extending Von Neumann – Reduced Data Movement

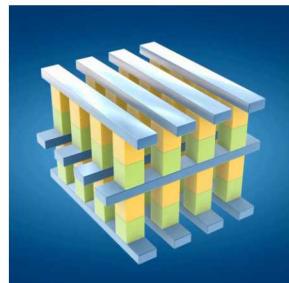
Optical Interconnects



2.5D & 3D Integration



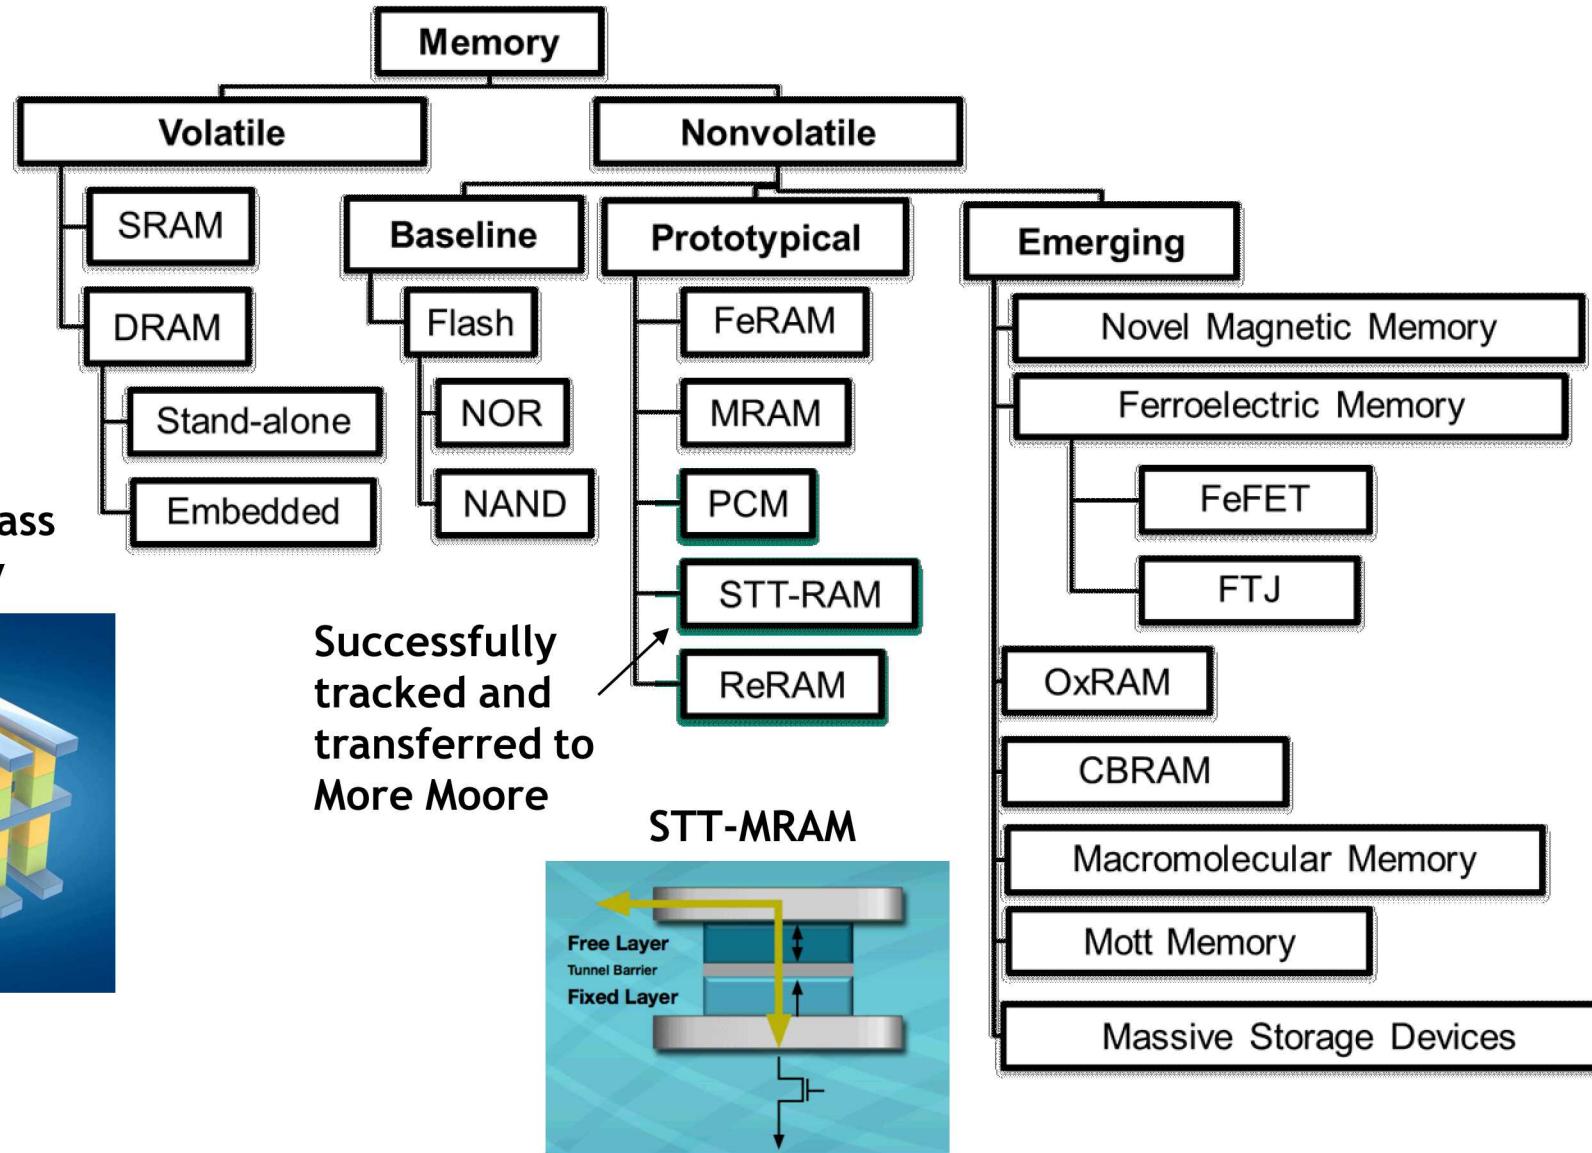
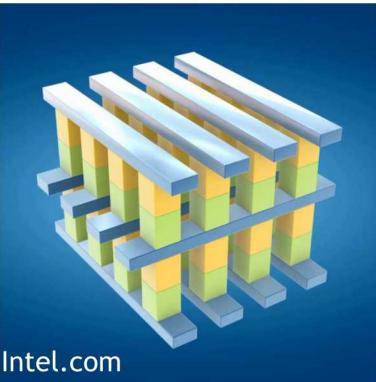
Embedded Memory



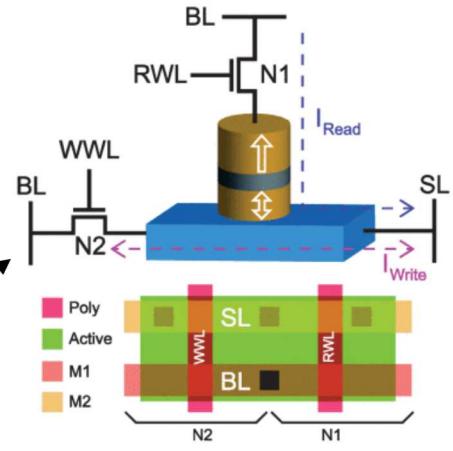
- Embed new denser nonvolatile memories in the processor
- Add simple processing to DRAM or disk drive controllers
“Processing in Memory”

Emerging Memory Devices

Storage Class Memory



SOT Memory



Z. Wang EDL 39, 2018

DNA Memory

STORAGE LIMITS

Estimates based on bacterial genetics suggest that digital DNA could one day rival or exceed today's storage technology.

	Hard disk	Flash memory	Bacterial DNA
Read-write speed (μs per bit)	> 3,000-5,000	~100	<100
Data retention (years)	>10	>10	>100
Power usage (watts per gigabyte)	~0.04	~0.01-0.04	<10 ⁻¹⁰
Data density (bits per cm ³)	~10 ¹³	~10 ¹⁶	~10 ¹⁹

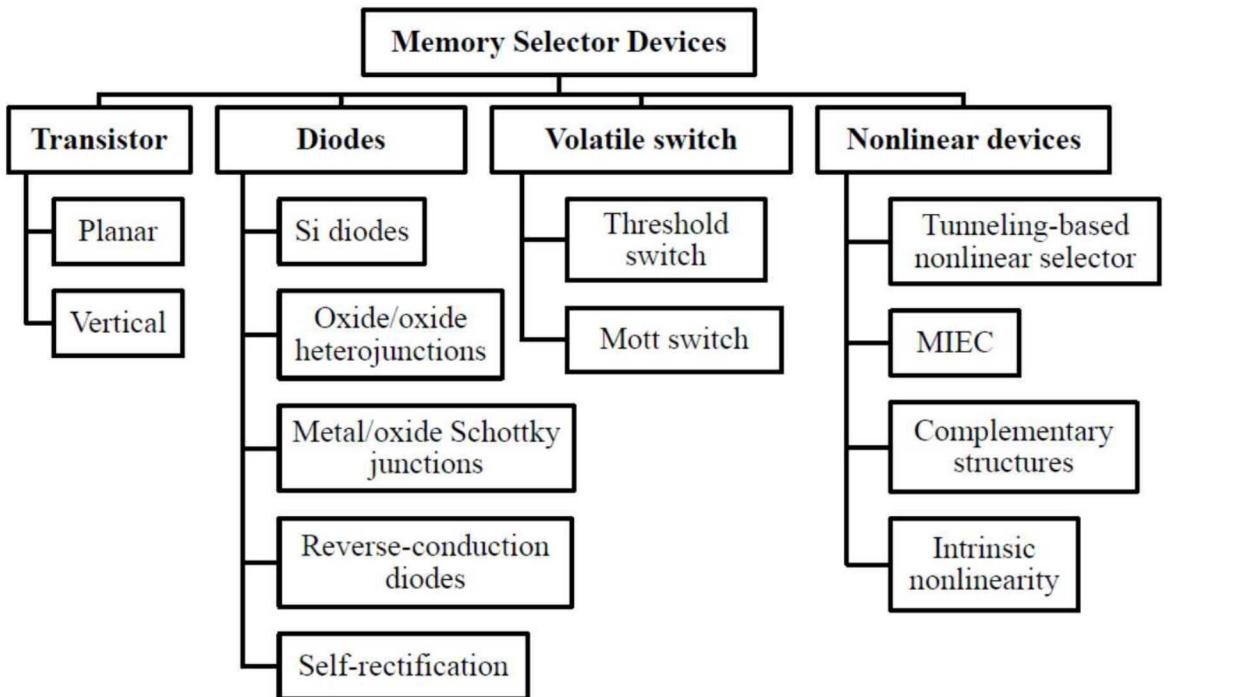
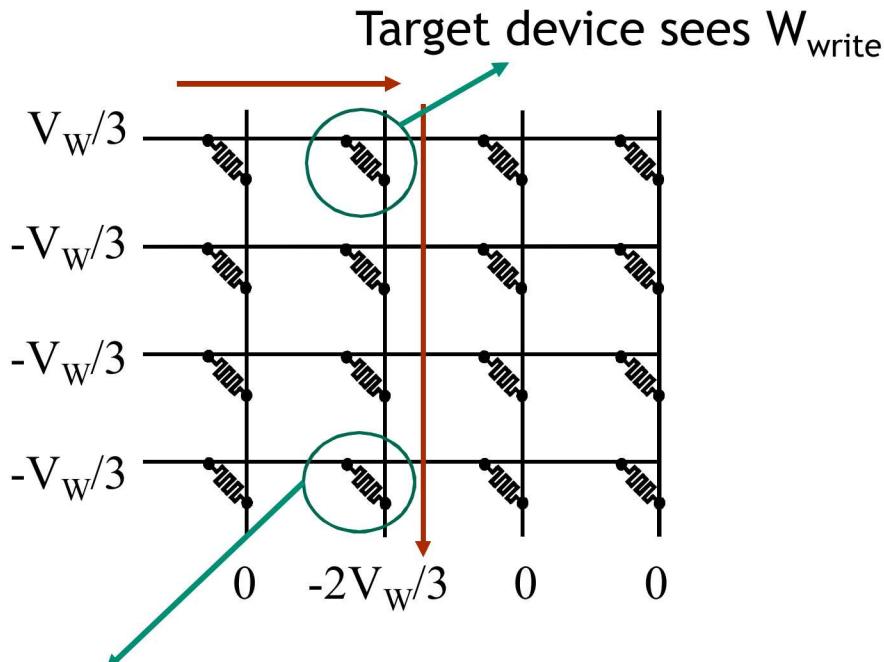
WEIGHT OF DNA NEEDED TO STORE WORLD'S DATA

©nature

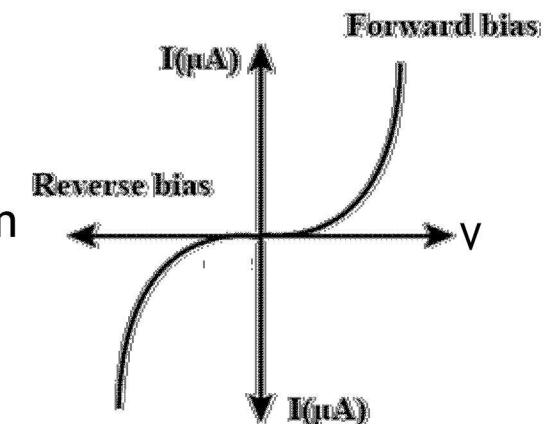
- Back end of line integration
 - Can integrate in the metal layers directly on top of logic
- 3D stackable
- Higher endurance (10^9 to 10^{12}) relative to flash (10^4), but not as good as DRAM ($>10^{16}$)
- Nonvolatile: >10 year retention, no standby leakage
- O(10ns) read and write times

The Need for a Select Device for Resistive Memories

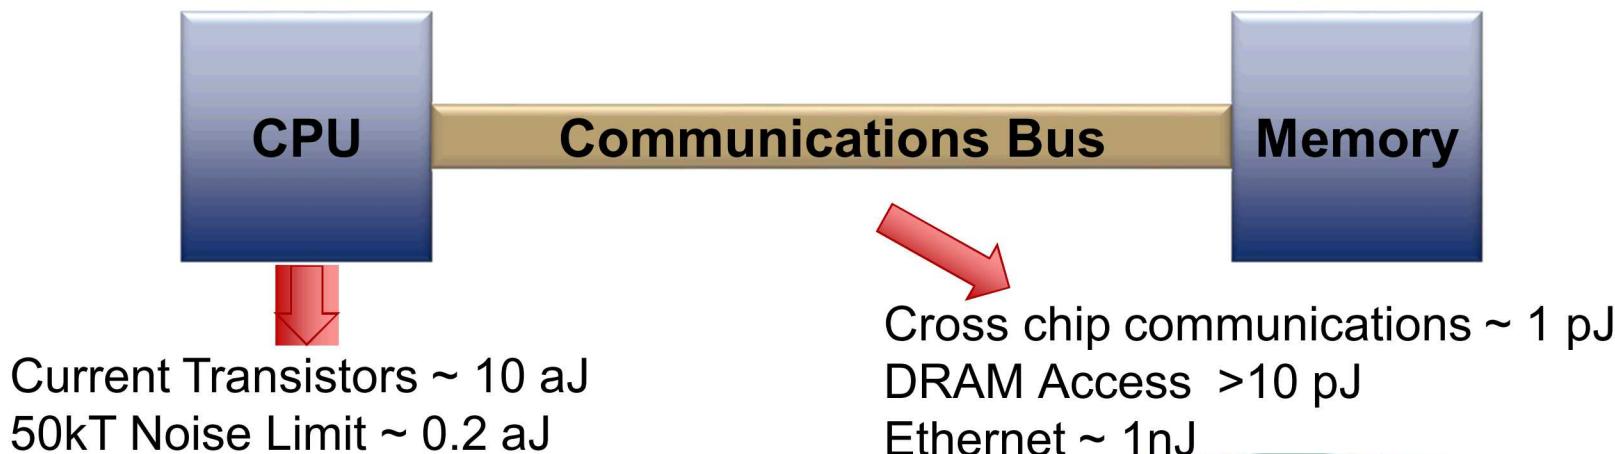
V/3 write scheme



- Need a nonlinear 2 terminal back end of line compatible device to block current from half selected and unselected devices
- The larger the on/off ratio, the larger the possible array and therefore density
- Typically needs to be bi-directional



Beyond Moore Technologies



Extending Von Neumann

- Low Voltage or Novel Transistors
- Optical Communications
- Reduced Data Movement
 - New On-Chip Memory
 - Processing near Memory

Alternate Computing Paradigms

- Neuromorphic
- Analog
- Computing with memory devices
- Quantum
- Stochastic
- Approximate

Going Below 50 kT

- Error Correction
- Reversible Computing
 - Adiabatic Computing / Energy Recycling
- Superconducting

Crossbar Based Computing Architectures

- Vector Matrix Multiplication
- Outer Product Update
- Crossbar Based Matrix Solvers
- Ternary Content Addressable Memory

Neuro-inspired Computing

- Hyperdimensional Computing
- Local Learning Rules
- Spiking Neural Networks

Probabilistic and Stochastic Circuits

Computing With Dynamical Systems

- Simulated Annealing
- Coupled Oscillator based energy minimization

Alternate Computing Paradigms

Crossbar Based Computing Architectures

- Vector Matrix Multiplication
- Outer Product Update
- Crossbar Based Matrix Solvers
- Ternary Content Addressable Memory



Neural Network
Training Accelerators

Neuro-inspired Computing

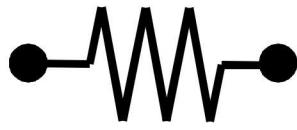
- Hyperdimensional Computing
- Local Learning Rules
- Spiking Neural Networks

Probabilistic and Stochastic Circuits

Computing With Dynamical Systems

- Simulated Annealing
- Coupled Oscillator based energy minimization

Use Resistive Memories for Local Computation



$$V = I \times R$$

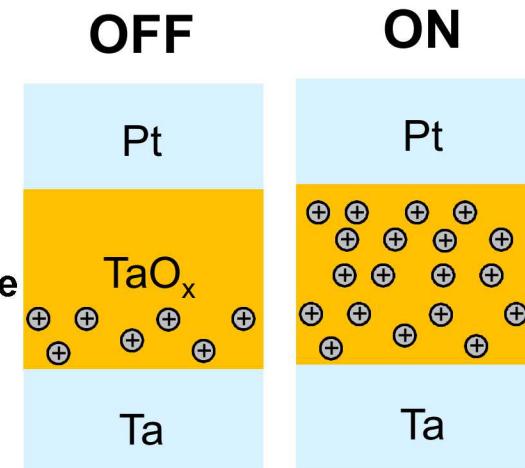
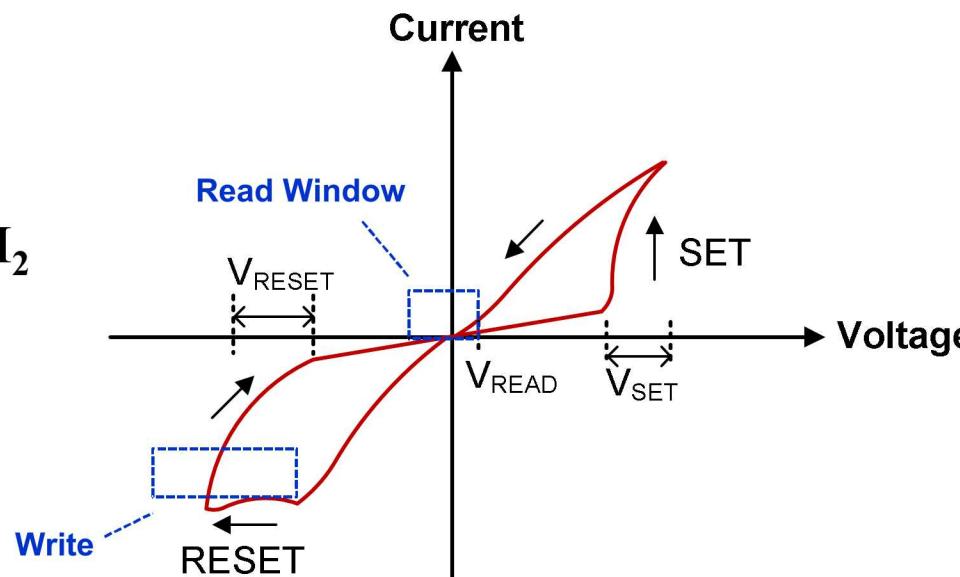
$$I = G \times V$$

multiplication



$$\text{Addition: } I = I_1 + I_2$$

- A resistive memory or ReRAM is a programmable resistor
 - Apply small voltages allows the conductance to be read: $I = G \times V$
 - Apply large voltages to change the resistance



Vector Matrix Multiply: Directly Process in the Memory Itself



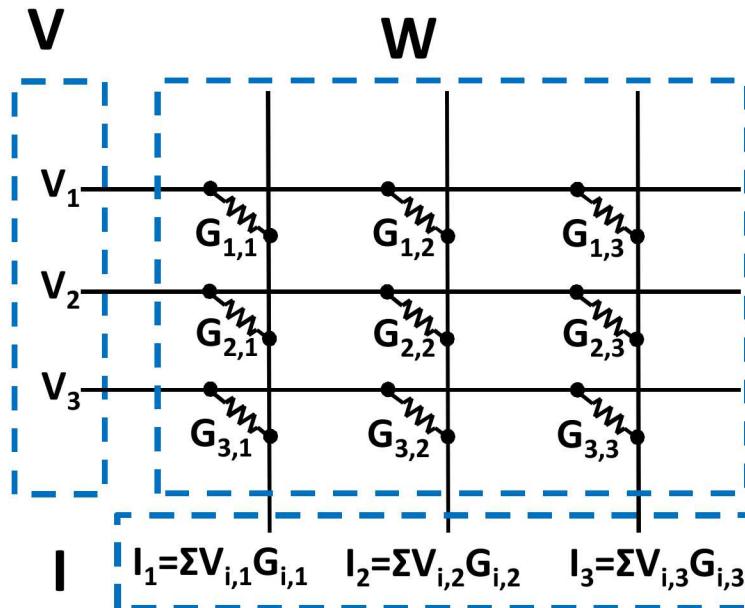
Mathematical

$$V^T W = I$$

$$\begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix} \begin{bmatrix} w_{1,1} & w_{1,2} & w_{1,3} \\ w_{2,1} & w_{2,2} & w_{2,3} \\ w_{3,1} & w_{3,2} & w_{3,3} \end{bmatrix} =$$

$$\begin{bmatrix} I_1 = \sum v_i w_{i,1} & I_2 = \sum v_i w_{i,2} & I_3 = \sum v_i w_{i,3} \end{bmatrix}$$

Electrical



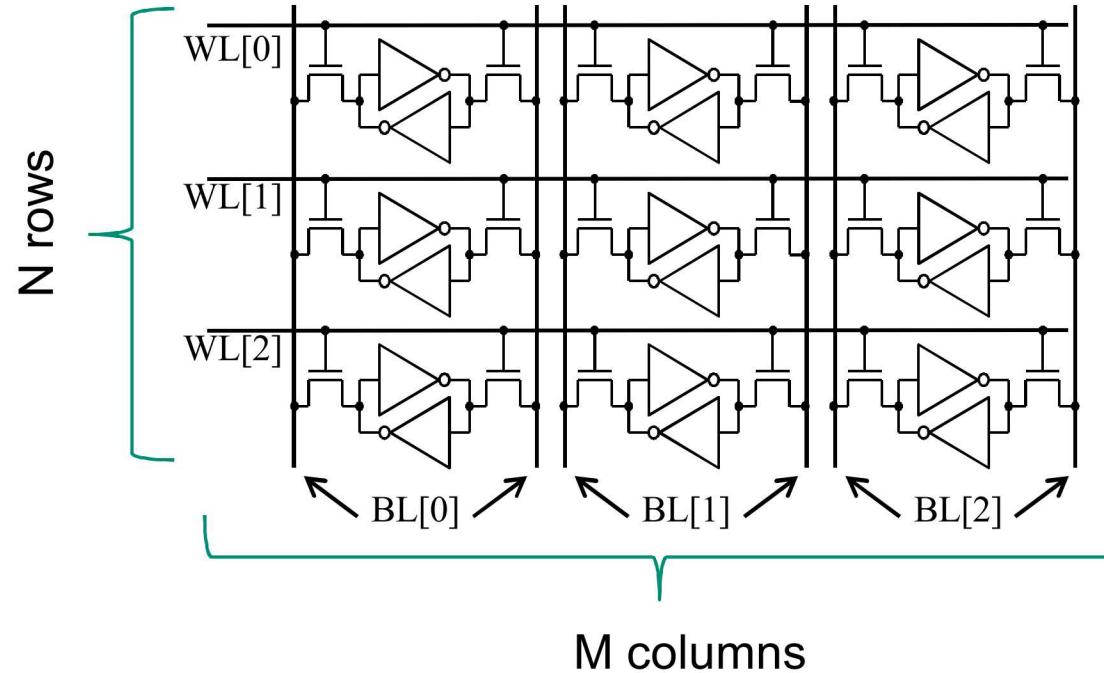
<10 fJ MAC
<10 fJ Update
>100 TOPS/W

Analog is efficiently and naturally able to combine computation and data access

Large-scale processing in memory with a multiplier and adder at each real-valued memory location

- Energy to charge the crossbar is CV^2
- $E \propto C \propto \text{number of RRAMs} \propto N \times M$

SRAM Arrays Require Charging Columns Multiple Times



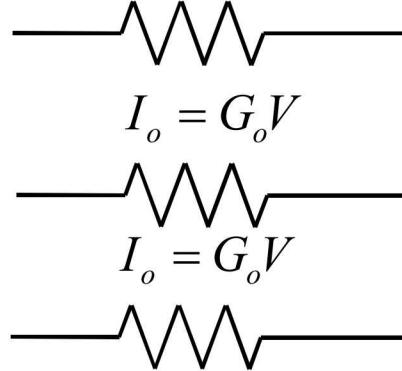
SRAMs must be read one row at a time, charging M columns
Each column wire length is $O(N)$.

$$\text{Energy} = N \text{ Rows} \times M \text{ Columns} \times O(N) \text{ wire length}$$

$$\text{Energy} \sim O(N^2 \times M)$$

$O(N)$ times worse than a crossbar!

The Noise Limited Energy to Read a Crossbar Column is Independent of Crossbar Size

$$I_o = G_o V$$


$$\text{Thermal Noise} = \langle \Delta I^2 \rangle$$
$$= N \times (4k_b T \times G_o \times \Delta f)$$

$$SNR^2 = \frac{(NI_o)^2}{\langle \Delta I^2 \rangle}$$

$$\frac{1}{\Delta f} = 4k_b T \times SNR^2 \times \frac{1}{V^2 G_o \times N}$$

Measure N resistors and determine the total output current with some signal to noise ratio (SNR)*

What is the minimum energy?

$$Energy = \underbrace{V^2 G_o \times N}_{\text{Power in each resistor} \times \text{number of resistors}} \times \frac{1}{\Delta f}$$

Power in each resistor \times
number of resistors

Determined by
noise and SNR

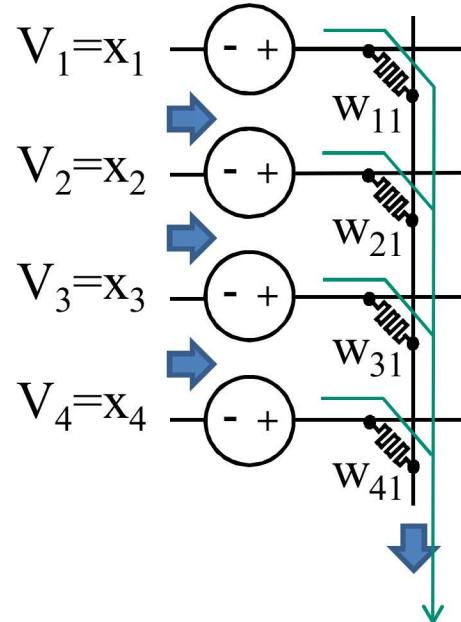
If we double the number of resistors, we can double the speed to get the same energy and SNR.

This is because the noise scales as \sqrt{N} while the signal scales as N

$$Energy = 4k_b T \times SNR^2$$

*we are assuming we need some fixed precision on the output, and don't need full floating point accuracy

Need to Use Analog to Efficiently Discard Precision



Sum 1024 8 bit weights X 8 bit inputs:

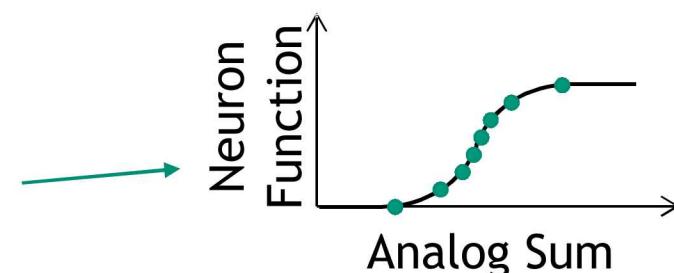
- Result has 26 bits of information!
- A 26 bit ADC would eliminate any analog advantage!

The sum can be done at full precision in analog, but a lower precision approximation is needed when digitizing

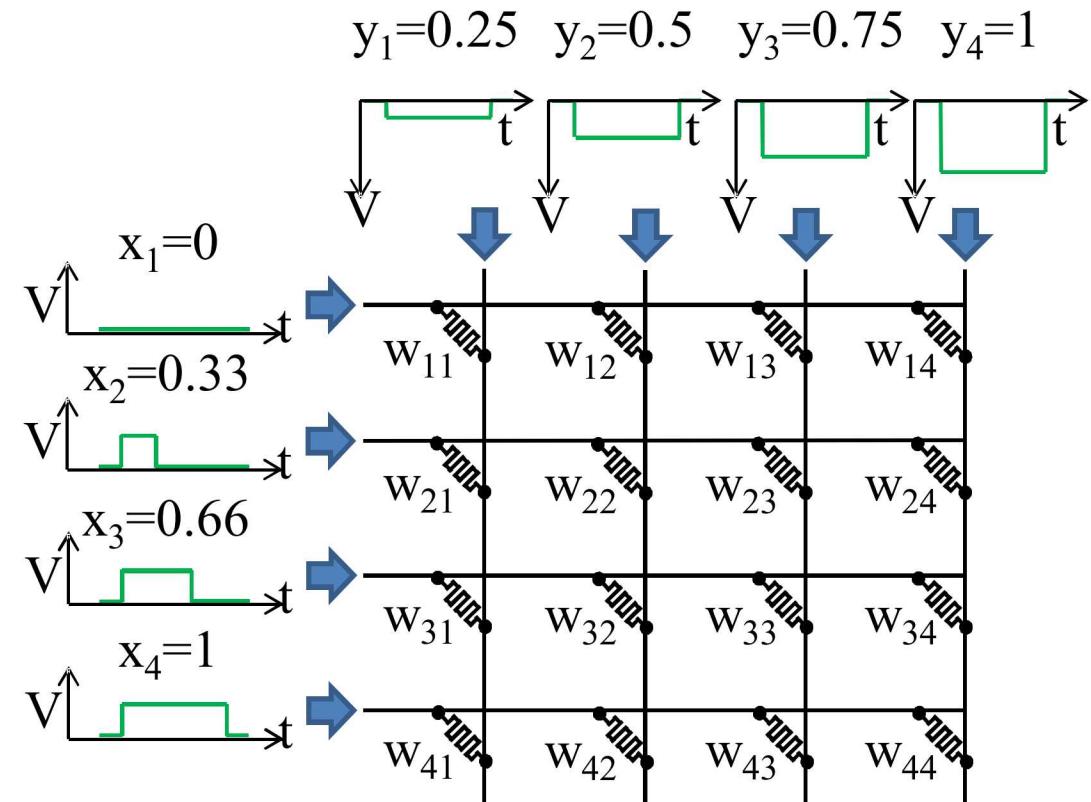
- i.e. digitize only 8 bits or fewer

To get the highest 8 bits of information, digital would need to keep a 26 bit intermediate result

Can design an ADC to choose non uniform values to digitize



Outer Product Update: Parallel Write

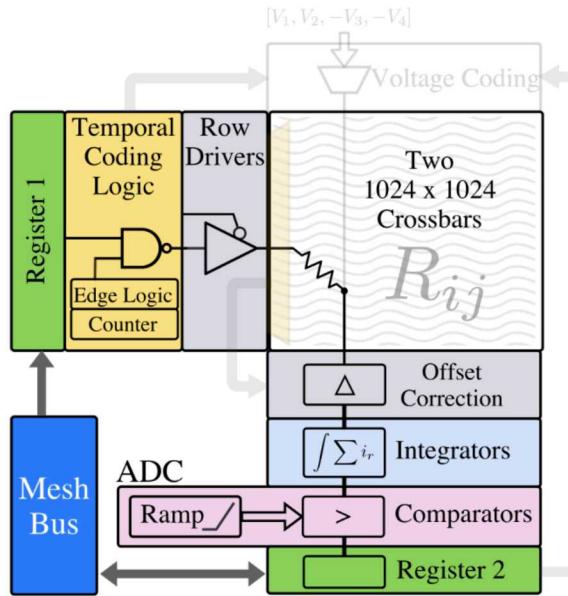


Energy to charge the crossbar is CV^2
 $E \propto C \propto \text{number of RRAMs} \propto N \times M$

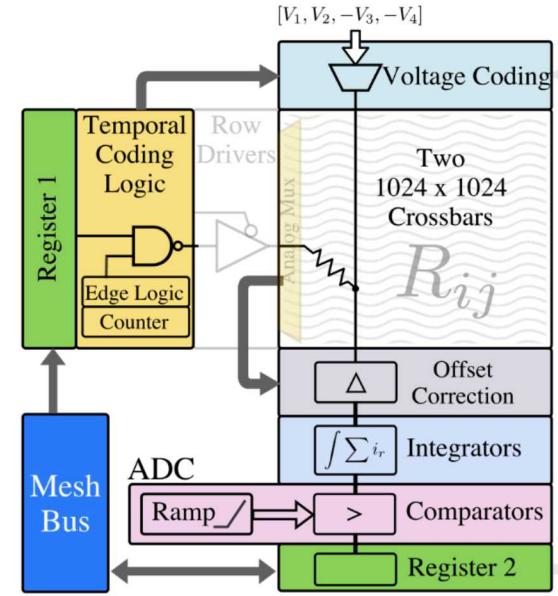
$$E \sim O(N \times M)$$

Example: Design an Neural Network Training Accelerator

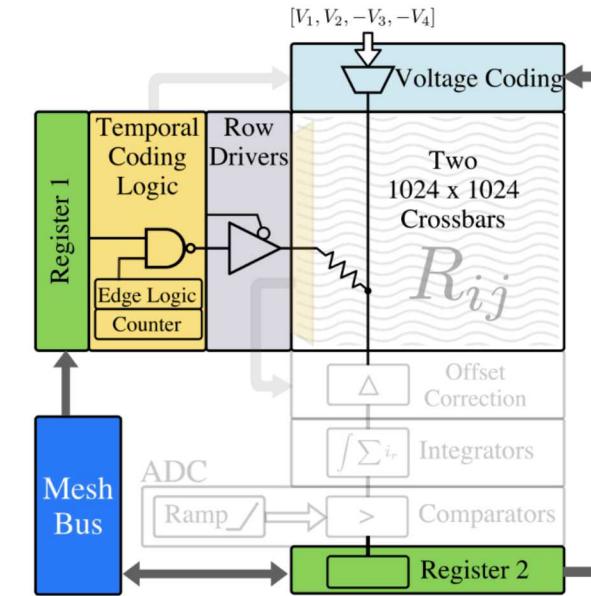
Vector Matrix Multiply



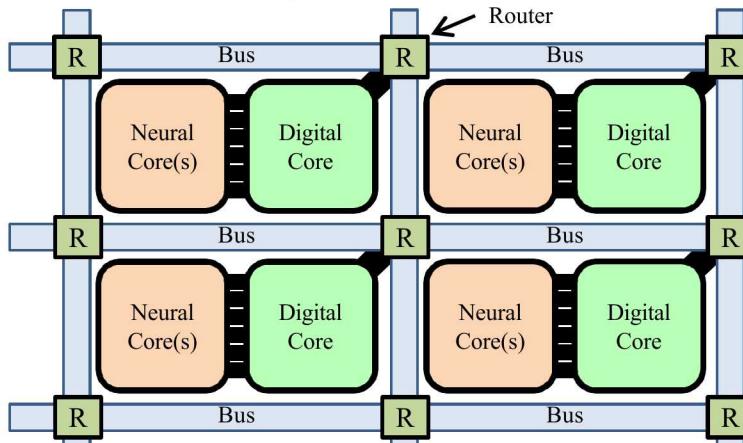
Matrix Vector Multiply



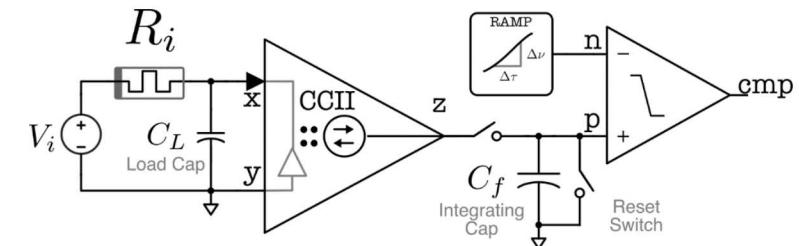
Outer product Update



Chip Architecture



Neuron Circuitry

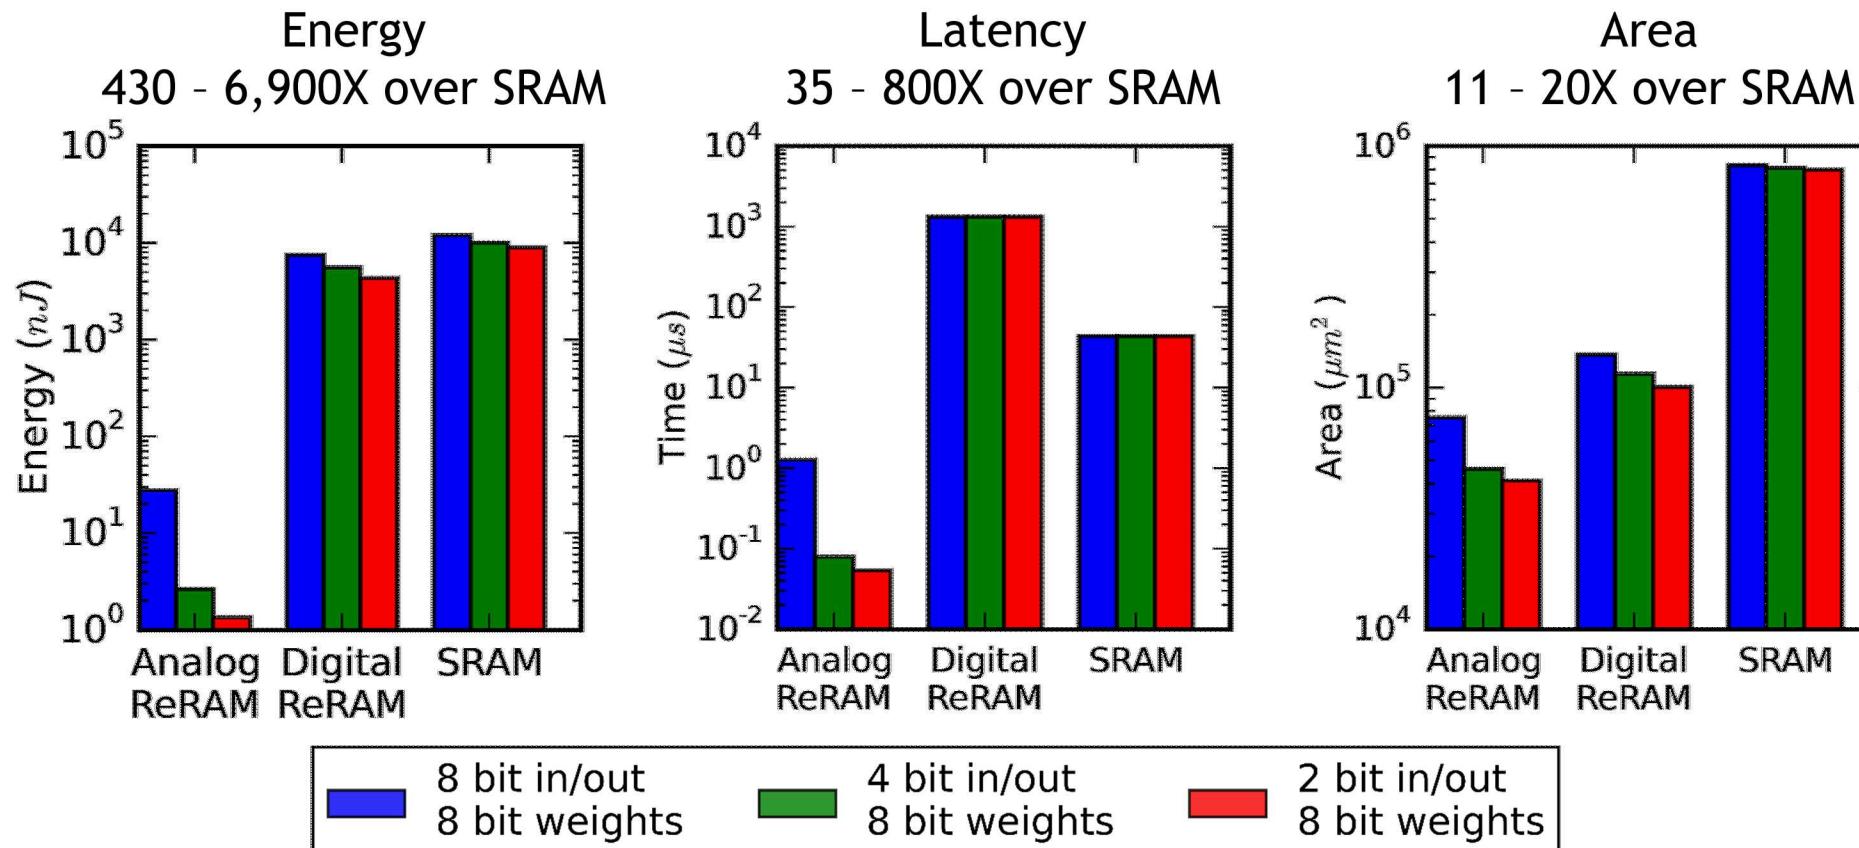


Energy Area and Latency Advantages of an Analog Accelerator

29

$1024 \times 1024 = 1M$ array operations, sum over 1 training cycle, 3 operations:

- Vector Matrix Multiply
- Matrix Vector Multiply
- Outer Product Update

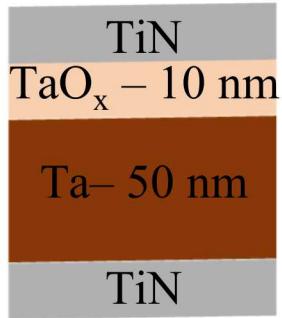


Used a commercial 14/16 nm PDK

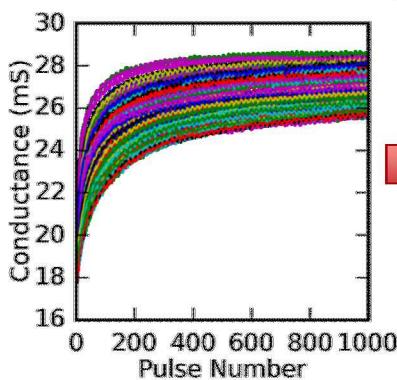
***Requires 100 $M\Omega$ on state devices

Go from Measurement to Accuracy

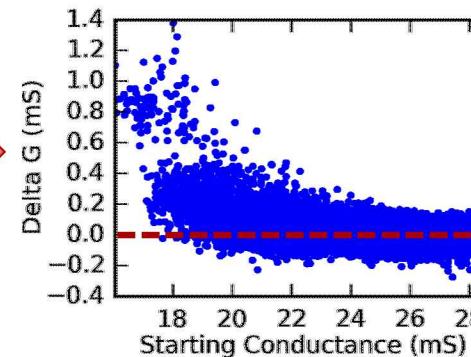
Fabricate Device



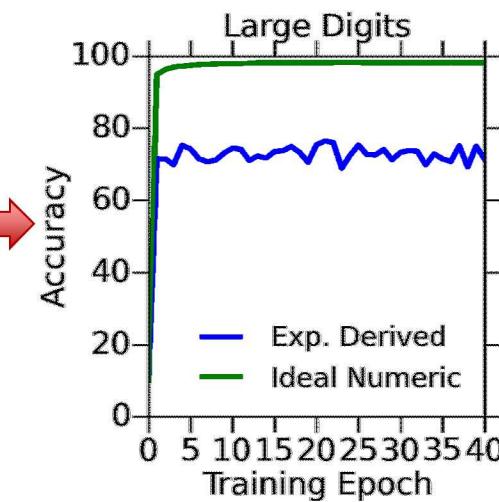
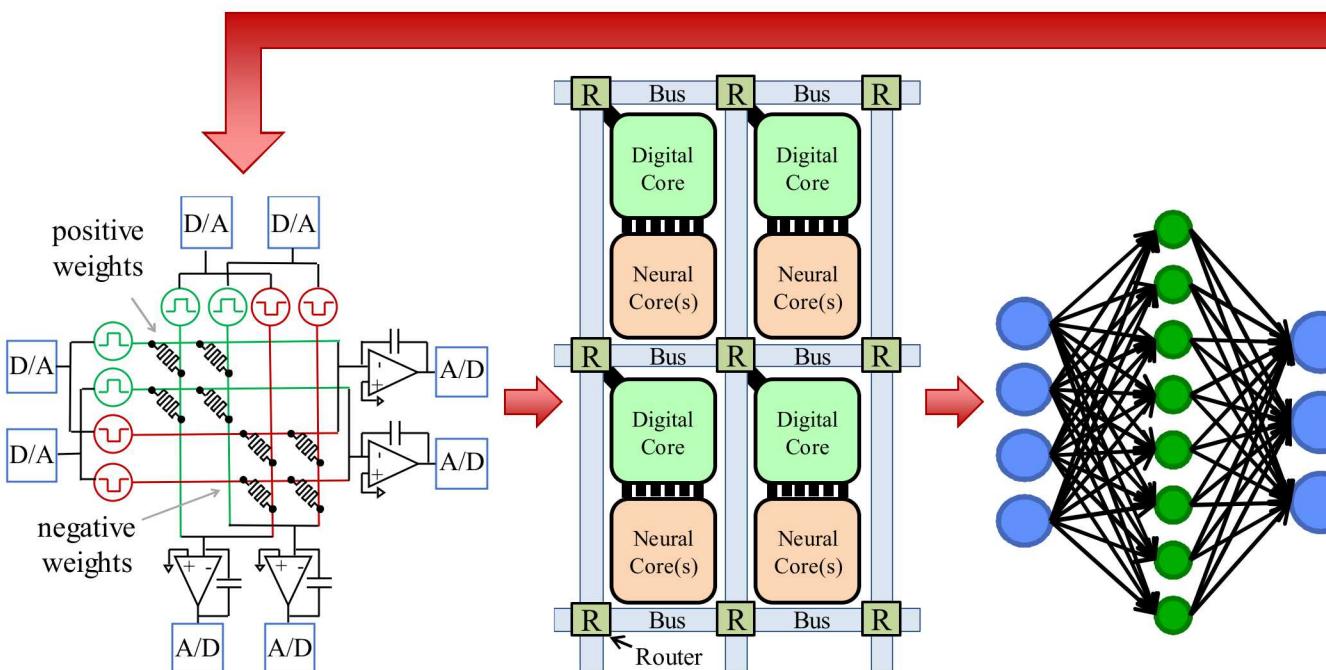
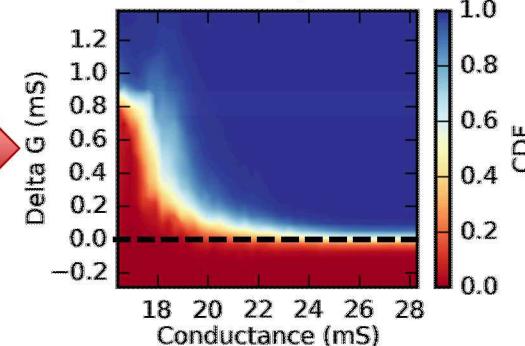
Measured Pulsing



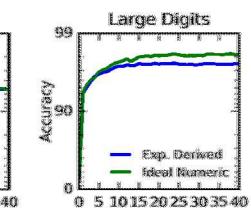
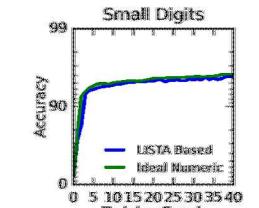
ΔG Scatterplot



Cumulative Probability of ΔG



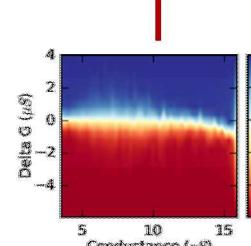
Multiscale Model of a Neural Training Accelerator



Target Algorithms

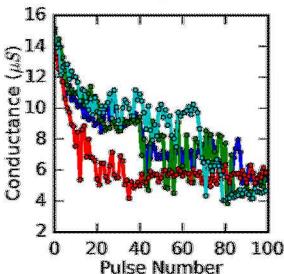
- Deep Learning
- Sparse Coding
- Liquid State Machines

Algorithms

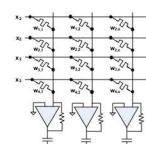


#ROSS SIM

Sandia Cross-Sim:
Translates device measurements and crossbar circuits to algorithm-level performance



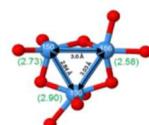
Memristor fabrication and measurements in MESAFab



Circuits

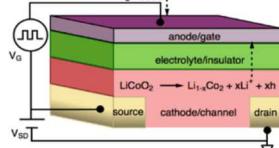


Sandia's Xyce Circuit Sim: Simulate crossbar circuits based on our devices

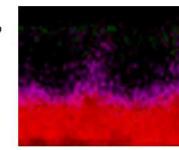
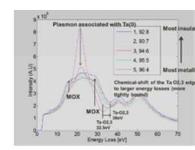
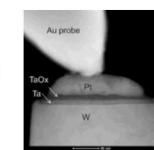
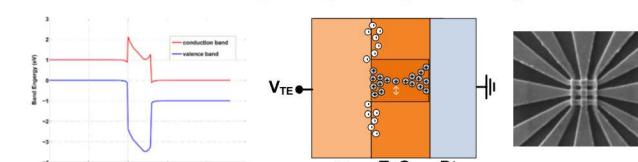


DFT of model of oxide physics, bands

Devices

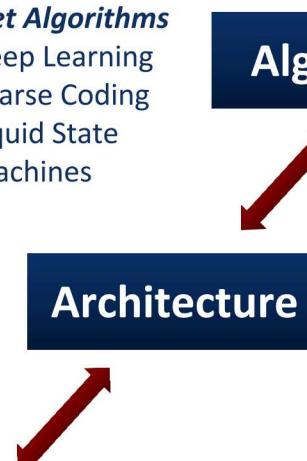


Drift-diffusion model of ReRAM band diagram & transport (REOS, Charon)

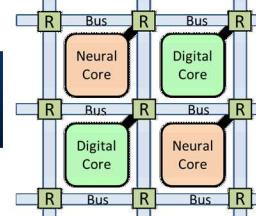


Materials

In situ TEM of filament switching: Use DFT model to interpret EELS signature

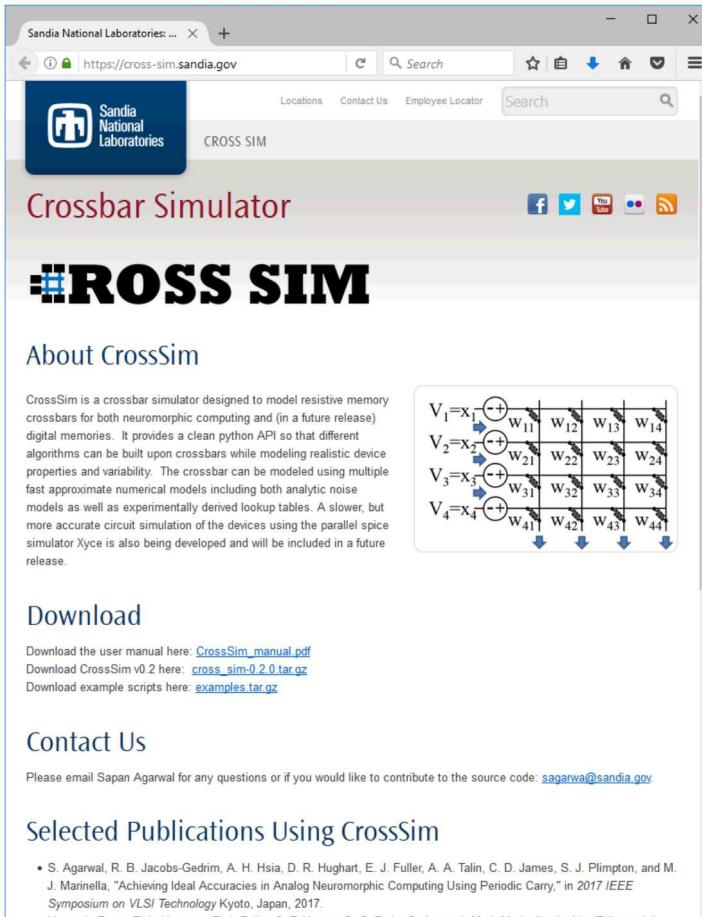


Modified McPAT/CACTI:
Model performance and energy requirements



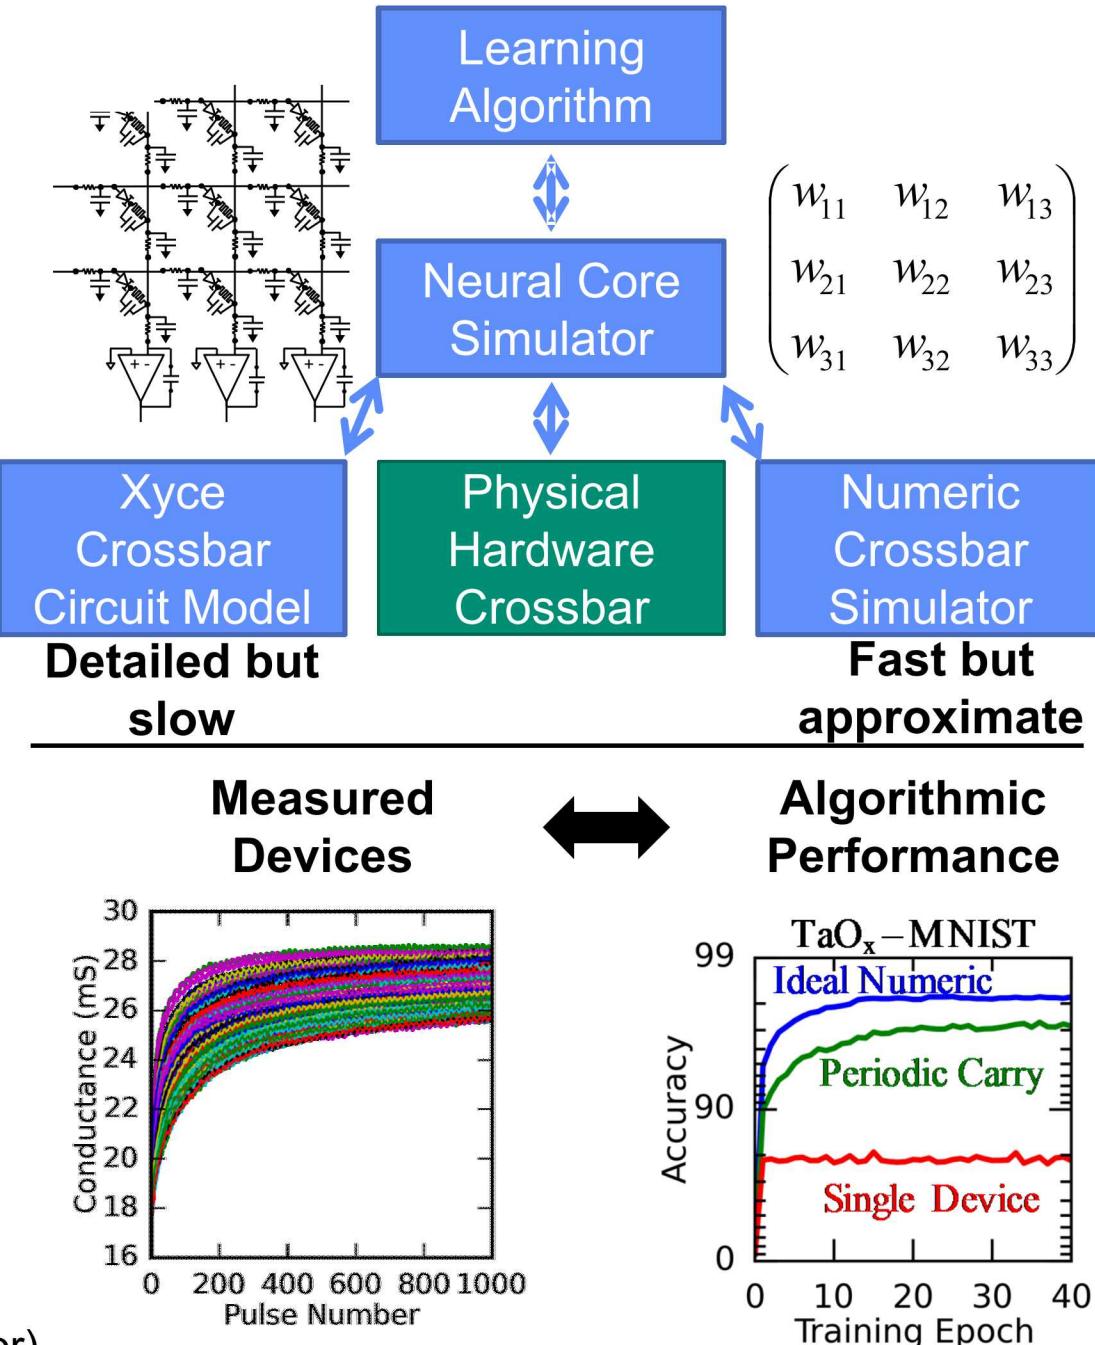
CROSS SIM

<https://cross-sim.sandia.gov>



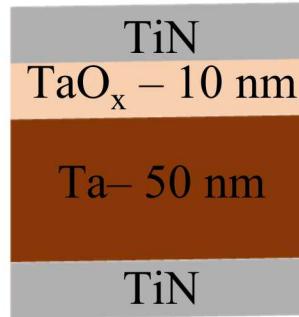
Simple Python API:

Do a matrix vector multiplication
result = neural_core.run_xbar_mvm(vector)

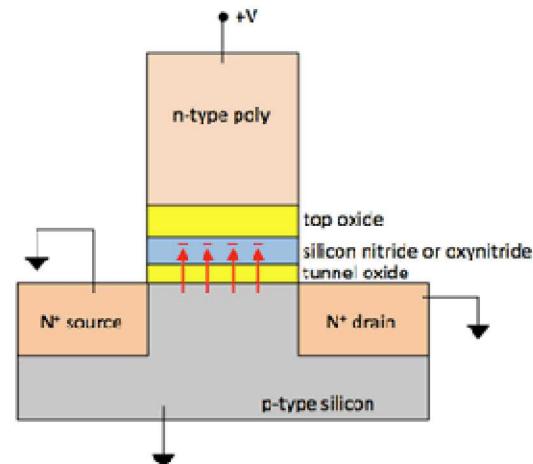


Compare Analog Devices

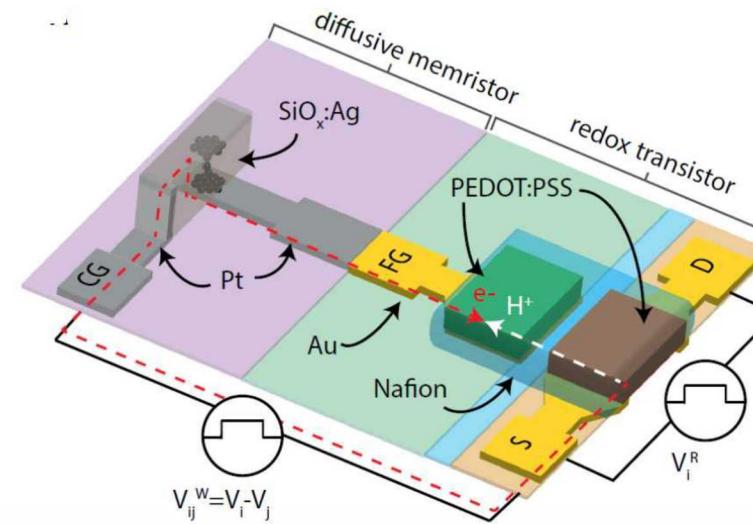
ReRAM



SONOS Silicon-Oxygen- Nitrogen-Oxygen-Silicon



Ionic Floating-Gate Memory

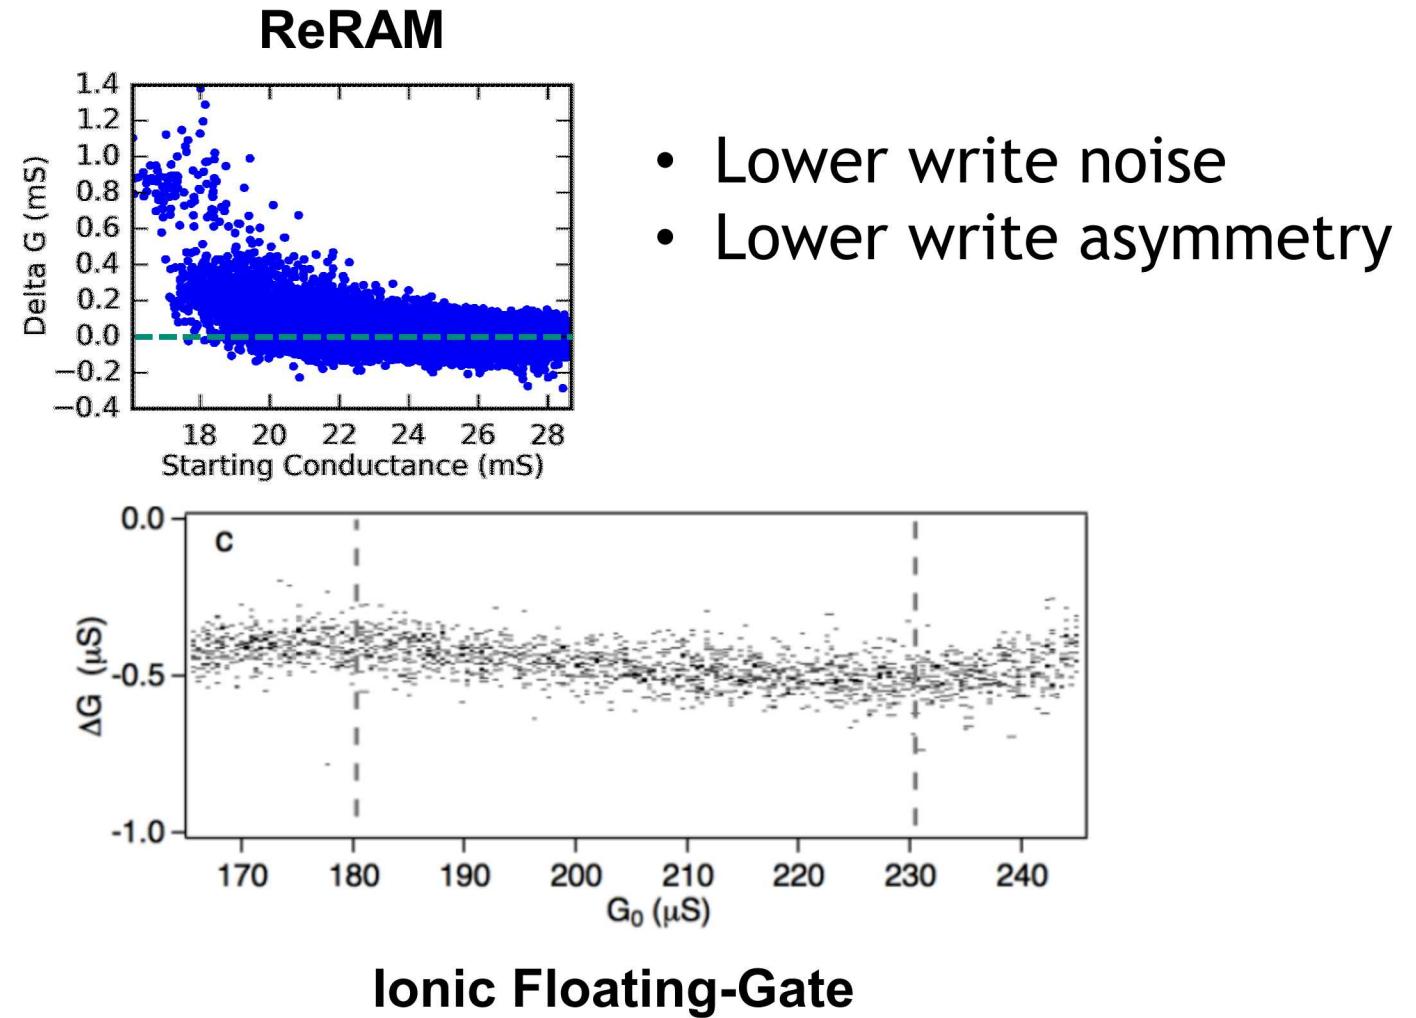
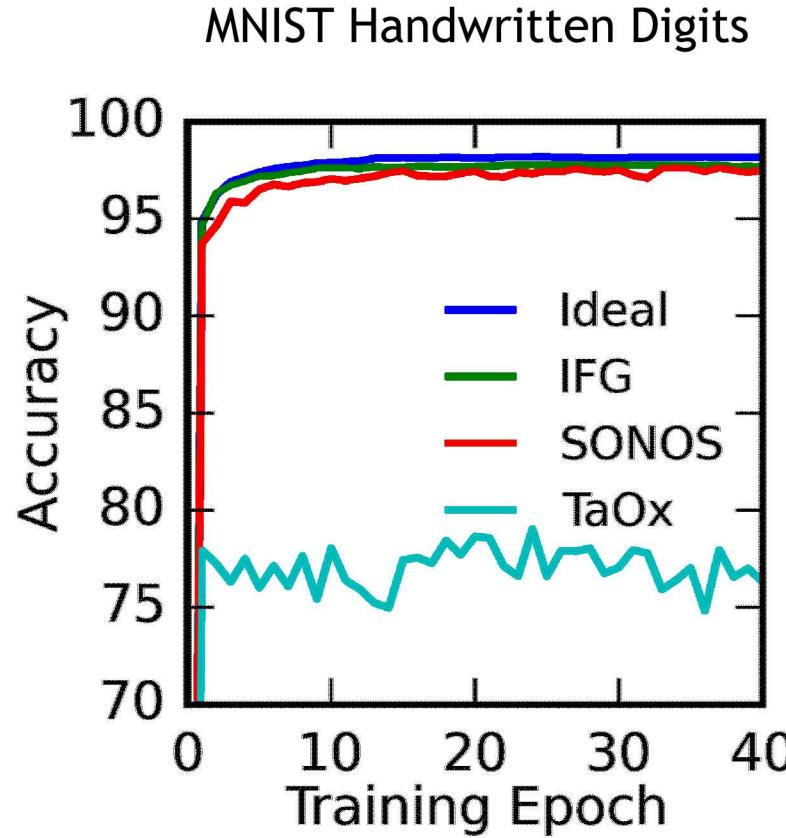


R. B. Jacobs-Gedrim *et al.*, "Impact of Linearity and Write Noise of Analog Resistive Memory Devices in a Neural Algorithm Accelerator," IEEE International Conference on Rebooting Computing (ICRC) Washington, DC, November 2017.

S. Agarwal *et al.*, "Using Floating Gate Memory to Train Ideal Accuracy Neural Networks," *IEEE Journal of Exploratory Solid-State Computational Devices and Circuits*, 2019

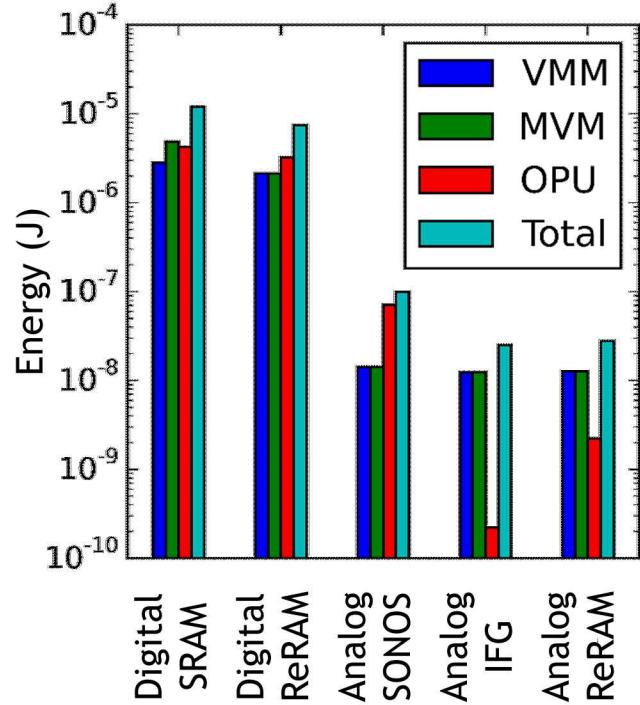
E. J. Fuller *et al.*, "Parallel programming of an ionic floating-gate memory array for scalable neuromorphic computing," *Science*, vol. 364, no. 6440, pp. 570-574, 2019.

Three Terminal Devices Tend to Have Higher Accuracy

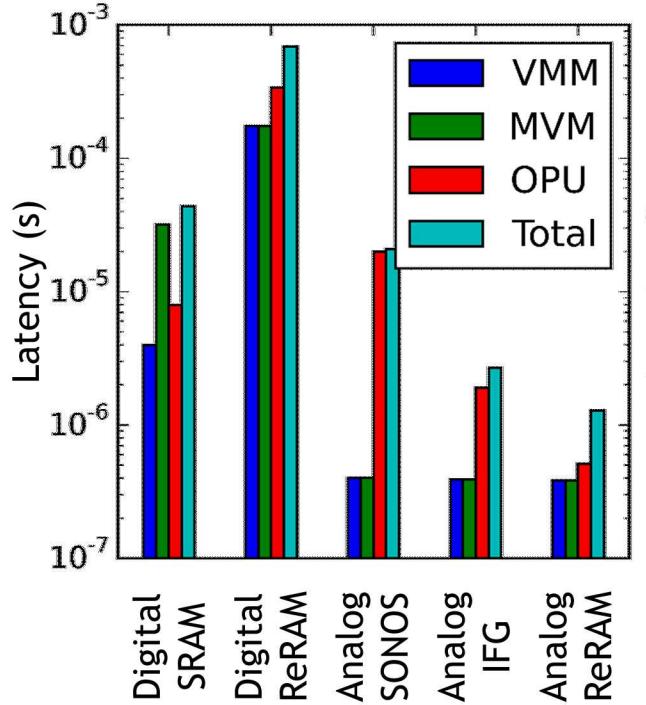


Compare Architectural Advantages

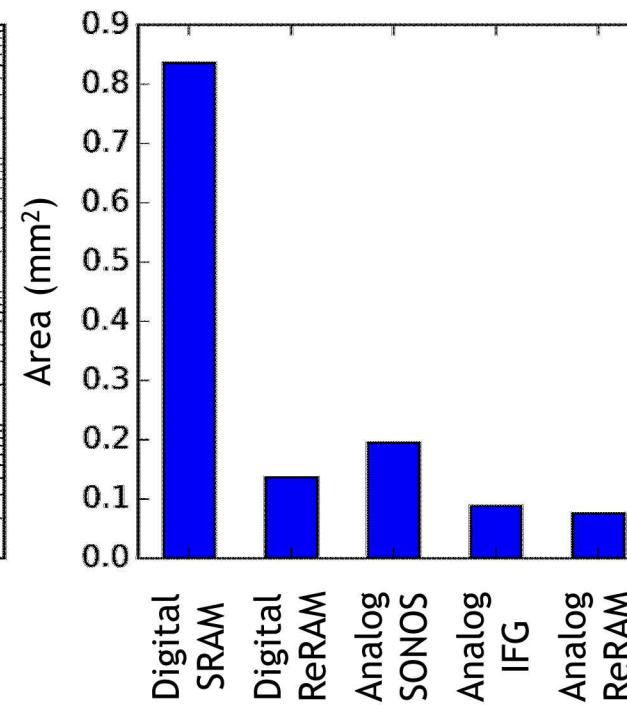
120-430X Energy Advantage



2-34X Latency Advantage



5-11X Area Advantage

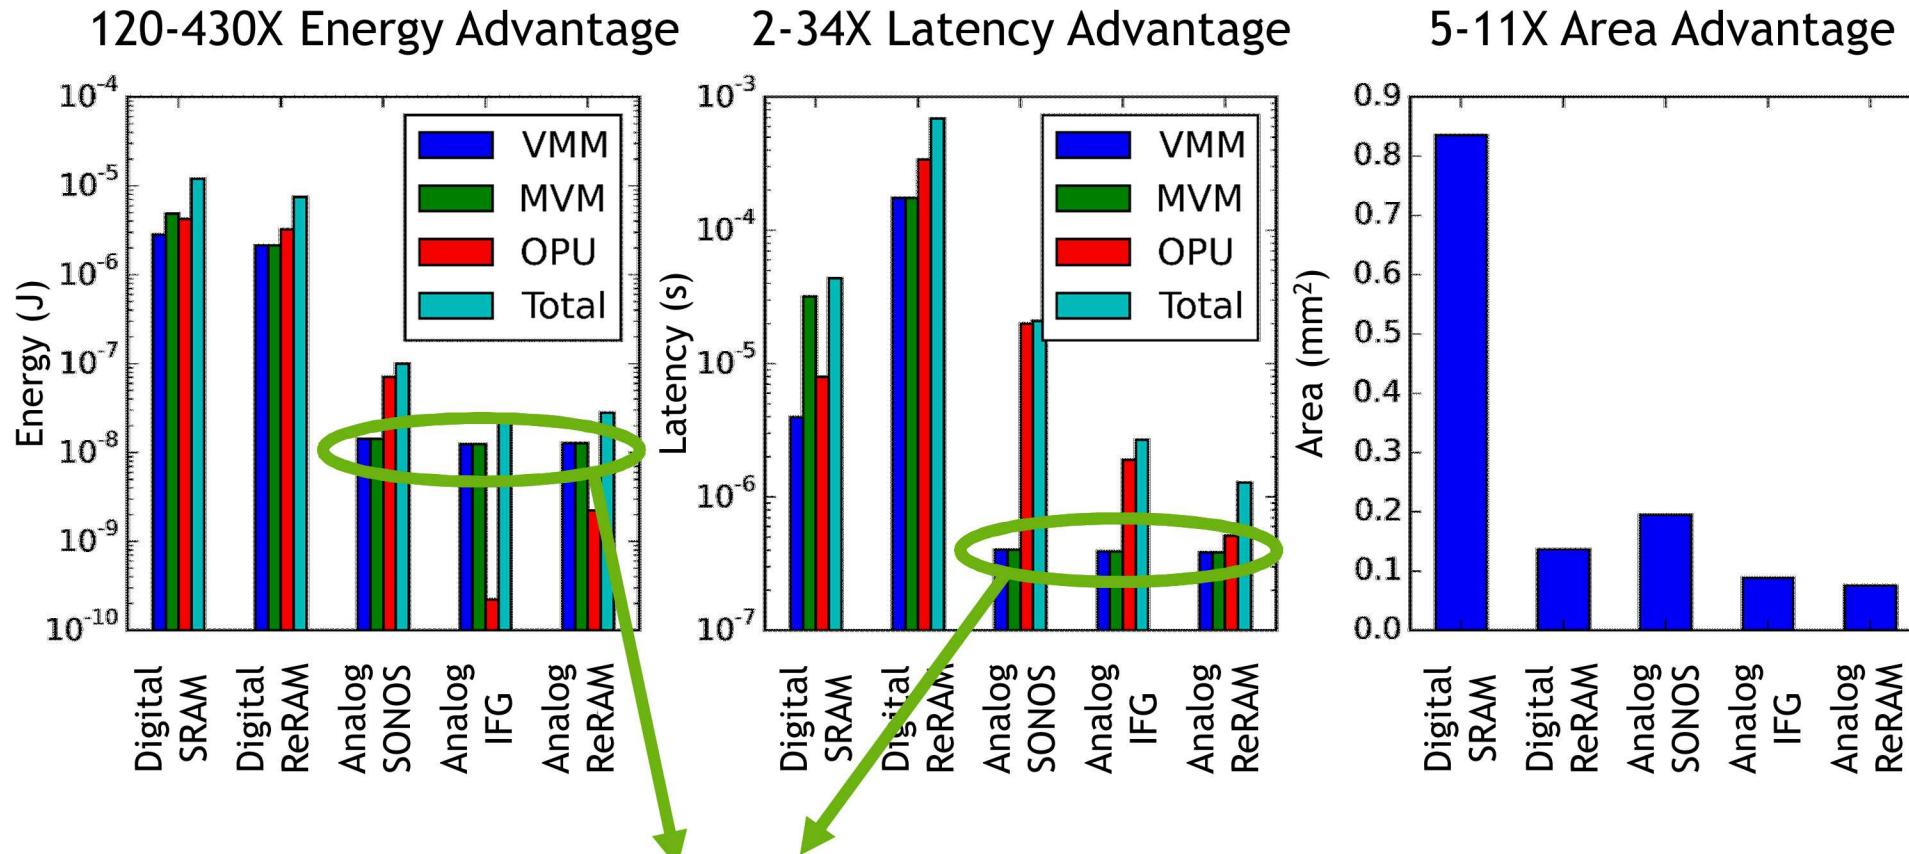


$1024 \times 1024 = 1M$ array operations, sum over 1 training cycle, 3 operations:
- Vector Matrix Multiply - Matrix Vector Multiply - Outer Product Update

Used a commercial 14/16 nm PDK

***Requires 100 MΩ on state devices

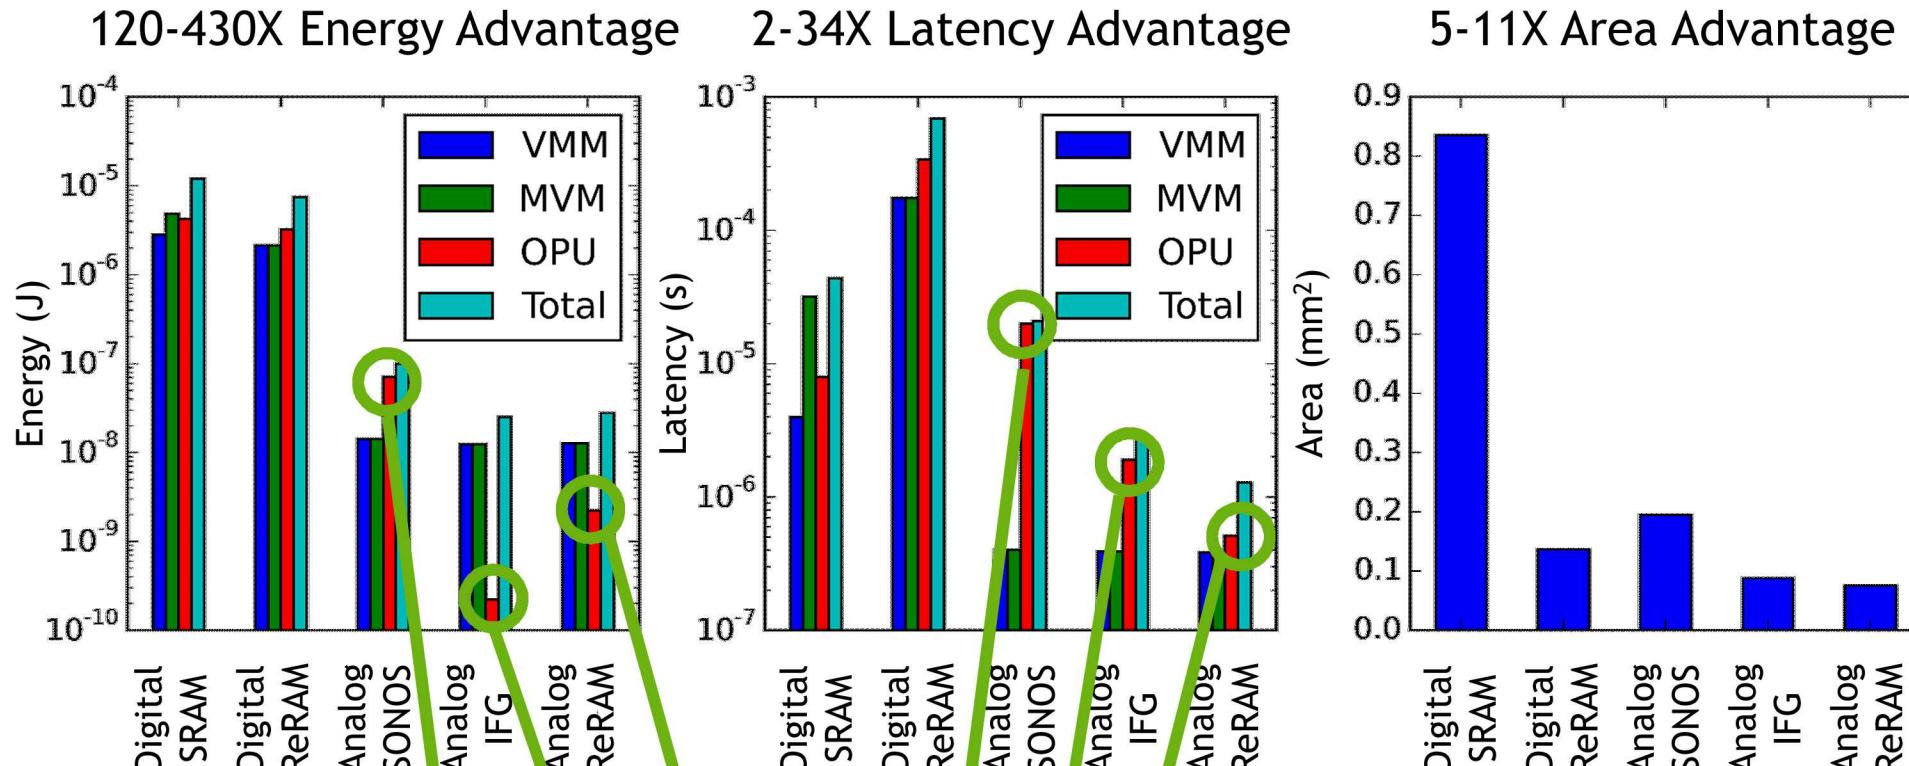
Compare Architectural Advantages: Vector Matrix Multiply



All Analog Vector Matrix Multiply and Matrix Vector Multiply have same energy and latency

- Entirely dominated by ADC, device properties irrelevant

Compare Architectural Advantages: Outer Product Update

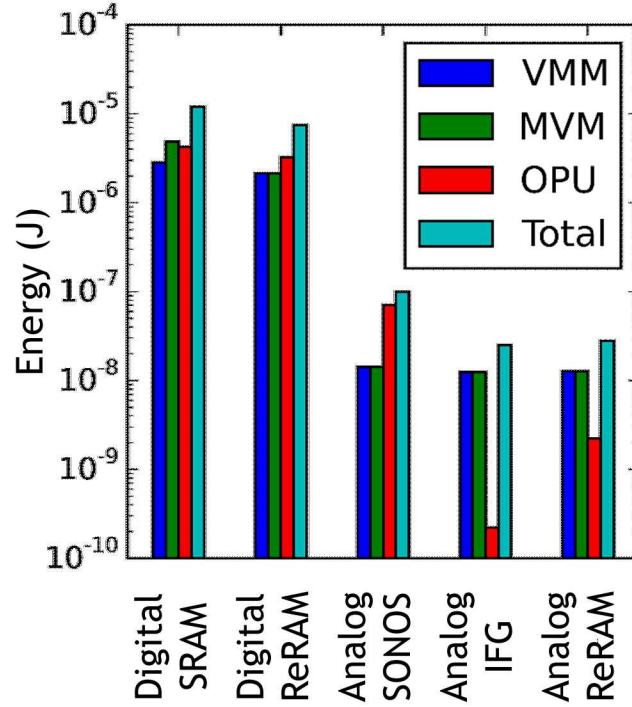


Outer Product Update is device dependent

- SONOS has slow write (~1 ms) and high write voltage (11V)
- IFG and ReRAM write energy negligible compared to VMM
- IFG has extra delay over ReRAM for access device to turn off

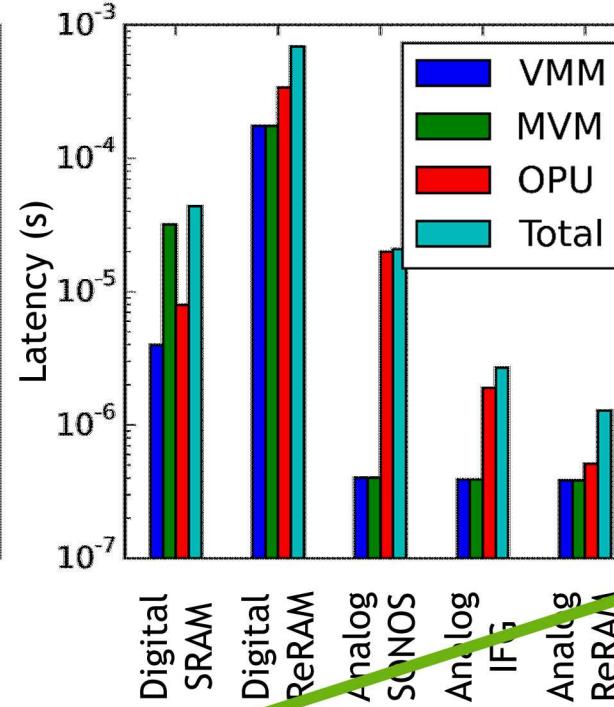
Compare Architectural Advantages: Area

120-430X Energy Advantage



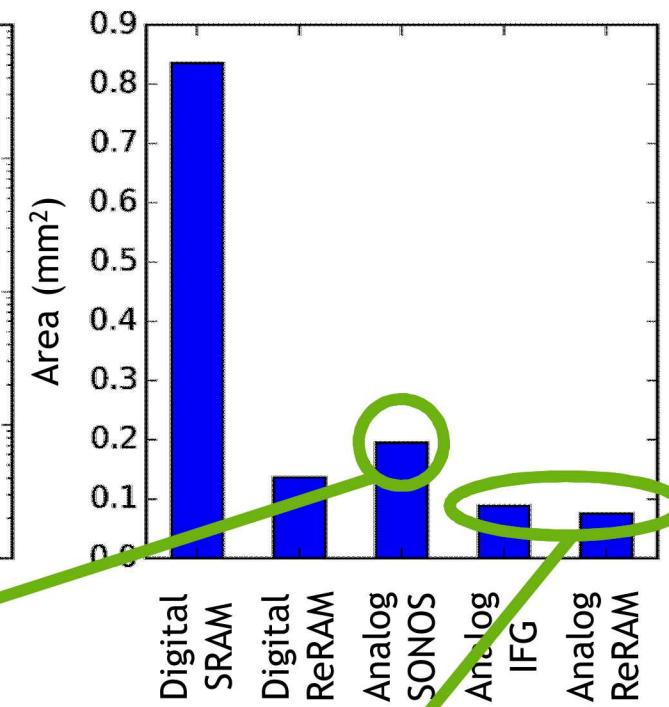
SONOS area cost reasonable, roughly doubles area

2-34X Latency Advantage

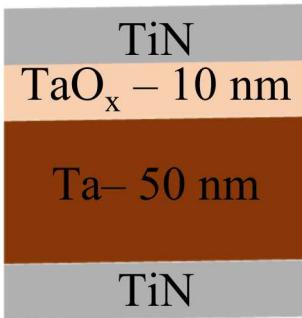


IFG and ReRAM go over transistors, area dominated by ADC and DAC

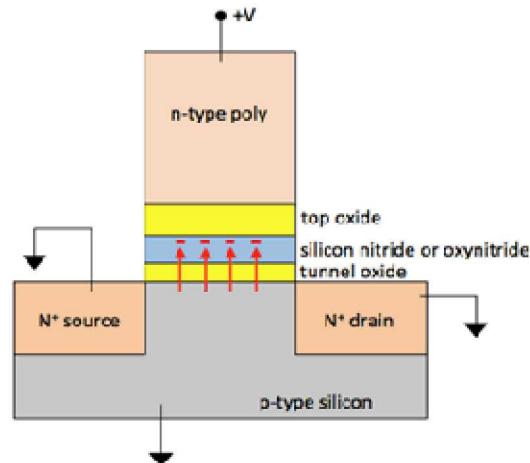
5-11X Area Advantage



ReRAM



SONOS Silicon-Oxygen- Nitrogen-Oxygen-Silicon

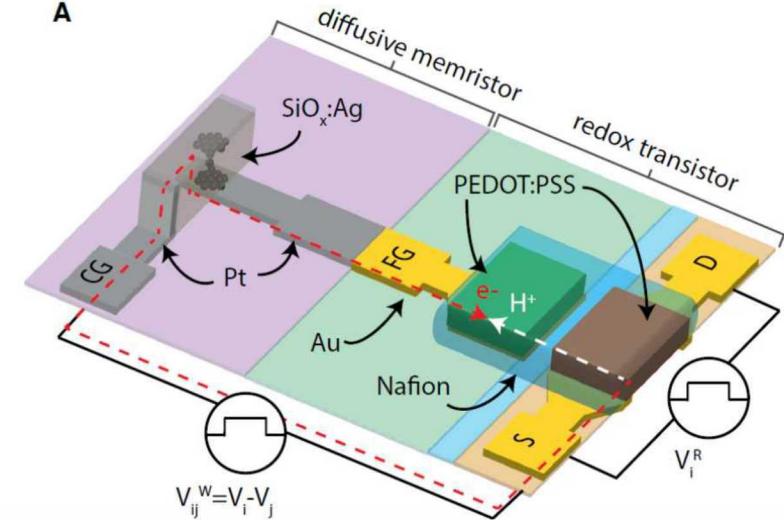


- Large Energy/Area/Latency advantage over digital
- Accuracy not good enough
- Back end of line compatible
- Under commercial development

- Moderate Energy/Area/Latency advantages over digital
- High Accuracy
- Commercially available
- Need to prove endurance and device to device variability

Ionic Floating-Gate Memory

A



- Large Energy/Area/Latency advantages over digital
- High Accuracy
- Not clear how to integrate
- Has retention challenges

Alternate Computing Paradigms

40

Crossbar Based Computing Architectures

- Vector Matrix Multiplication
- Outer Product Update
- Crossbar Based Matrix Solvers
- Ternary Content Addressable Memory

Neuro-inspired Computing

- Hyperdimensional Computing
- Local Learning Rules
- Spiking Neural Networks

Probabilistic and Stochastic Circuits

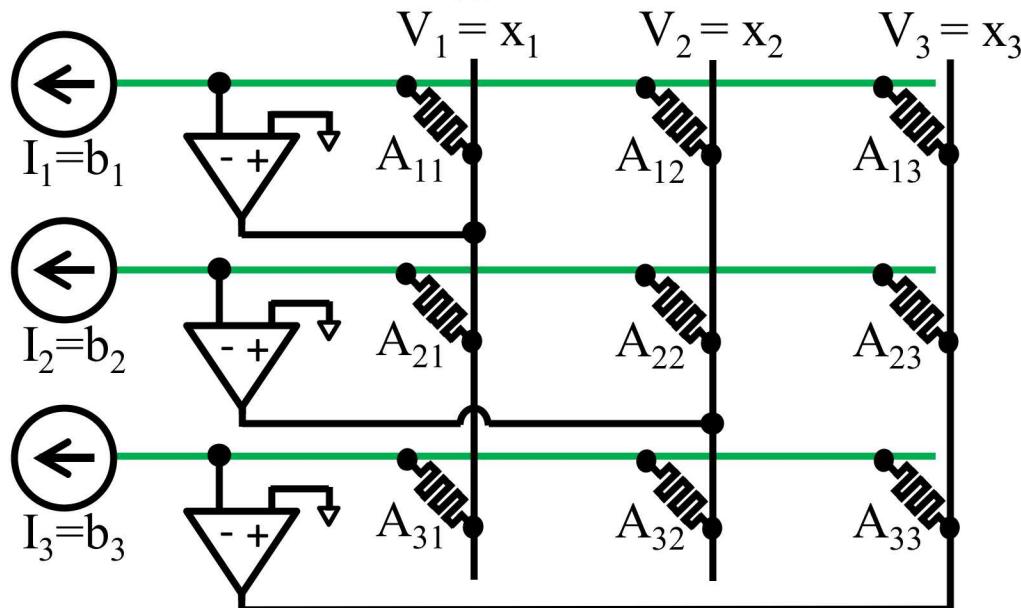
Computing With Dynamical Systems

- Simulated Annealing
- Coupled Oscillator based energy minimization

Analog Matrix Inversion

Analog matrix inversion can perform a dense approximate matrix solve

$$b = Ax \Rightarrow I_{row} = \sum_{col} G_{row,col} (V_{col} - V_{row}) \rightarrow 0$$

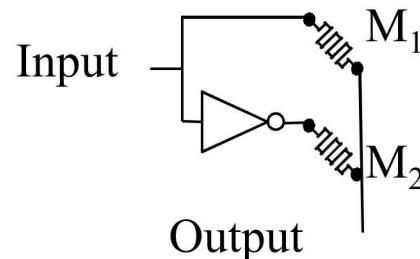


Challenges:

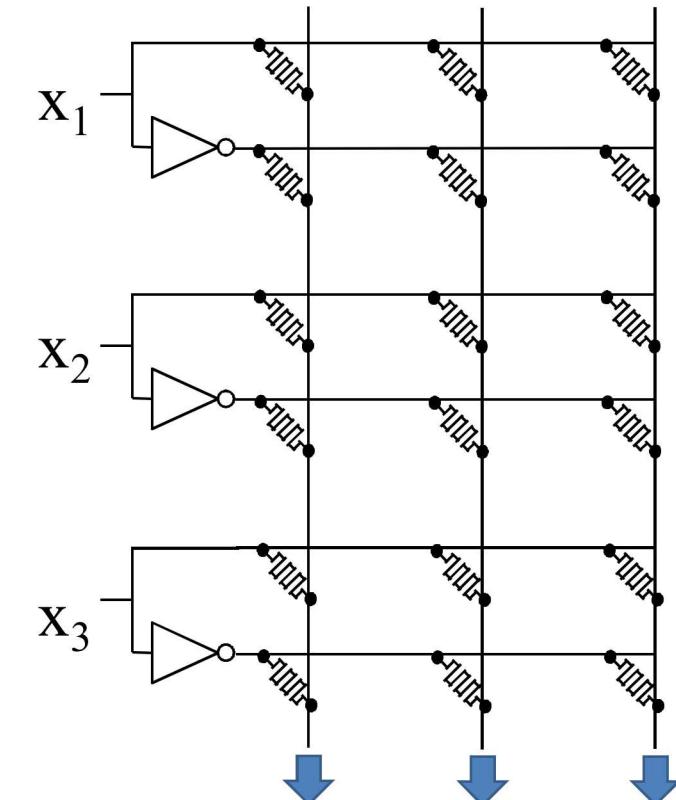
- Matrix inversion is non-linear, limiting how the computation can be split for large matrices
- Analog non-idealities can cause significant errors

Ternary Content Addressable Memory (TCAM)

- Can do very efficient fast pattern matching to search stored data
 - Data analytics, k-nearest neighbors machine learning
 - Sparse matrix multiplication
 - Associative Computing
- Crossbars can implement extremely efficient TCAMs



Input	Stored State	M ₁ State	M ₁ Output	M ₂ State	M ₂ Output	Total Output
0	0	0	0	1	1	1
1	0	0	0	1	0	0
0	1	1	0	0	0	0
1	1	1	1	0	0	1
0	x	1	0	1	1	1
1	x	1	1	1	0	1



Crossbar Based Computing Architectures

- Vector Matrix Multiplication
- Outer Product Update
- Crossbar Based Matrix Solvers
- Ternary Content Addressable Memory

Neuro-inspired Computing

- Hyperdimensional Computing
- Local Learning Rules
- Spiking Neural Networks

Probabilistic and Stochastic Circuits

Computing With Dynamical Systems

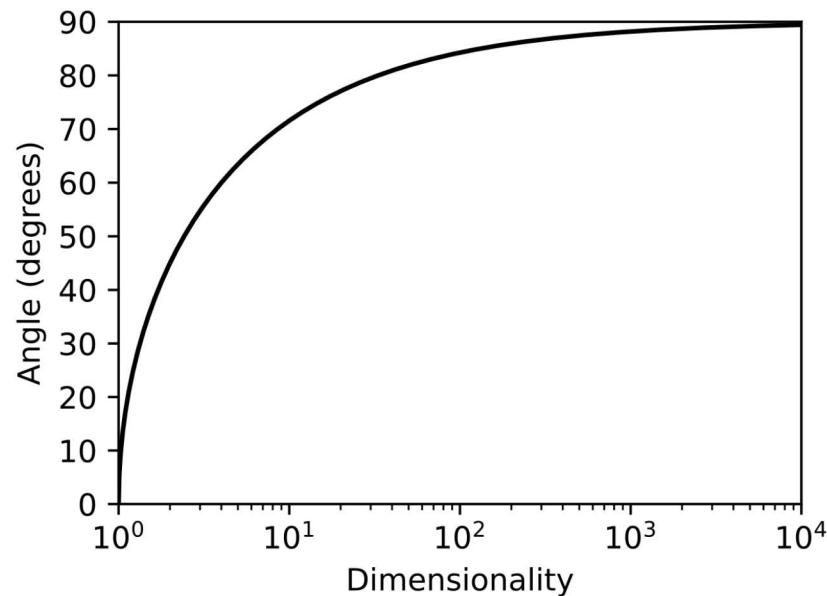
- Simulated Annealing
- Coupled Oscillator based energy minimization

Hyper-Dimensional Computing

Store data in redundant high dimensional vectors

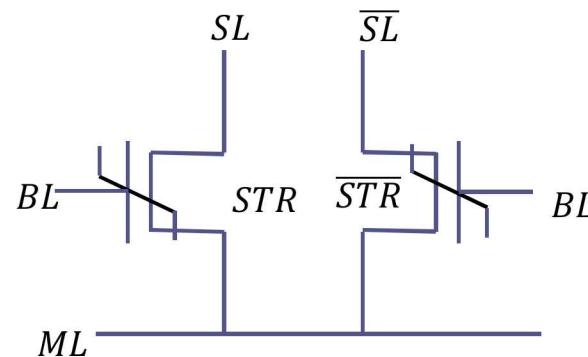
1 0 0 1 0 1 1 0 1 1 0 0 1 0 1 0 1 1 1 0 0 0 1 1 0 0 1 0

As dimensionality increases, two random vectors are nearly orthogonal



Encode data by combining vectors such that the more similar the data is, the smaller the angle between the vectors

Use ferroelectric content addressable memory to enable pattern matching

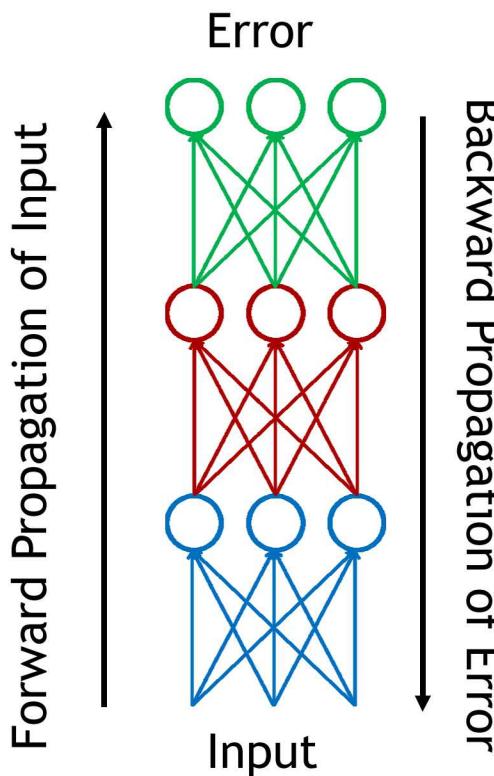


New hardware allows for processing large scale vectors and therefore new algorithms that would otherwise be computationally inefficient

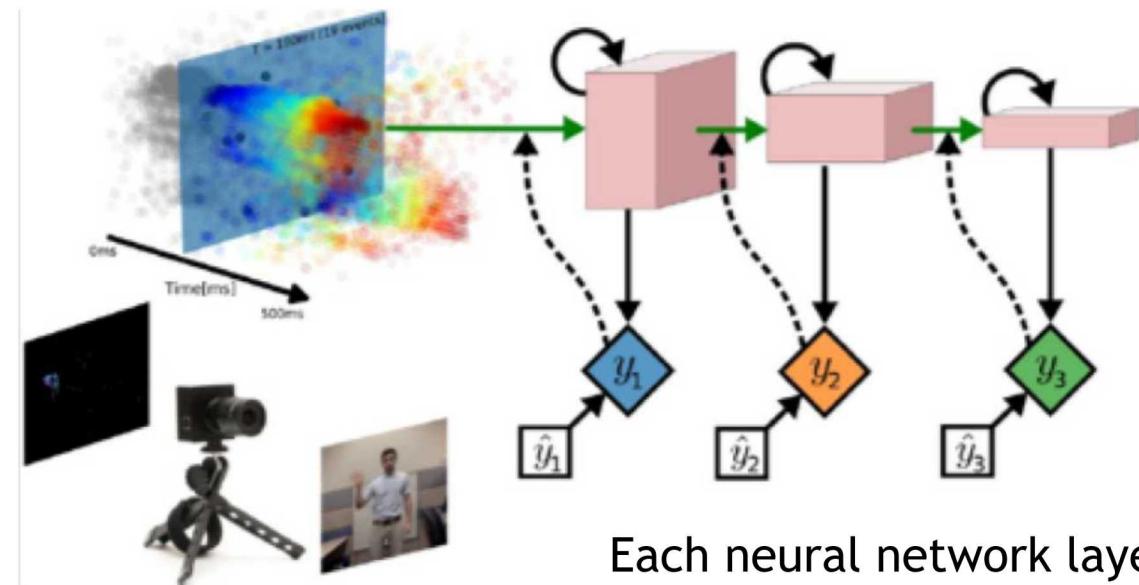
Local Learning Rules

Minimizing Data Movement Directly Minimizes Energy and Latency

Training neural networks requires backpropagating information across all layers resulting in long range communication and storage of intermediate values



Want learning rules that can train using only local information that is present at a given moment in time



Each neural network layer has its own local classifier

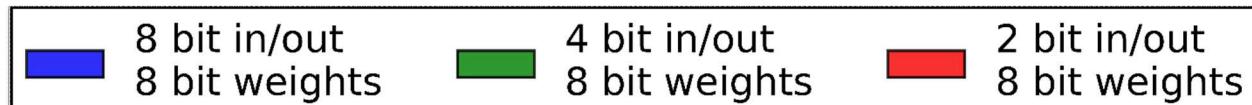
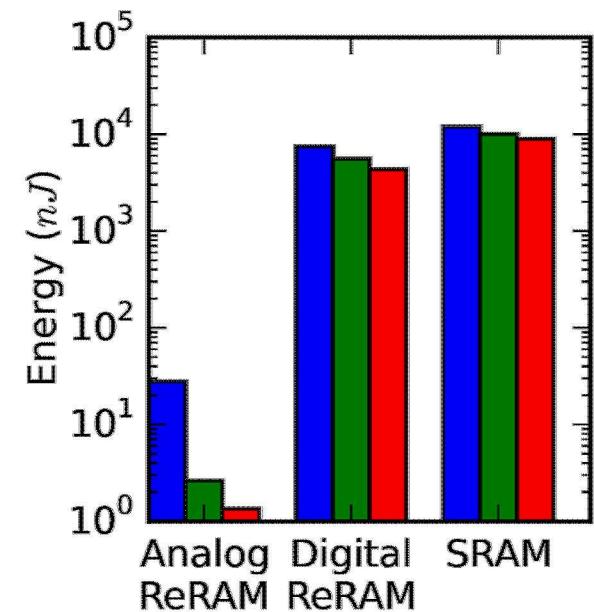
Kaiser, Jacques, Hesham Mostafa, and Emre Neftci. "Synaptic Plasticity Dynamics for Deep Continuous Local Learning (DECOLLE)." *Frontiers in Neuroscience* 14 (2020): 424.

Spiking Neural Networks

Minimizing Data Movement Directly Minimizes Energy and Latency

- For sparse data, communicating only non-zero values is more energy efficient than communicating all data
 - Need to account for overhead of including an address in flexible routing based networks
- Analog system energy is limited by analog to digital conversion
 - Binary outputs from an analog system are far more efficient
- Key challenge is developing high accuracy algorithms with binary inputs and outputs

Analog Neural Network Training Accelerator



Crossbar Based Computing Architectures

- Vector Matrix Multiplication
- Outer Product Update
- Crossbar Based Matrix Solvers
- Ternary Content Addressable Memory

Neuro-inspired Computing

- Hyperdimensional Computing
- Local Learning Rules
- Spiking Neural Networks

Probabilistic and Stochastic Circuits

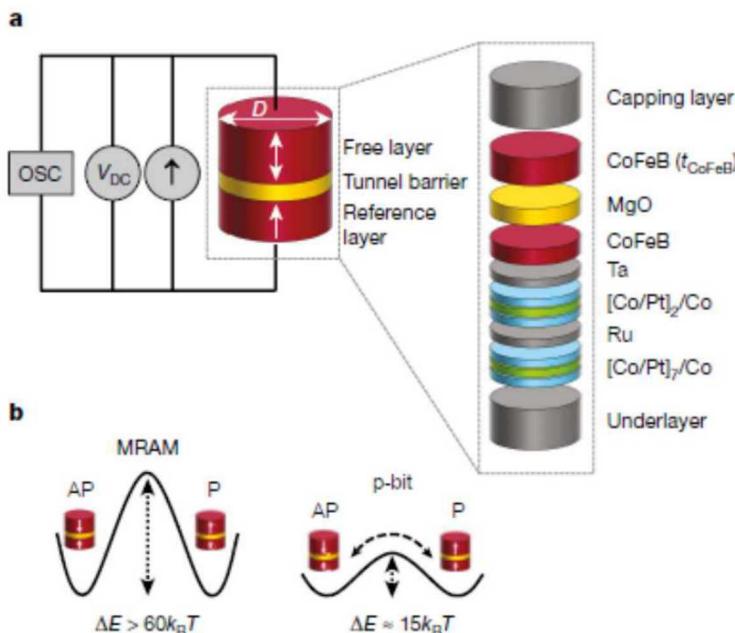
Computing With Dynamical Systems

- Simulated Annealing
- Coupled Oscillator based energy minimization

Probabilistic and Stochastic Circuits

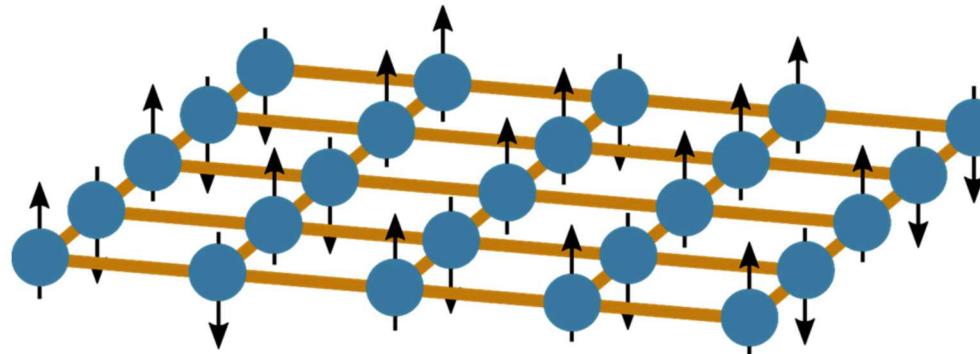
- Generating good random numbers is very computationally intensive
- Compact devices that provide true randomness with tunable probabilities enable new stochastic computing paradigms

Magnetic Tunnel Junction



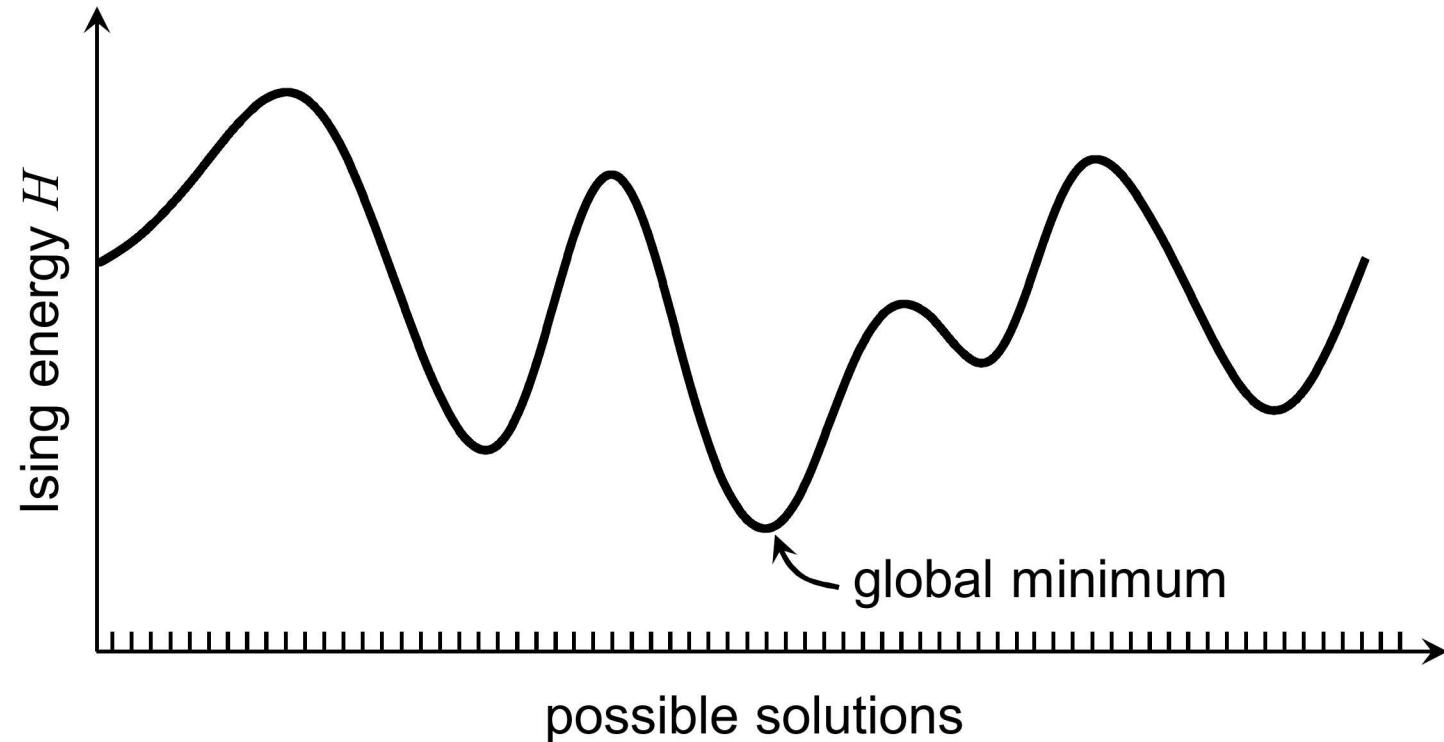
- Single Electron Bipolar Avalanche Transistor
 - Avalanche breakdown is stochastic
- ReRAM
 - The intrinsic variability of memristive switching provides a source of randomness
- Contact-Resistive RAM
- CMOS - ring oscillator jitter
- Stochastic Josephson Junctions

Can map optimization problems to a set of connected stochastic bits: The Ising problem

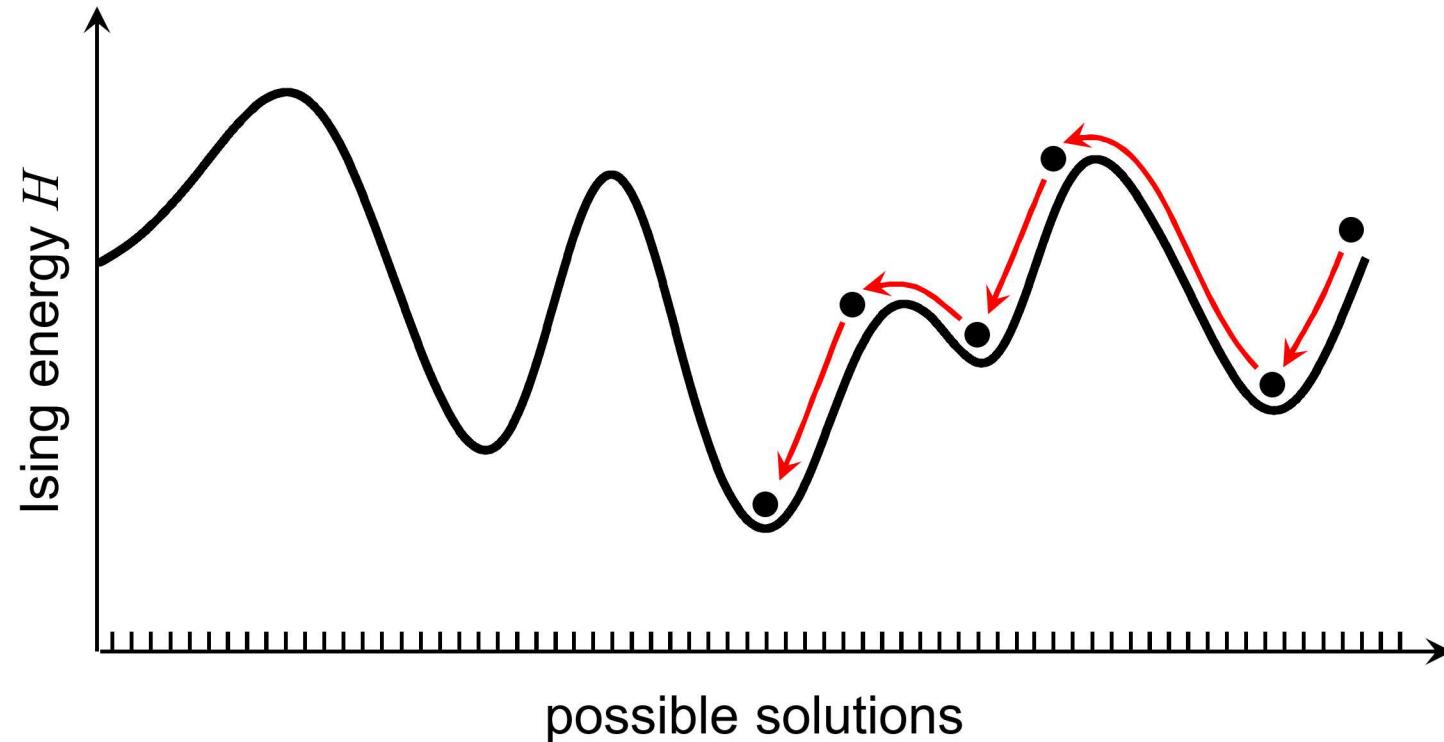


A. Lucas, “Ising formulations of many NP problems,” 2014

Landscape of the Ising problem

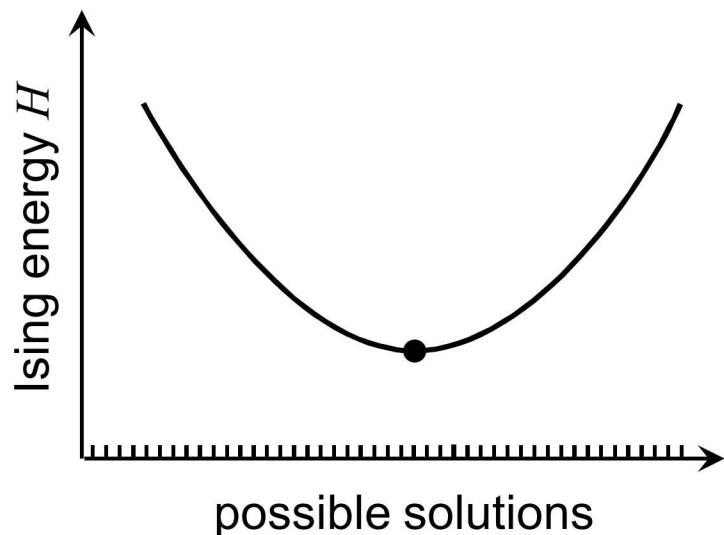


Method 1: Simulated annealing (and other digital heuristics)

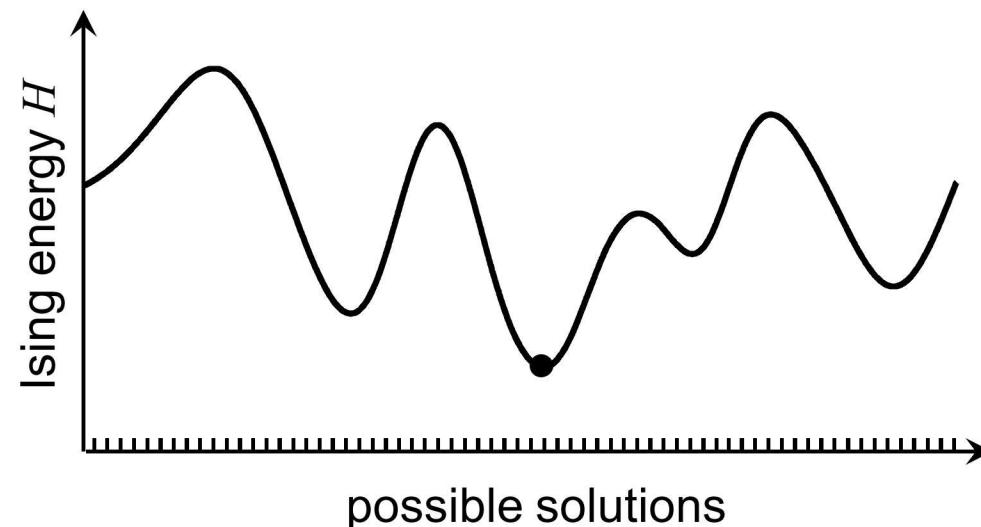


Method 2: Adiabatic quantum optimization

Prepare the ground state
of a simple problem

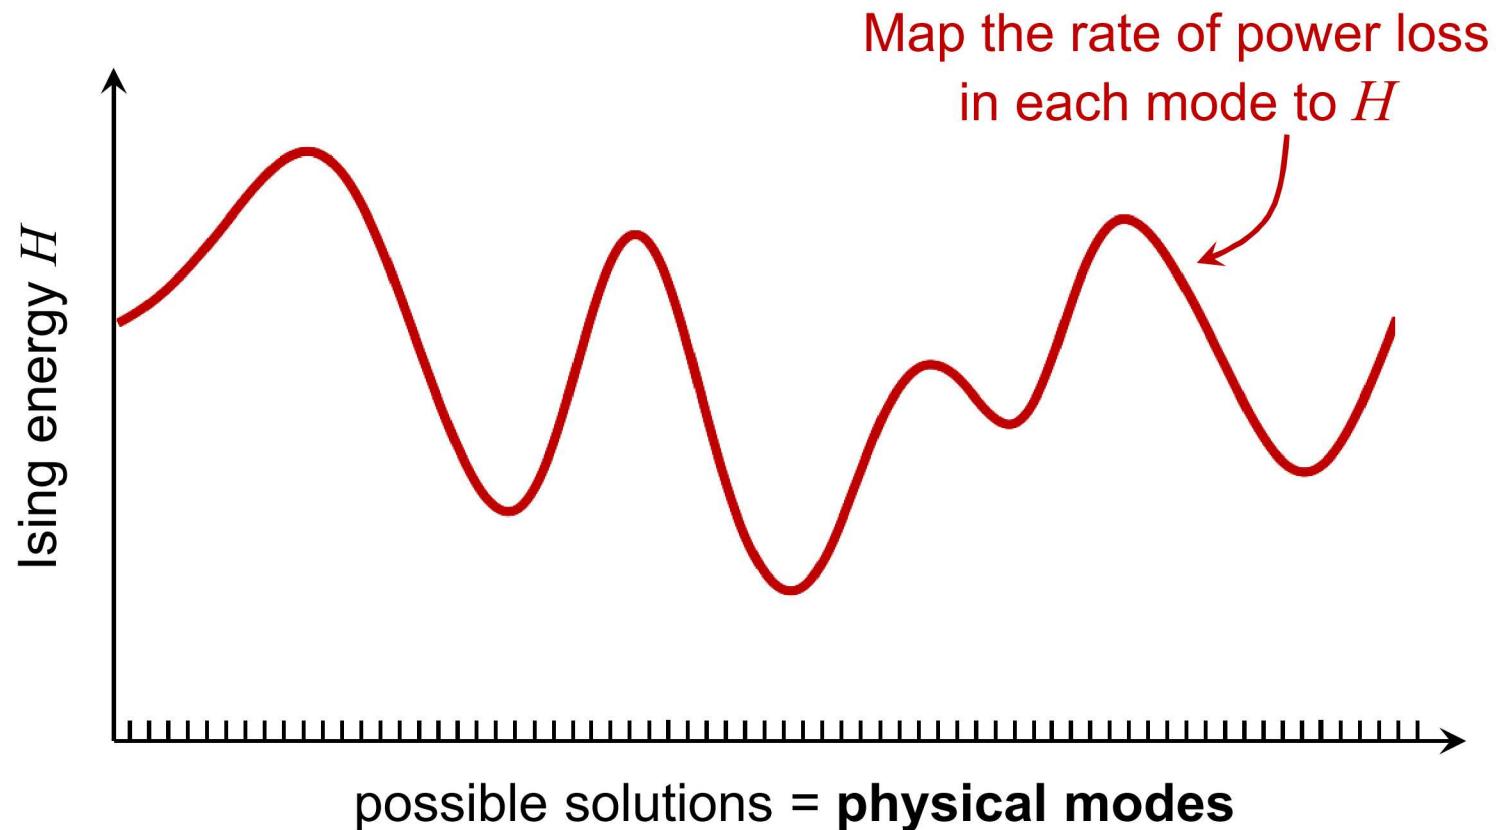


Transform the system into the
desired problem

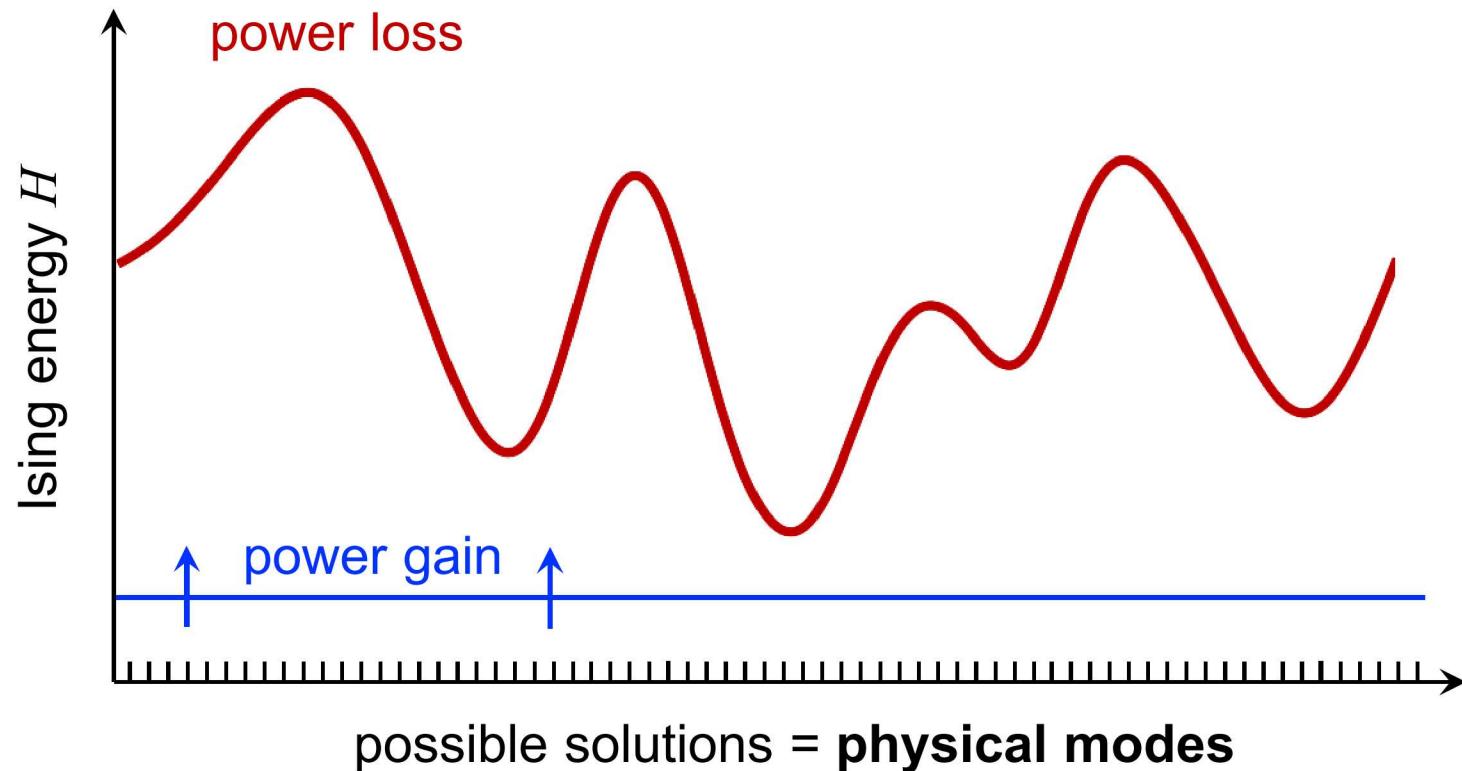


If done *slowly enough*, the system is guaranteed to remain in the
ground state during the full evolution

Method 3: First to threshold

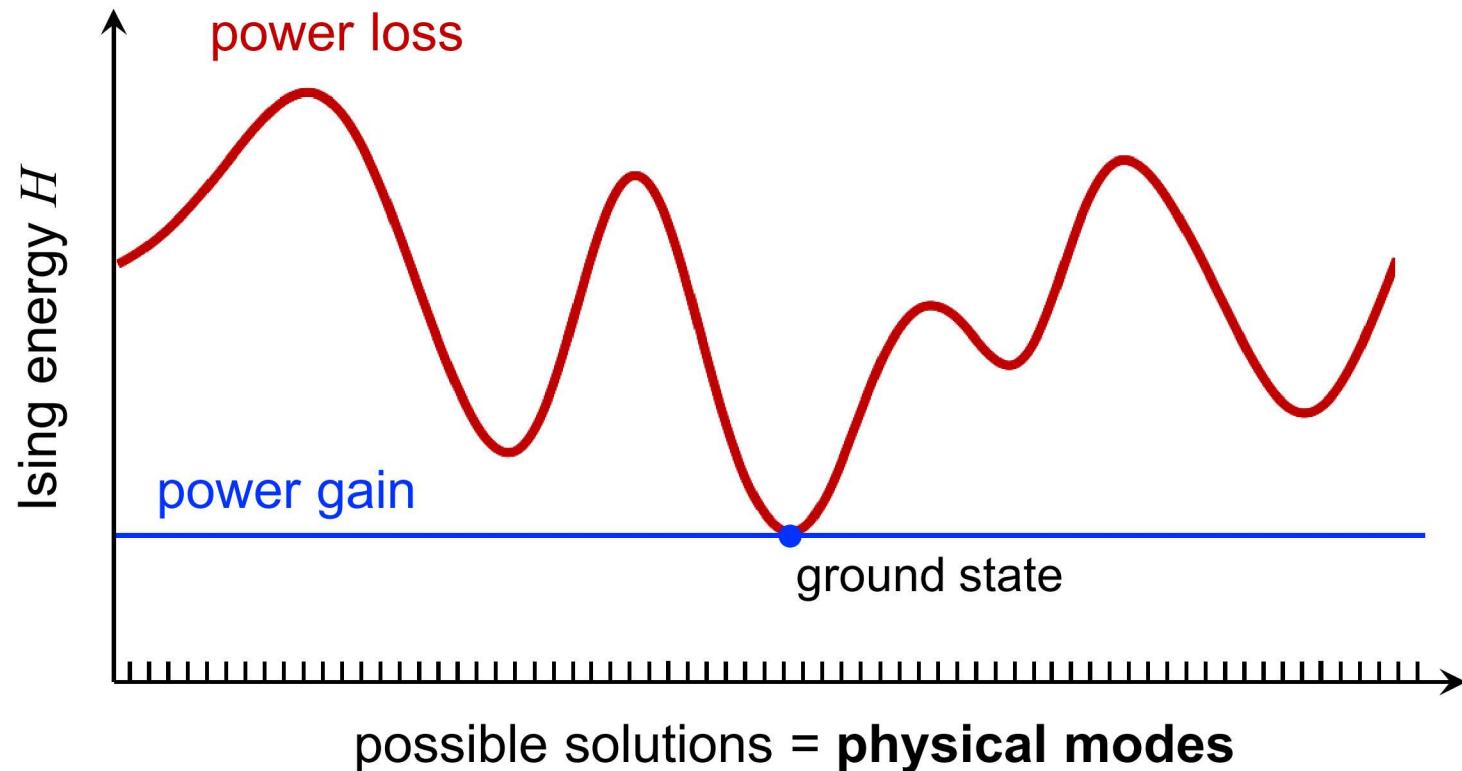


Method 3: First to threshold



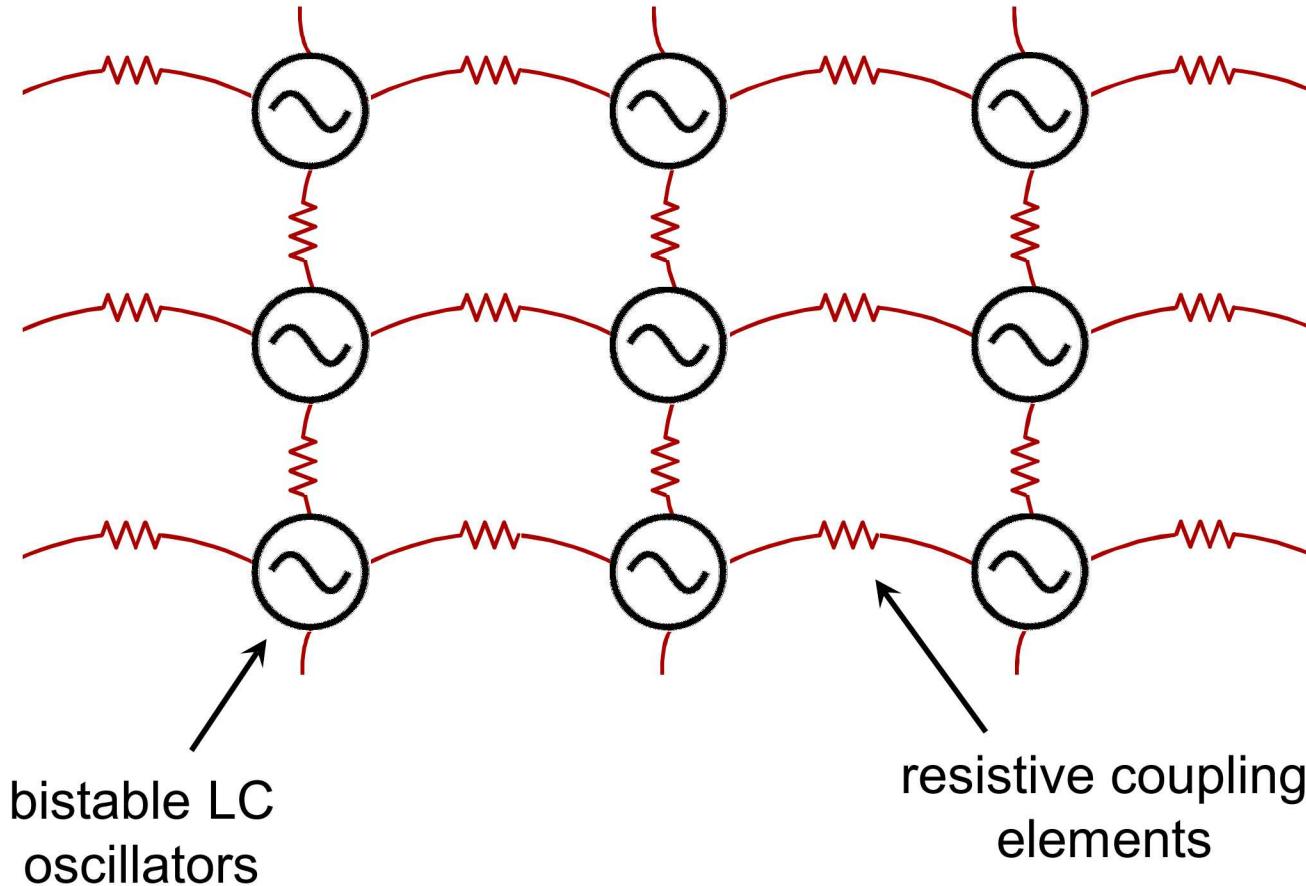
No mode is stable – noise dominates in circuit

Method 3: First to threshold

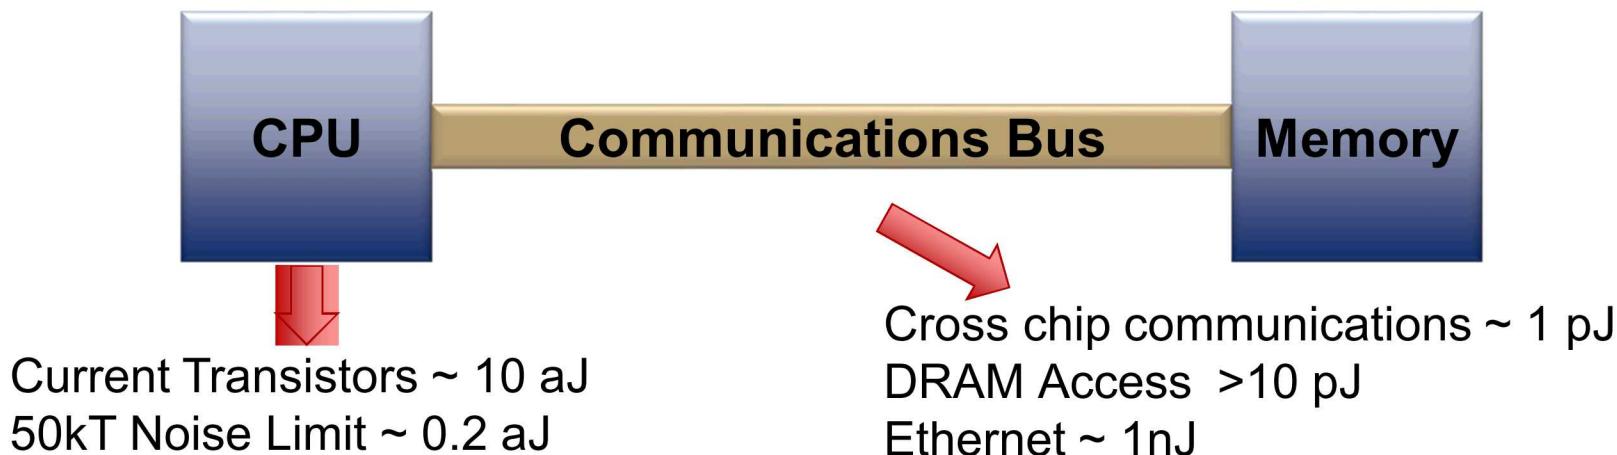


A stable mode emerges, representing the ground state!

Analog electronic Ising machine (high-level view)



Beyond Moore Technologies



Extending Von Neumann

- Low Voltage or Novel Transistors
- Optical Communications
- Reduced Data Movement
 - New On-Chip Memory
 - Processing in Memory

Alternate Computing Paradigms

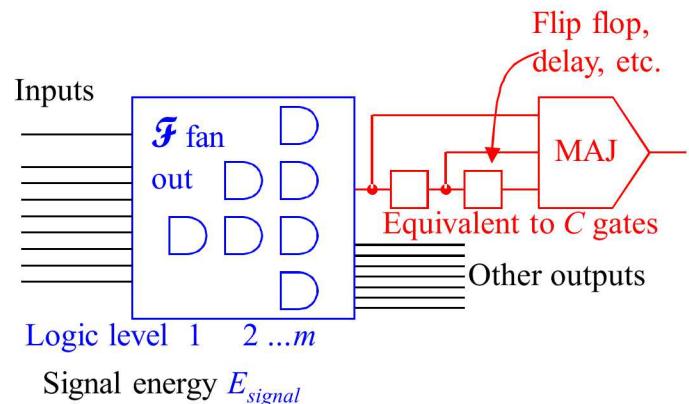
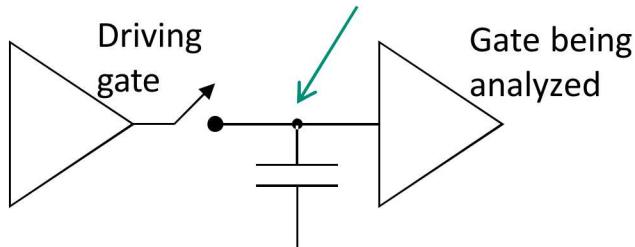
- Neuromorphic
- Analog
- Computing with memory devices
- Quantum
- Stochastic
- Approximate

Going Below 50 kT

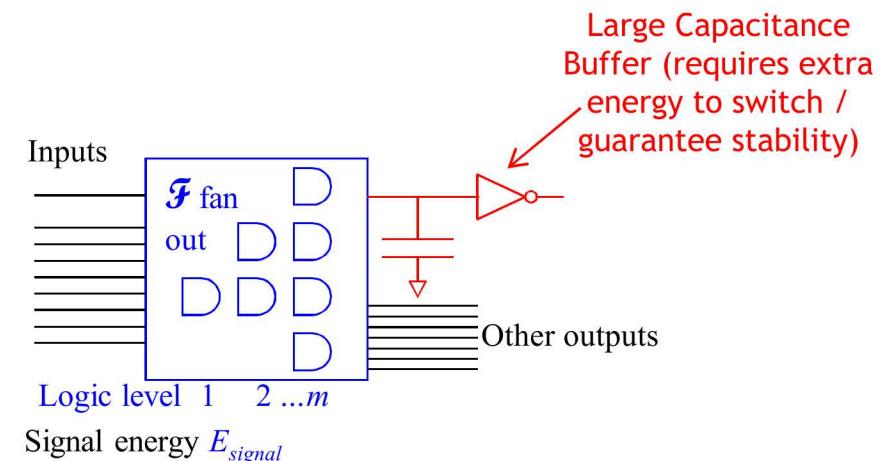
- Error Correction
- Reversible Computing
 - Adiabatic Computing / Energy Recycling
- Superconducting

Temporal Error Correction

$$E_{signal} = CV^2 \rightarrow P(\text{Error}) = e^{-E_{signal}/kT}$$



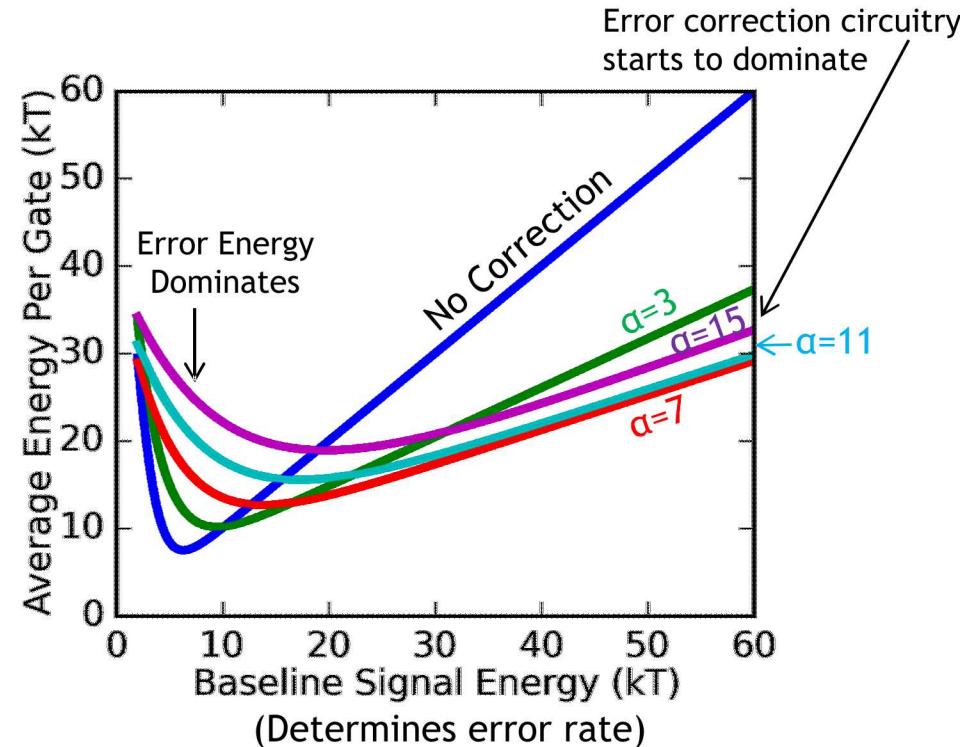
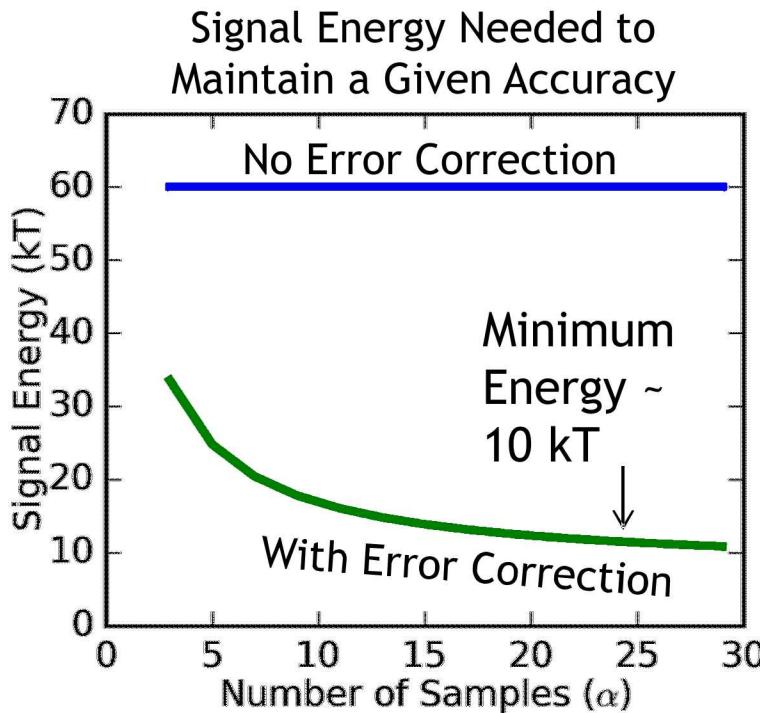
Use multiple samples and take the majority of the result



Use high capacitance stage / low pass filter to integrate out transient thermal errors

How Much Energy is Saved by a Majority Gate?

Take the majority of α samples



Consider a 16x16 multiplier

- 48 levels of logic depth
- 32 inputs/outputs (any input can affect any output)

Can get around a 2X reduction in energy



Consider a signal Energy E_{signal}

The probability of an error due to thermal noise is:

$$P(\text{Error}) = e^{-E_{signal}/kT}$$

In order to ensure a full system with billions of transistors is reliable, we need:

$$E_{signal} \sim 50 \text{ kT}$$

Landauer – Shannon Limit

What is E_{signal} ?

It could be the energy on a single irreversible gate

It could also be the energy in a reversible system that computes a complex logic function the comprises many logical functions

In both cases the signal energy is the same!

Adiabatic Computing



Conventional

- supply charge through a resistor, R , with voltage V_{dd} across it.
- The time it takes is RC

$$E = \frac{V_{dd}^2}{R} \times RC = CV_{dd}^2 = V_{dd} \times Q$$

Adiabatic

- Reduce the power burned in the resistor, by minimizing the voltage across it
 - Charge the circuit with a lower current, I_{low}
- This takes a longer time, τ , to get the required charge: $Q = C \times V_{dd} = I_{low} \times \tau$

$$\begin{aligned} E &= I_{low}^2 \times R \times \tau \\ &= I_{low} \times R \times (I_{low} \times \tau) \\ &= I_{low} \times R \times Q \\ &= V_{low} \times Q \end{aligned}$$

Energy Reduced by:

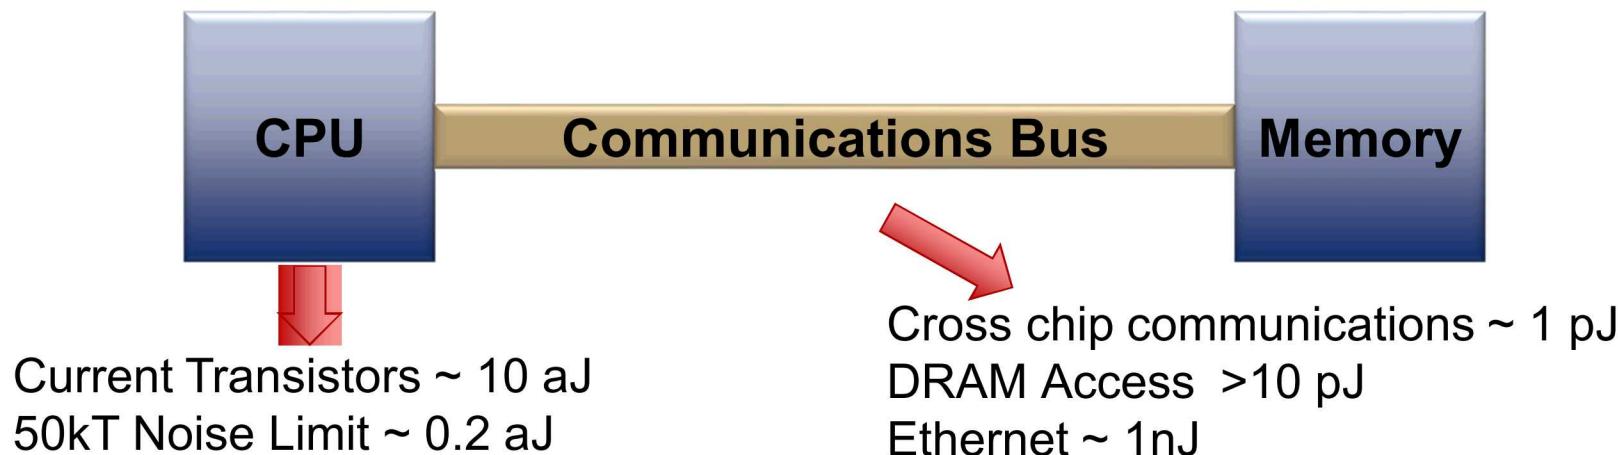
$$\frac{E_{adiabatic}}{E_{classical}} = \frac{V_{low}}{V_{dd}} = \frac{I_{low}R}{V_{dd}} = \frac{\frac{CV_{dd}}{\tau}R}{V_{dd}} = \frac{RC}{\tau}$$

$$\frac{\tau}{RC}$$

Delay Increased By:

(ignoring factors of 2 for simplicity)

Summary: There are many ways to extend Moore's Law!



Extending Von Neumann

- Low Voltage or Novel Transistors
- Optical Communications
- Reduced Data Movement
 - New On-Chip Memory
 - Processing near Memory

Alternate Computing Paradigms

- Neuromorphic
- Analog
- Computing with memory devices
- Quantum
- Stochastic
- Approximate

Going Below 50 kT

- Error Correction
- Reversible Computing
 - Adiabatic Computing / Energy Recycling
- Superconducting

Sandia is Hiring! sandia.gov/careers



~12,500 people in
Albuquerque, NM



~1,800 people in
Livermore, CA

Fulfilling Our National Security Mission



Nuclear Deterrence



*Defense Nuclear
Nonproliferation*



*National Security
Programs*



*Energy & Homeland
Security*



*Advanced Science &
Technology*