Digital Computing Beyond Moore's Law

Supercomputing Frontiers, Singapore

March 16, 2017

John Shalf

Lawrence Berkeley National Laboratory

Post Exsascale Landscape
MIND THE GAP!

Moore's Law
Lithography Scaling
2x increased density
2x lower power
Every 2 years!

Now – 2025
Moore's Law continues through ~5nm -- beyond which diminishing returns are expected.

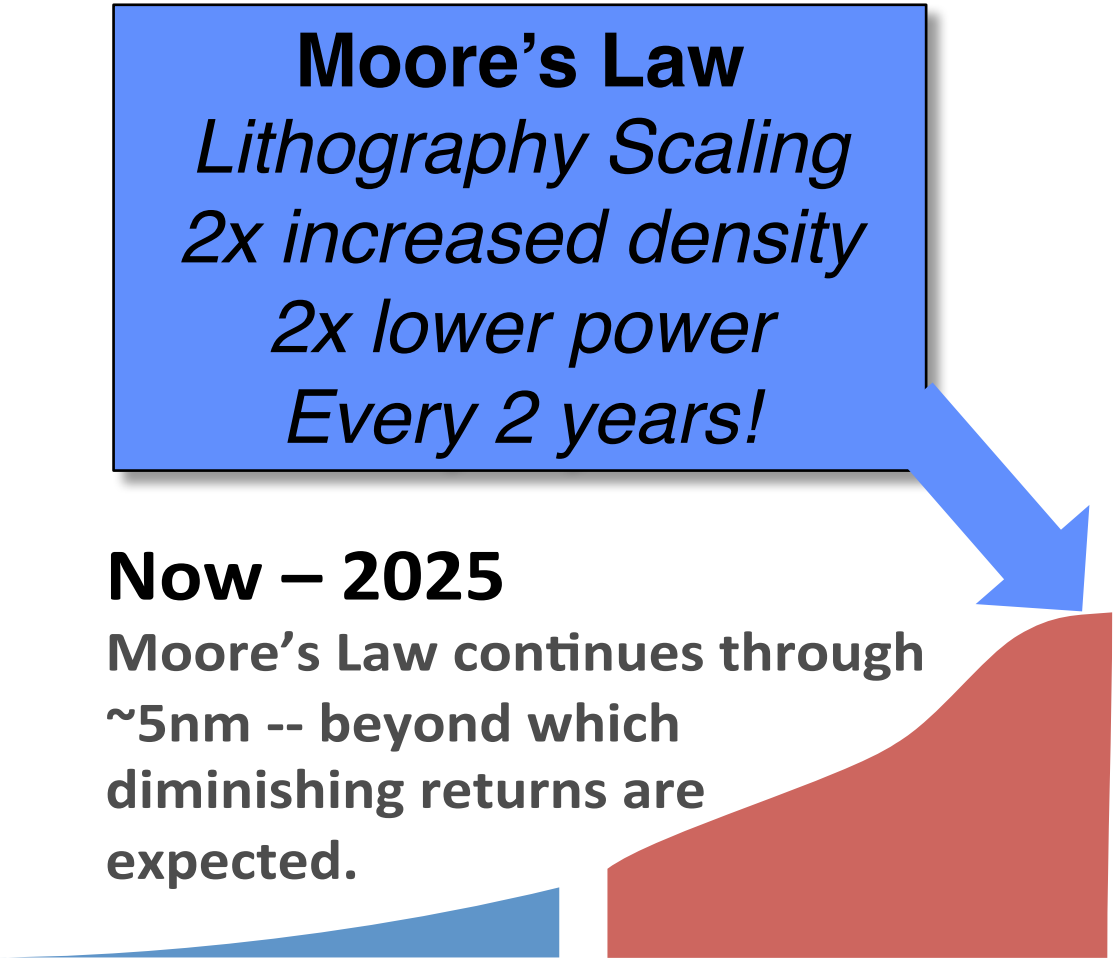

2016

2016-2025

End of Moore's Law
2025-2030?

Post Moore Scaling
New materials and devices introduced to enable continued scaling of digital electronics performance and efficiency.

2025+

Is it the end?

"I predict Moore's Law will never end.
That way I will only be wrong once!"

Alan Kay: Communications of the ACM 1989

Technology Scaling Trends

Exascale in 2021… and then what?

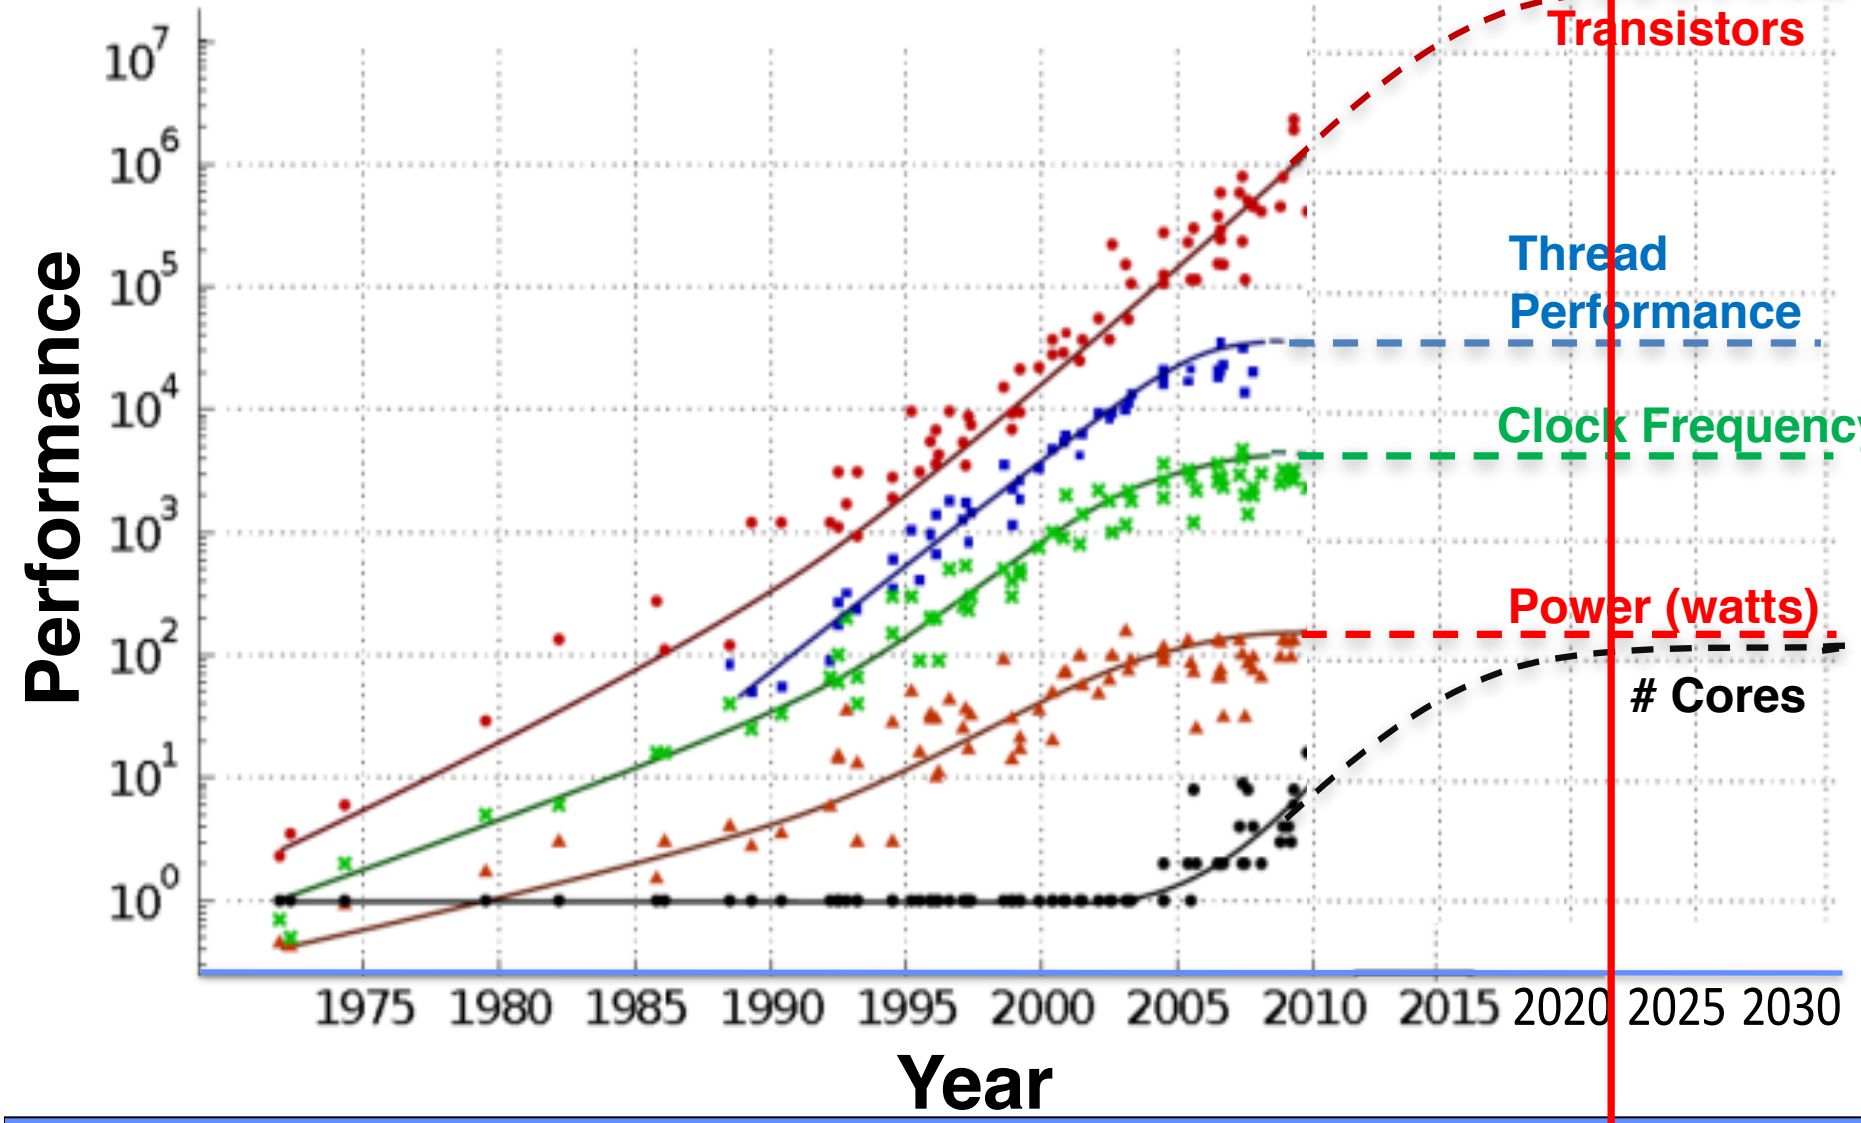

Figure courtesy of Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith

SO WHAT IF MOORE'S LAW ENDS ? WHY SHOULD I CARE?

Nothing lasts forever…

Especially an exponential trend!

IT challenge for future electricity supply

Global Semiconductor market size ~ $5Trillion by 2030

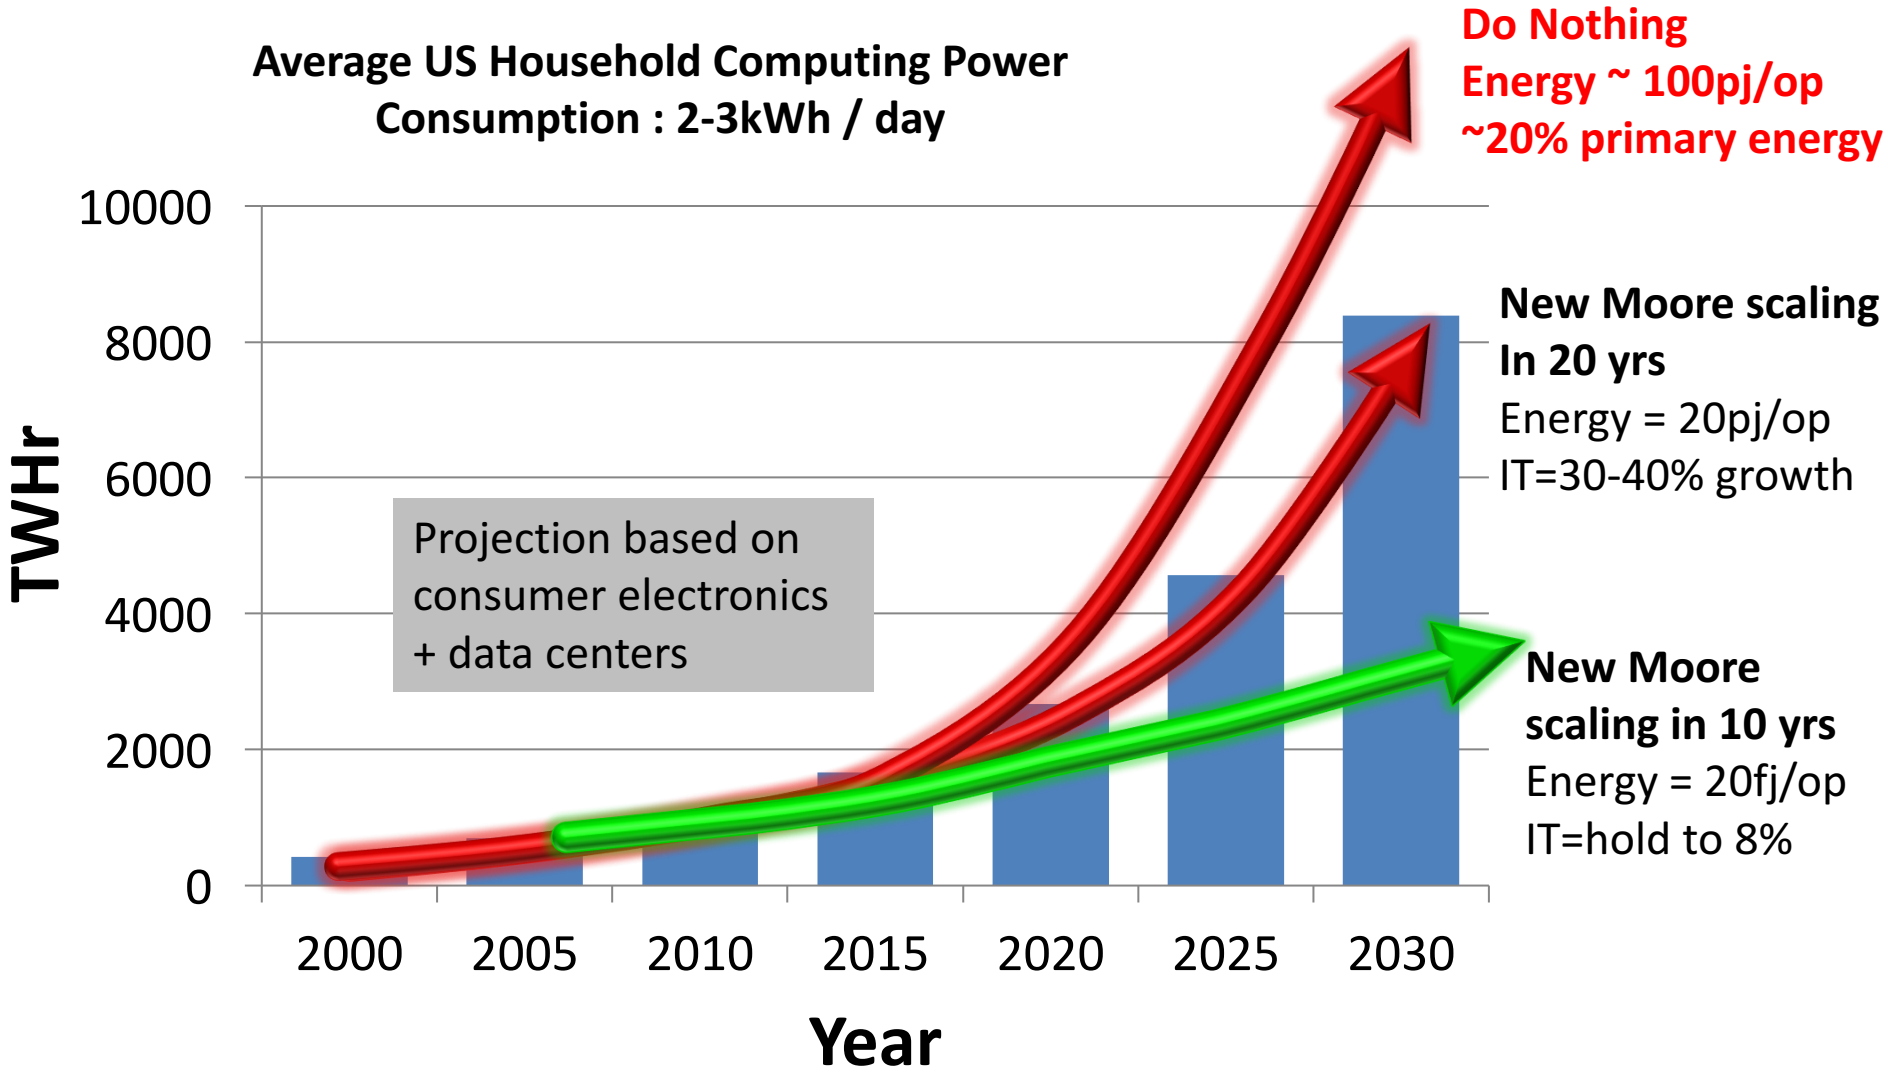

www.alliancetrustinvestments.com/sri-hub/posts/Energy-efficient-data-centres

www.iea.org/publications/freepublications/publication/gigawatts2009.pdf

Innovation is the Answer!

Moore's Law is an economic theory.

There are ways to continue scaling of digital technology after the end of classical lithographic scaling

(e.g. end of Dennard Scaling in ~2004
No more exponential clock frequency scaling
Move to exponentially increasing parallelism)

Post-Lithographic Scaling Options

There are other ways to continue Moore's Scaling!

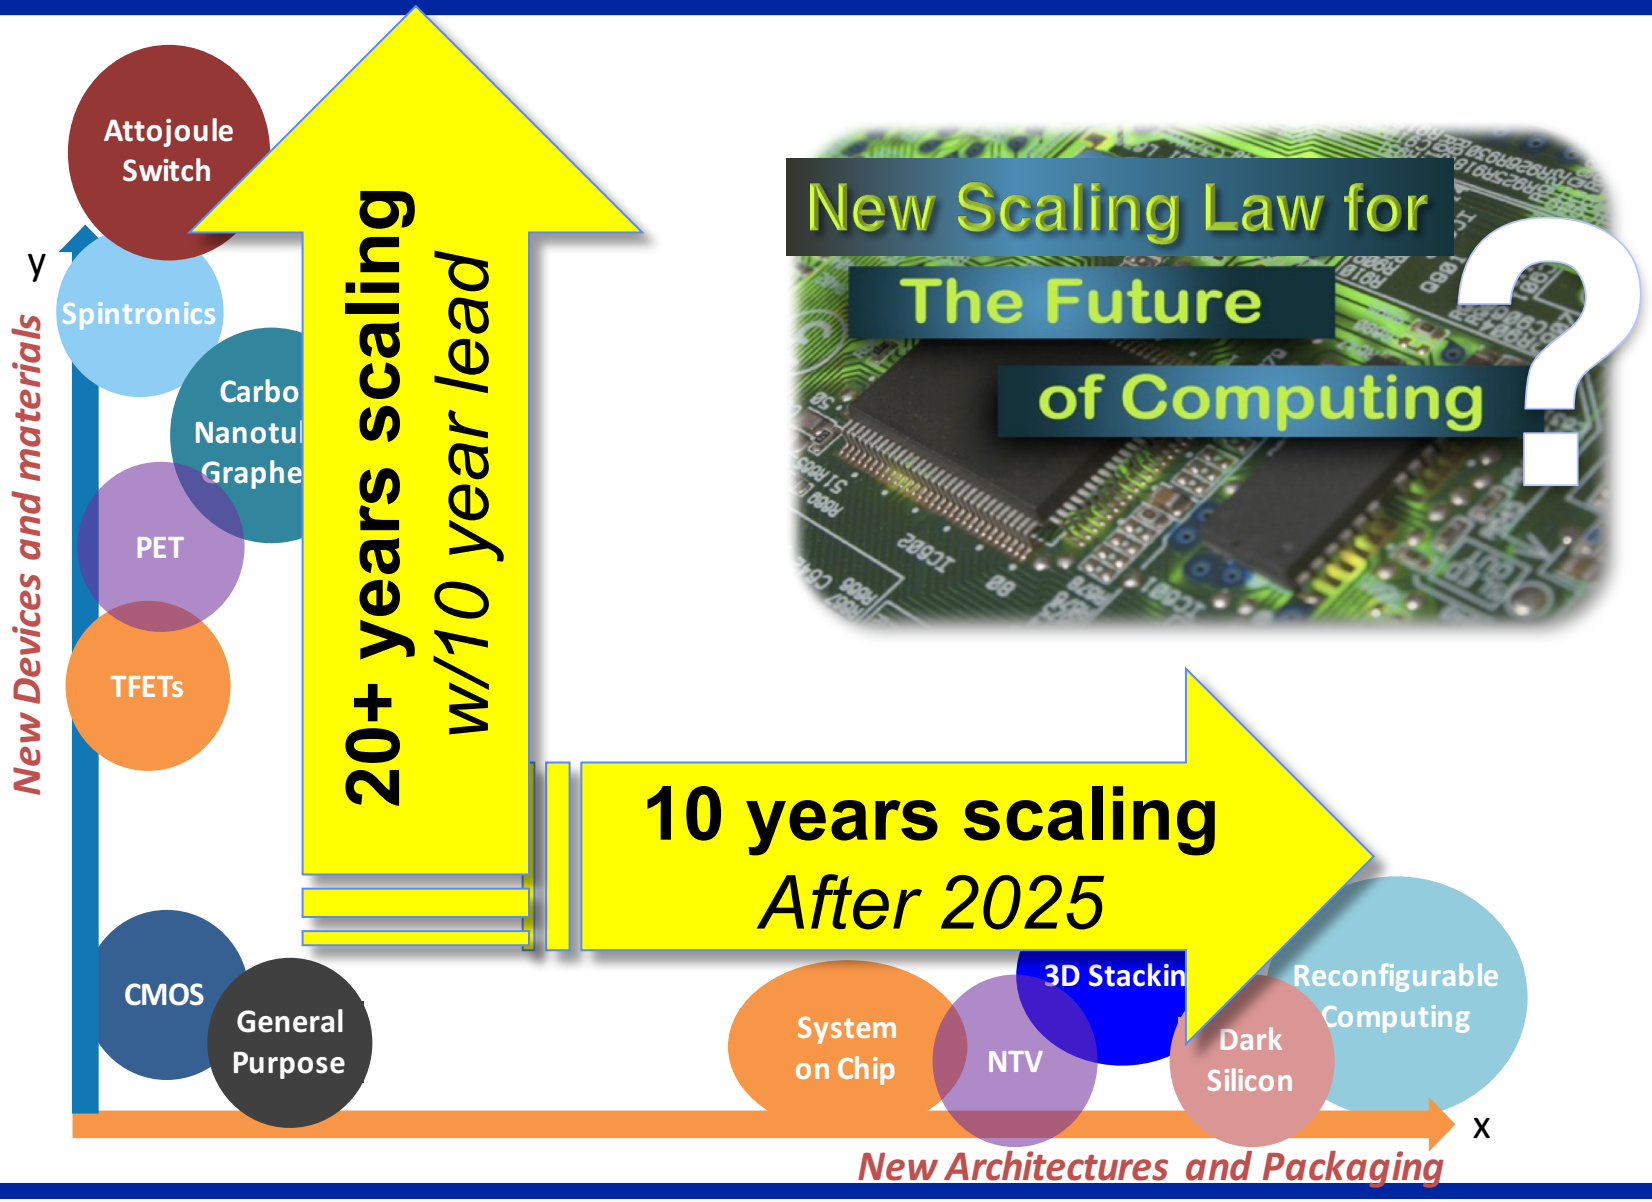

Beyond Moore Computing Taxonomy

Symbolic Computation, Arithmetic, Logic

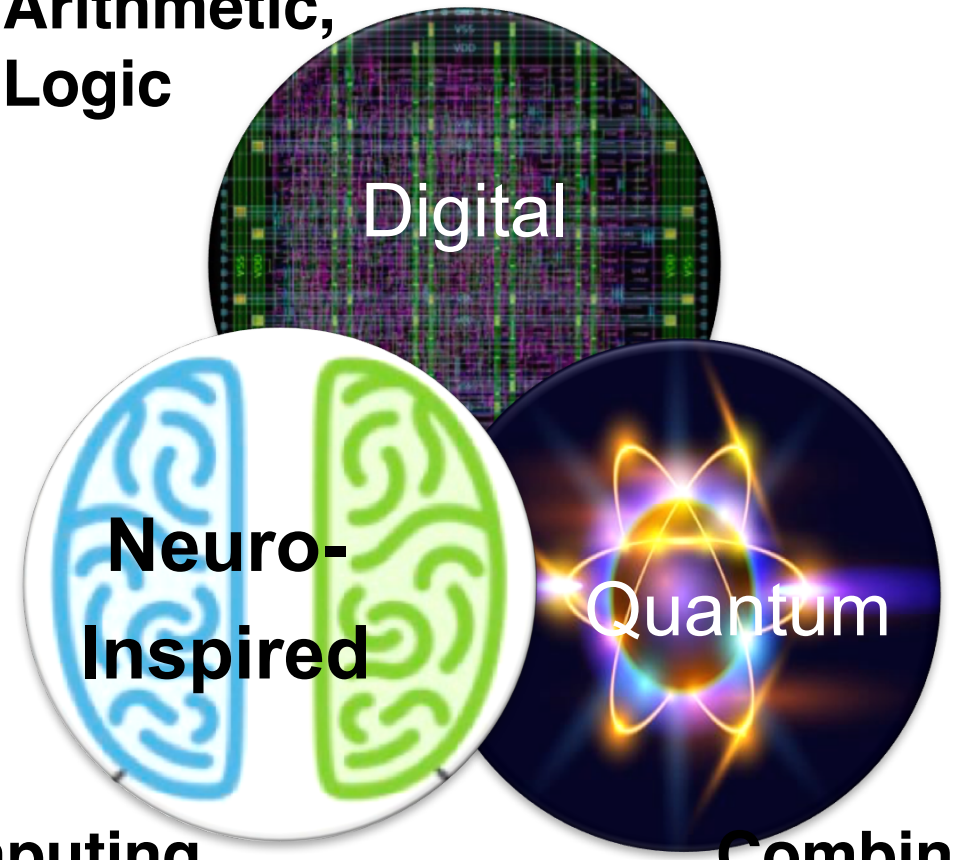

Cognitive Computing, Pattern Recognition

Combinatorial/NP, Annealing/Optimization, Simulated Atoms

Beyond Moore Computing Taxonomy

Symbolic Computation, Arithmetic, Logic

IDA/iARPA Study 2014

- Invest in extending reach of computing to new areas where digital is not efficient by studying Quantum, Neuromorphic
- But don't forget that you need digital *(it offers a kind of computation that is not well replicated alternatives)*

See Our Article in 2016 December Issue of IEEE Computer!

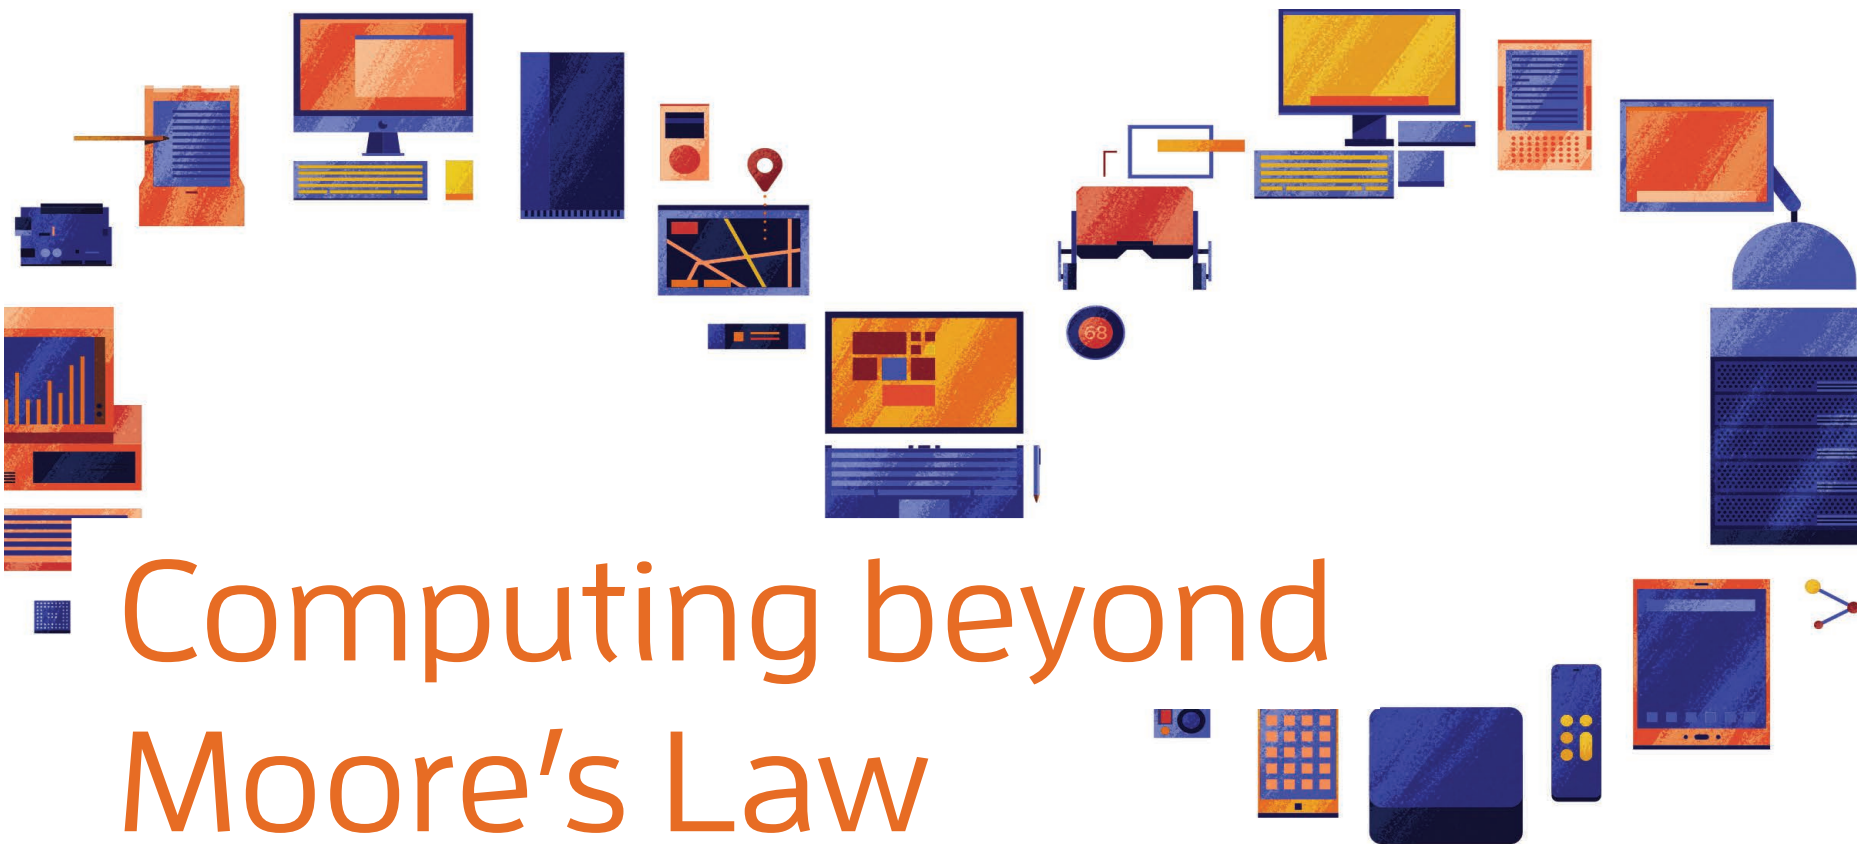

COVER FEATURE **REBOOTING COMPUTING**

Computing beyond Moore's Law

John M. Shalf, Lawrence Berkeley National Laboratory

Robert Leland, Sandia National Laboratories

TABLE 1. Summary of techology options for extending digital electronics.

Improvement Class	Technology	Timescale	Complexity	Risk	Opportunity
Architecture and software advances	Advanced energy management	Near-Term	Medium	Low	Low
	Advanced circuit design	Near-Term	High	Low	Medium
	System-on-chip specialization	Near-Term	Low	Low	Medium
	Logic specialization/dark silicon	Mid-Term	High	High	High
	Near threshold voltage (NTV) operation	Near-Term	Medium	High	High
3D integration and packaging	Chip stacking in 3D using thru-silicon vias (TSVs)	Near-Term	Medium	Low	Medium
	Metal layers	Mid-Term	Medium	Medium	Medium
	Active layers (epitaxial or other)	Mid-Term	High	Medium	High
Resistance reduction	Superconductors	Far-Term	High	Medium	High
	Crystaline metals	Far-Term	Unknown	Low	Medium
Millivolt switches (a better transistor)	Tunnel field-effect transistors (TFETs)	Mid-Term	Medium	Medium	High
	Heterogeneous semiconductors/strained silicon	Mid-Term	Medium	Medium	Medium
	Carbon nanotubes and graphene	Far-Term	High	High	High
	Piezo-electric transistors (PFETs)	Far-Term	High	High	High
Beyond transistors (new logic paradigms)	Spintronics	Far-Term	Medium	High	High
	Topological insulators	Far-Term	Medium	High	High
	Nanophotonics	Near/Far-Term	Medium	Medium	High
	Biological and chemical computing	Far-Term	High	High	High

Accelerated development & optimization For new Logic and Memory Devices

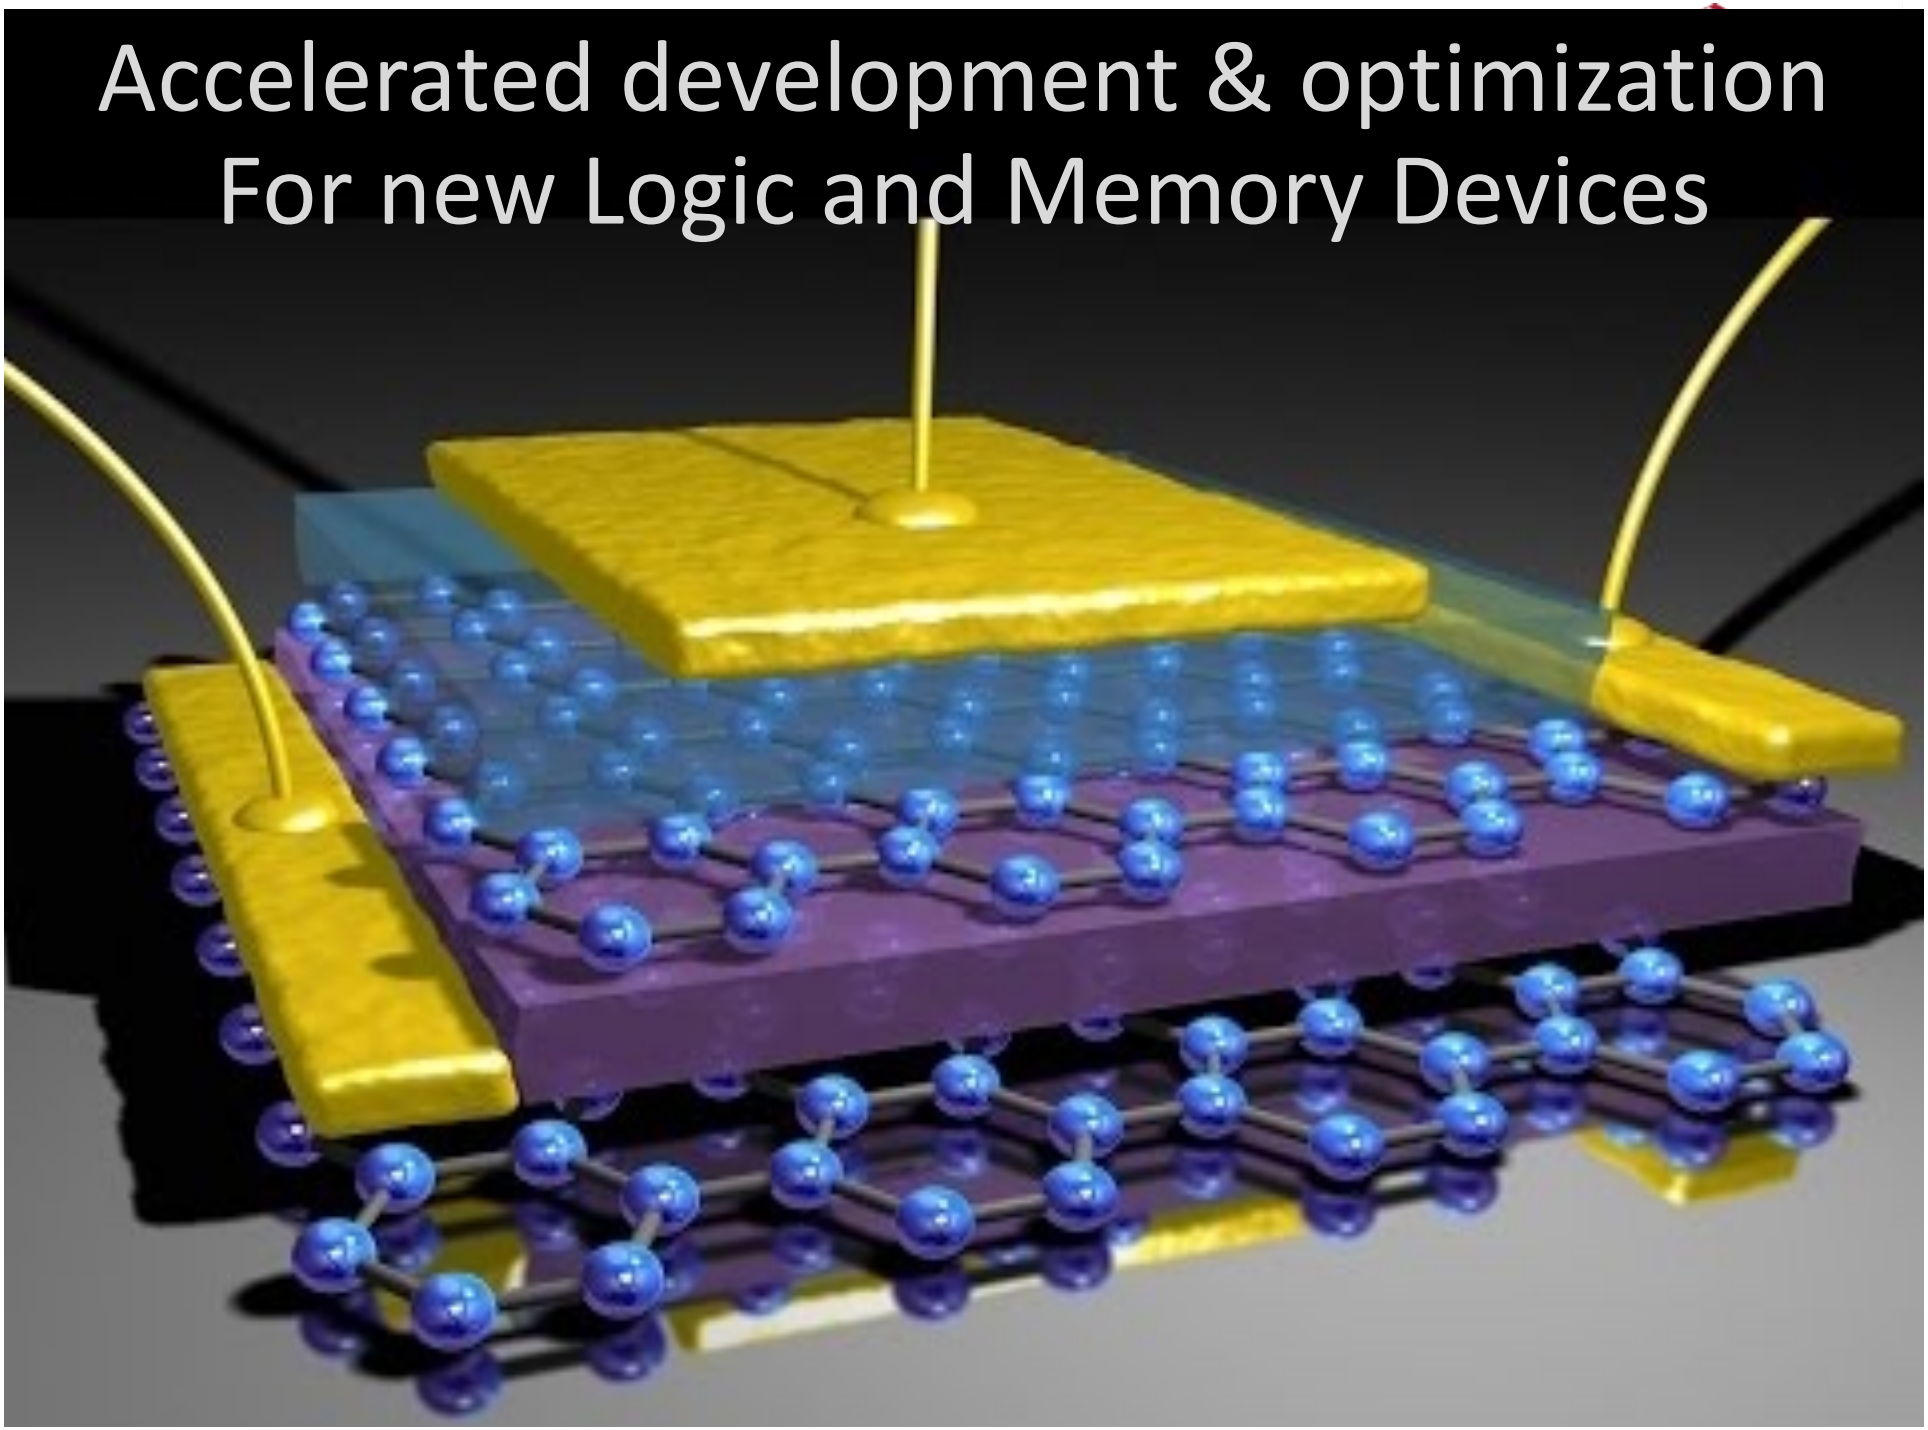

Long Term: New Materials

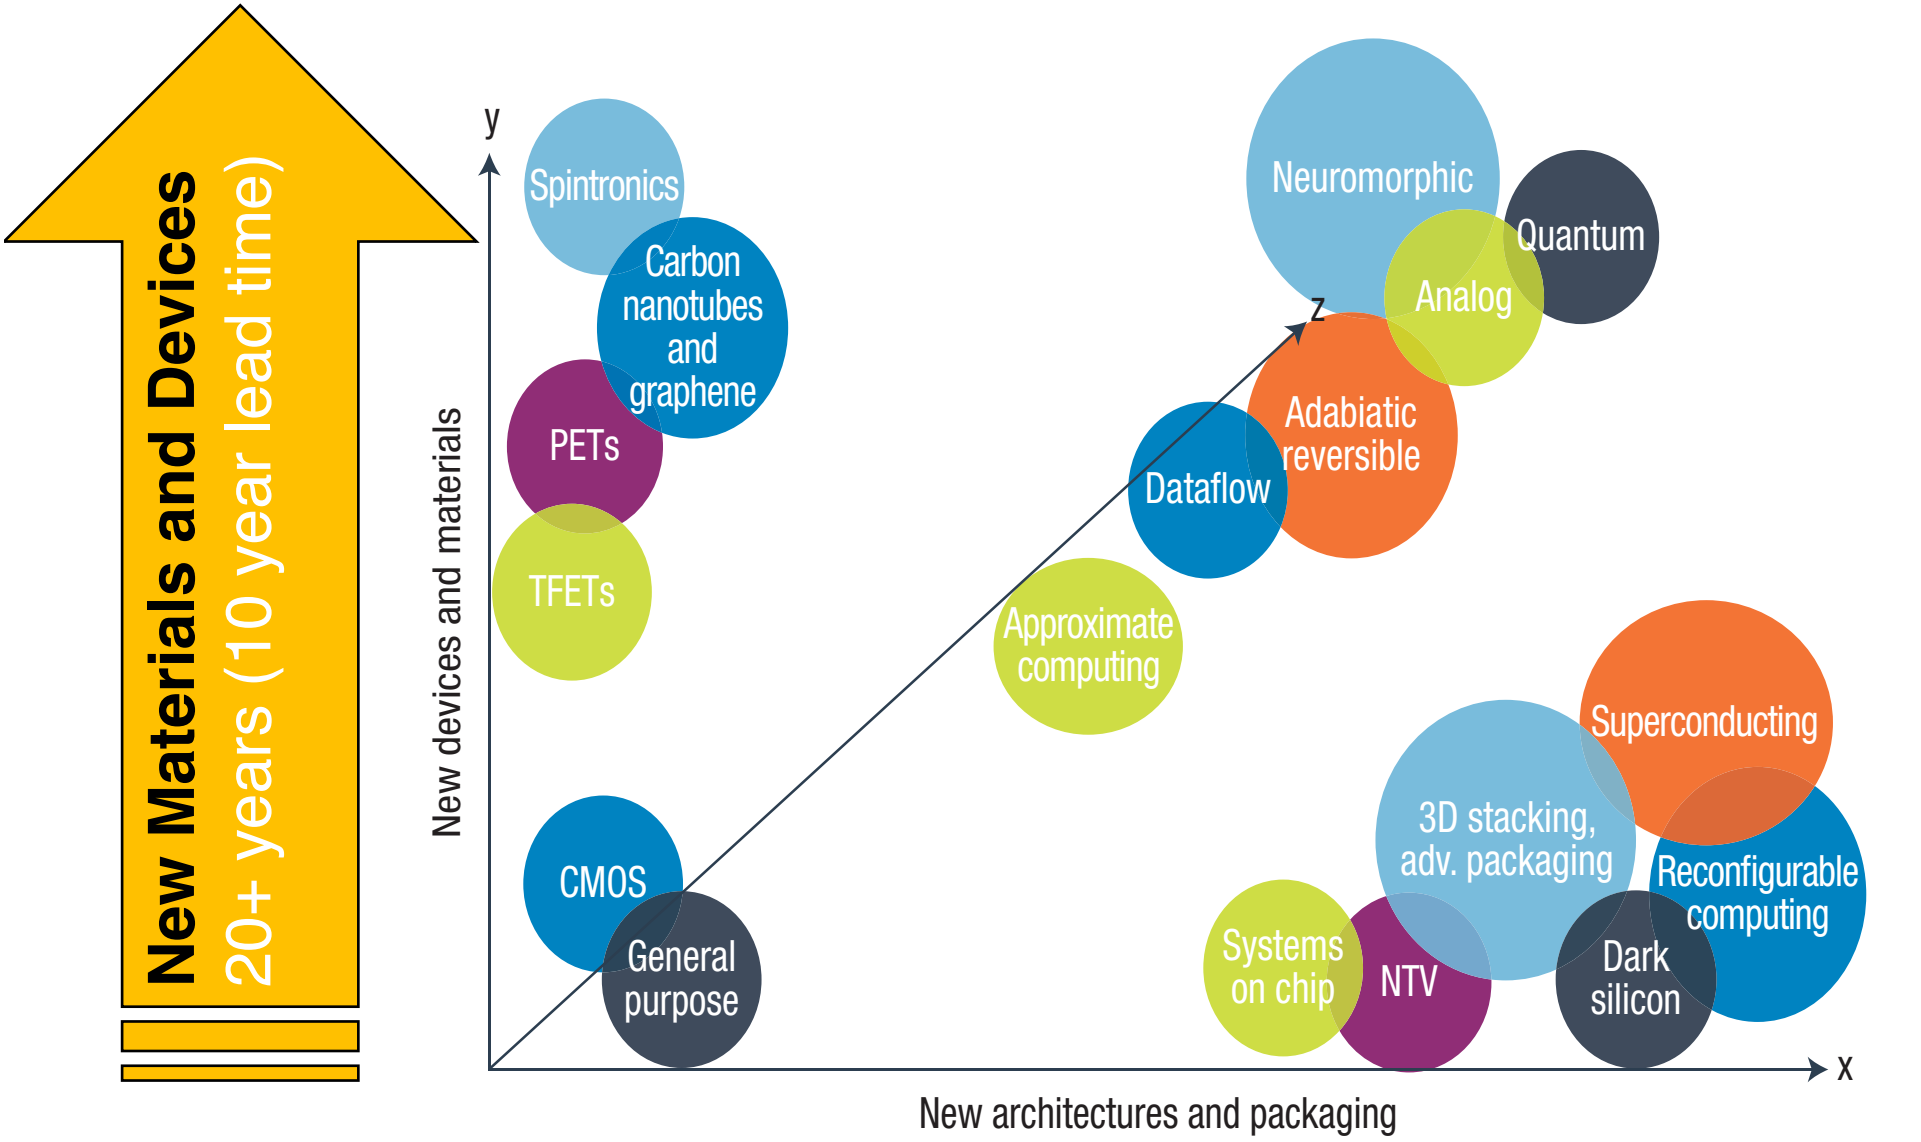

We might already be too late

Historically it is 10 years from lab to Fab…

But lets talk about it anyways.

Borkar-Shalf Criteria for New Device Technology

1. Gain
2. Signal to Noise
3. Scalability
4. Manufacturability

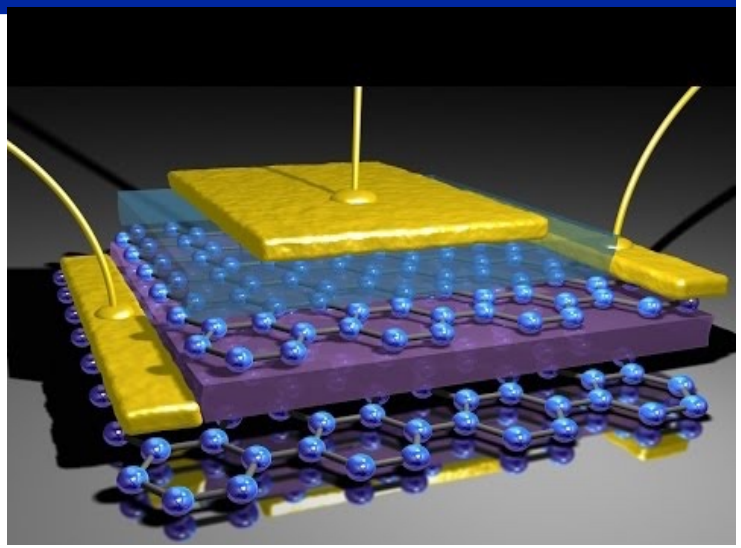

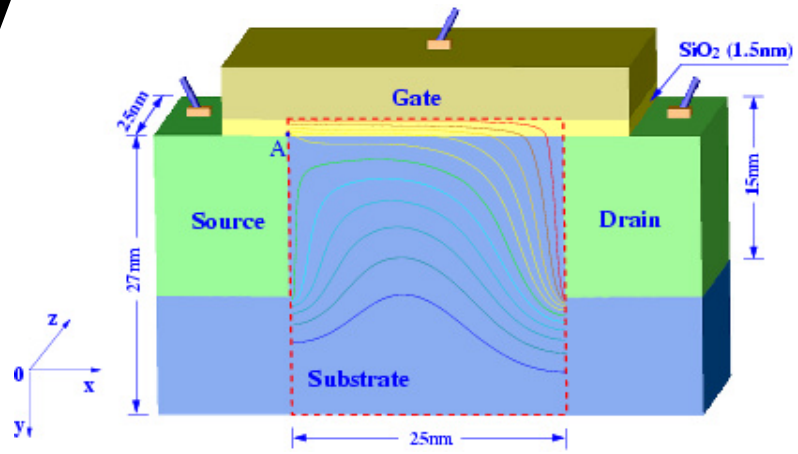

Alternatives to Conventional MOS Switches

(all require lower clock rate, and much more parallelism)

Energy-Performance Comparison

(30-stage fanout-4 inverter chains)

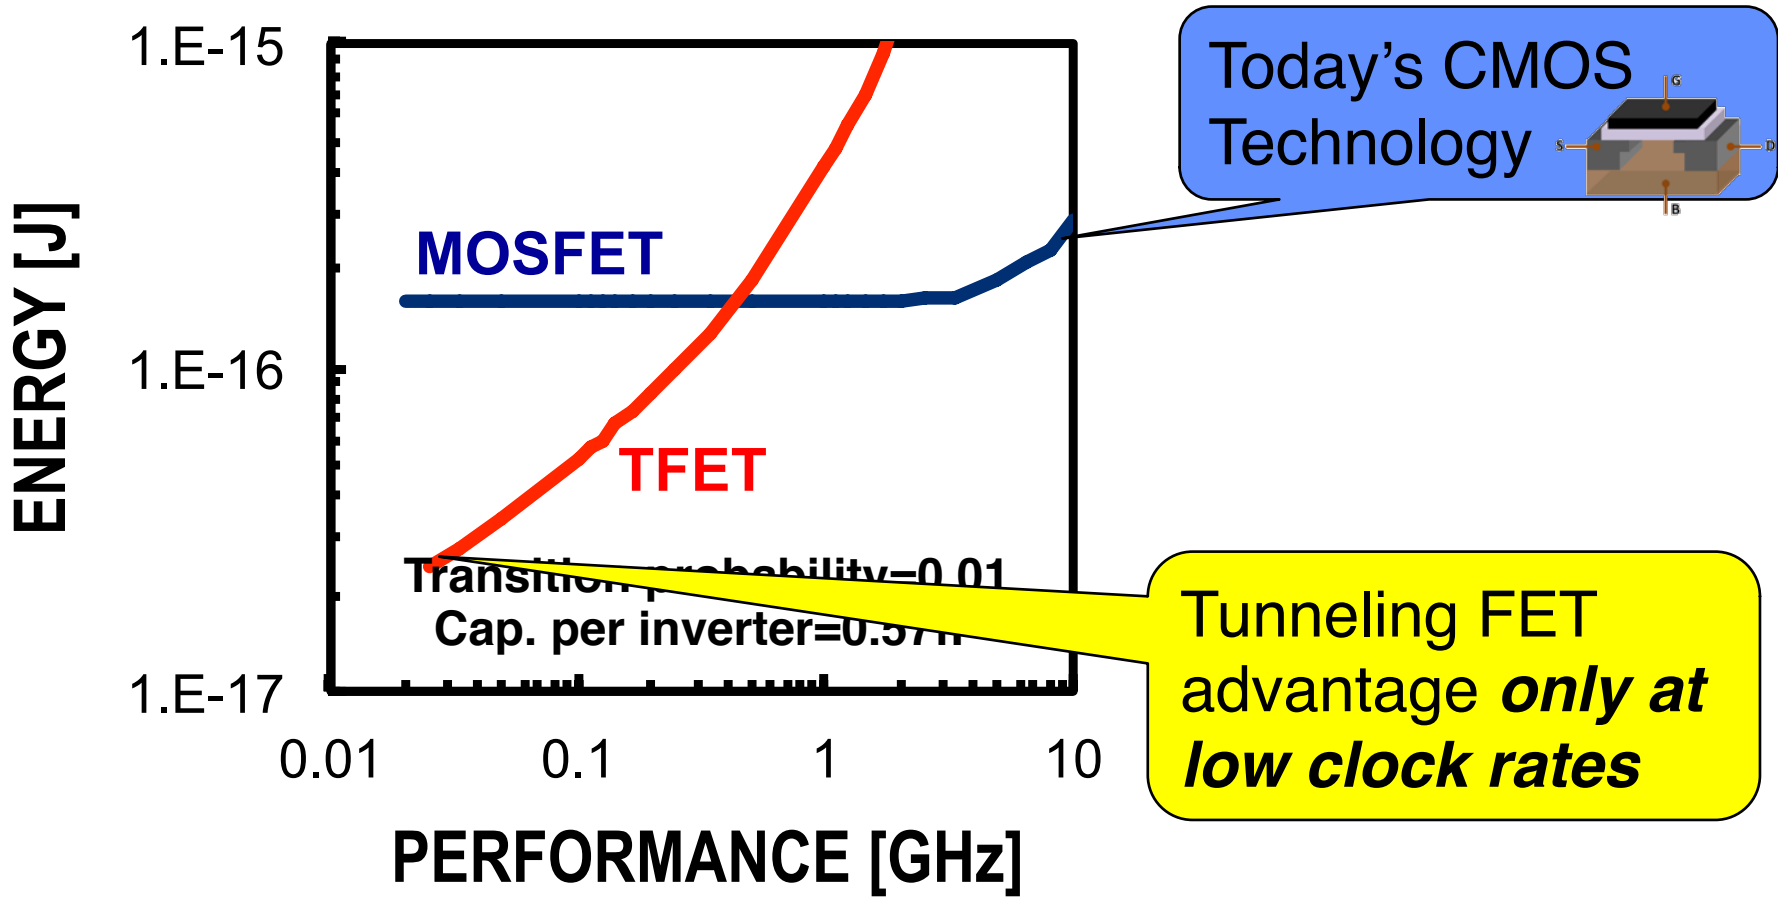

Systems, Packaging and Architecture

Now and Intermediate term: 3D Stacking and Advanced Packaging

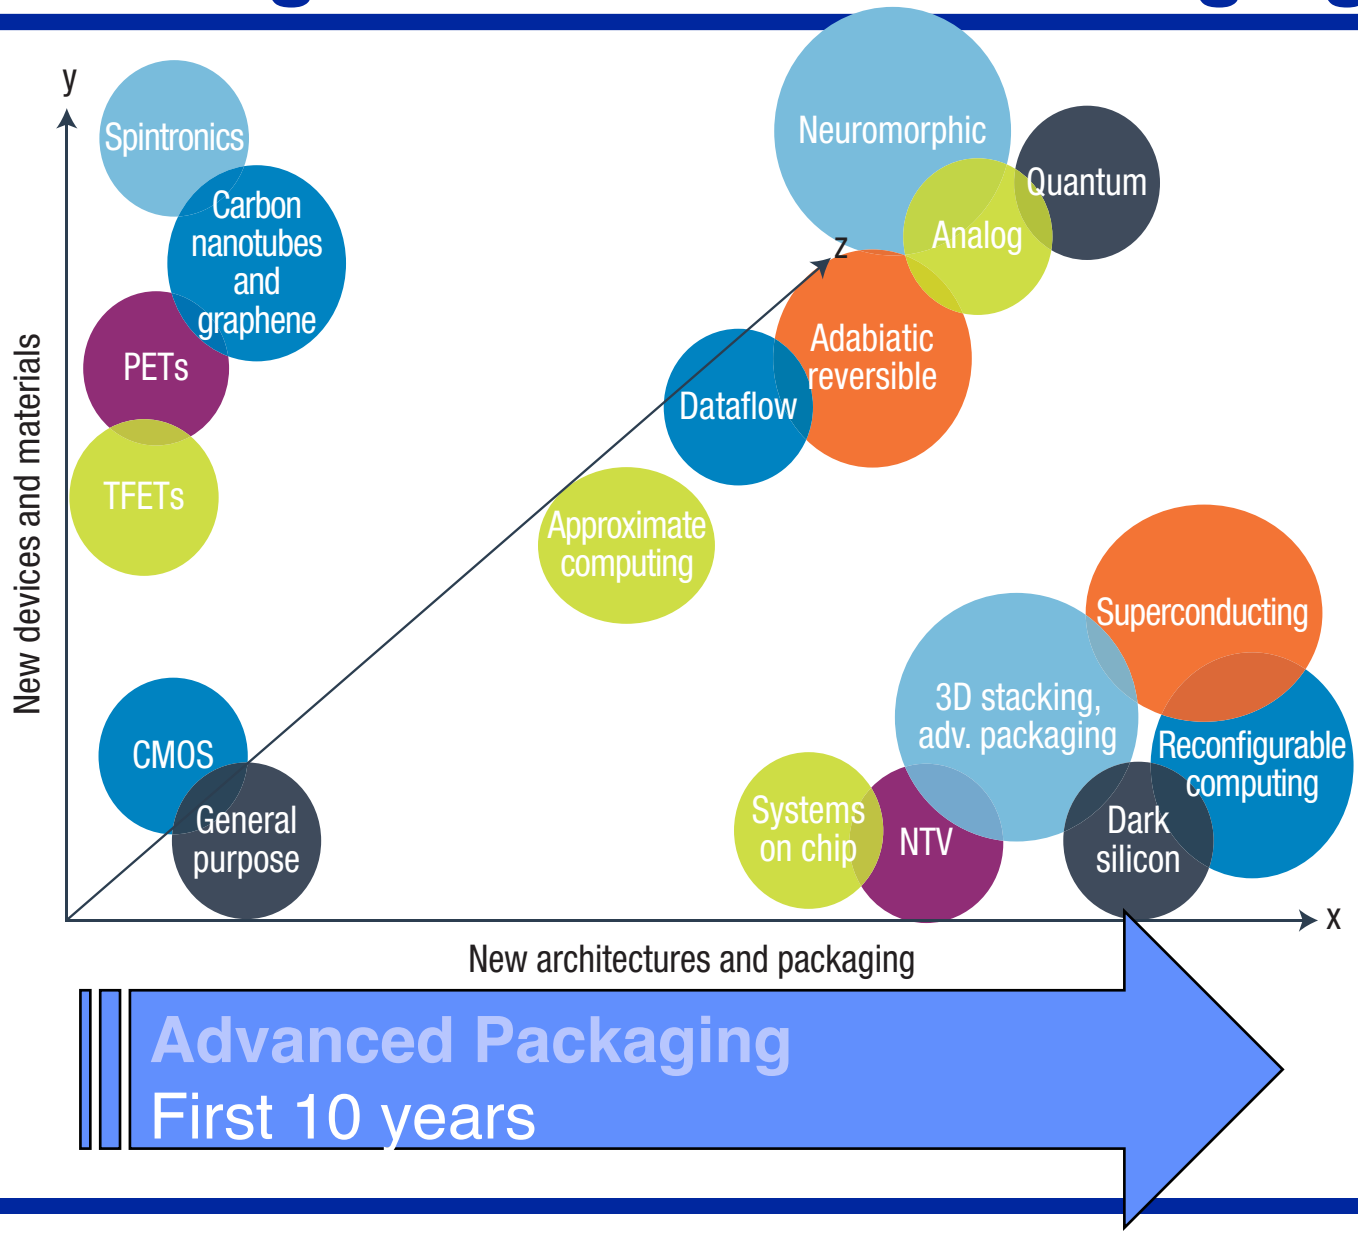

R&D in manufacturing at 2-nm node

CXRO EUV Test Facility at LBNL

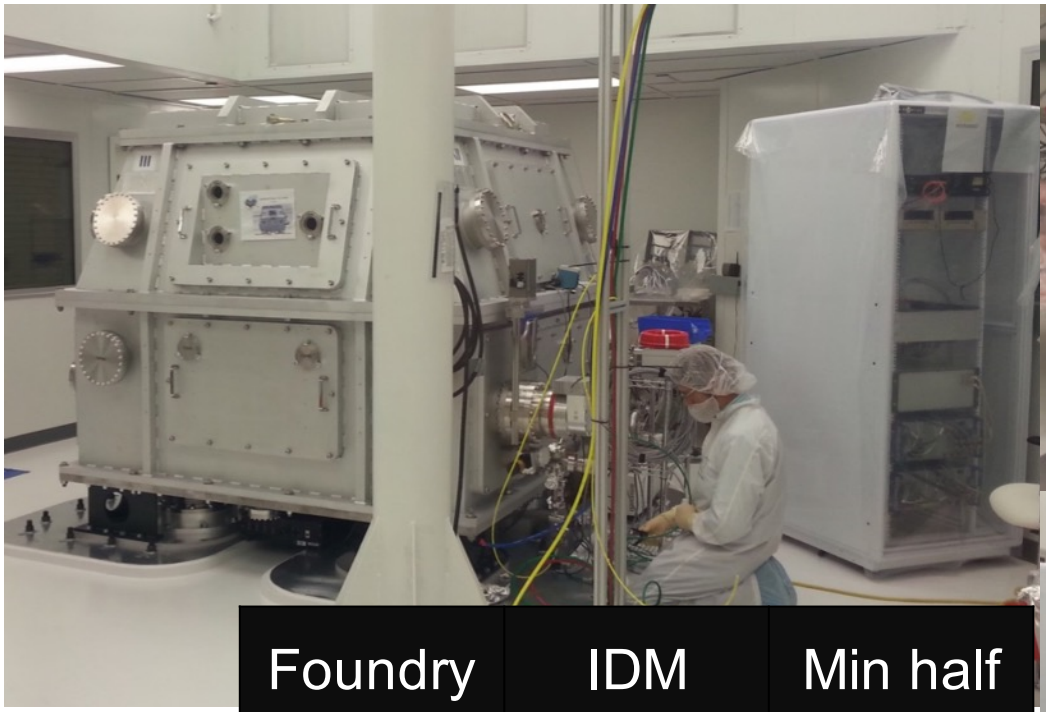

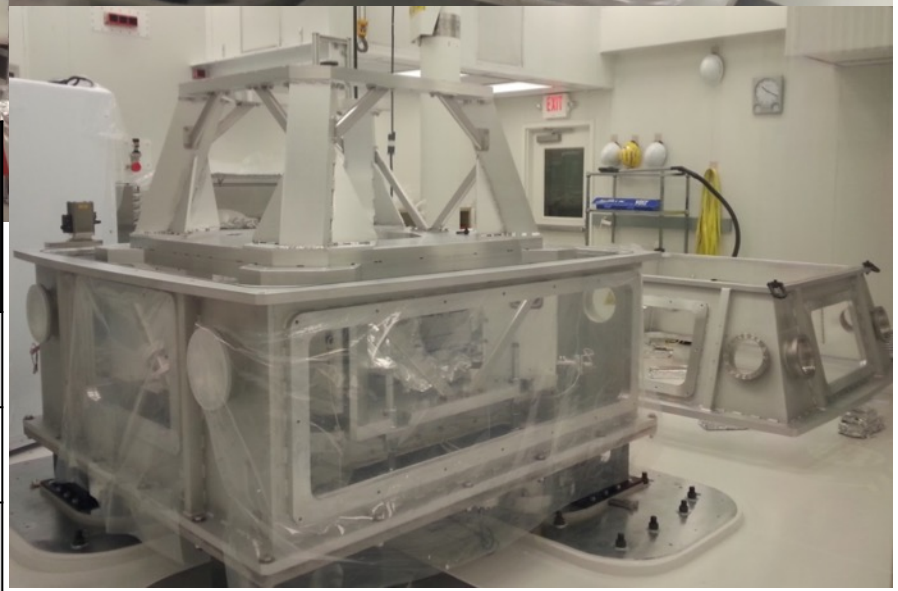

Foundry Node	IDM Node	Min half pitch
7 nm	10 nm	22 nm
5 nm	7 nm	16 nm
3 nm	5 nm	12 nm

Choose to Scale Something Else

(The future of Moore's Law Might not be about logic density…)

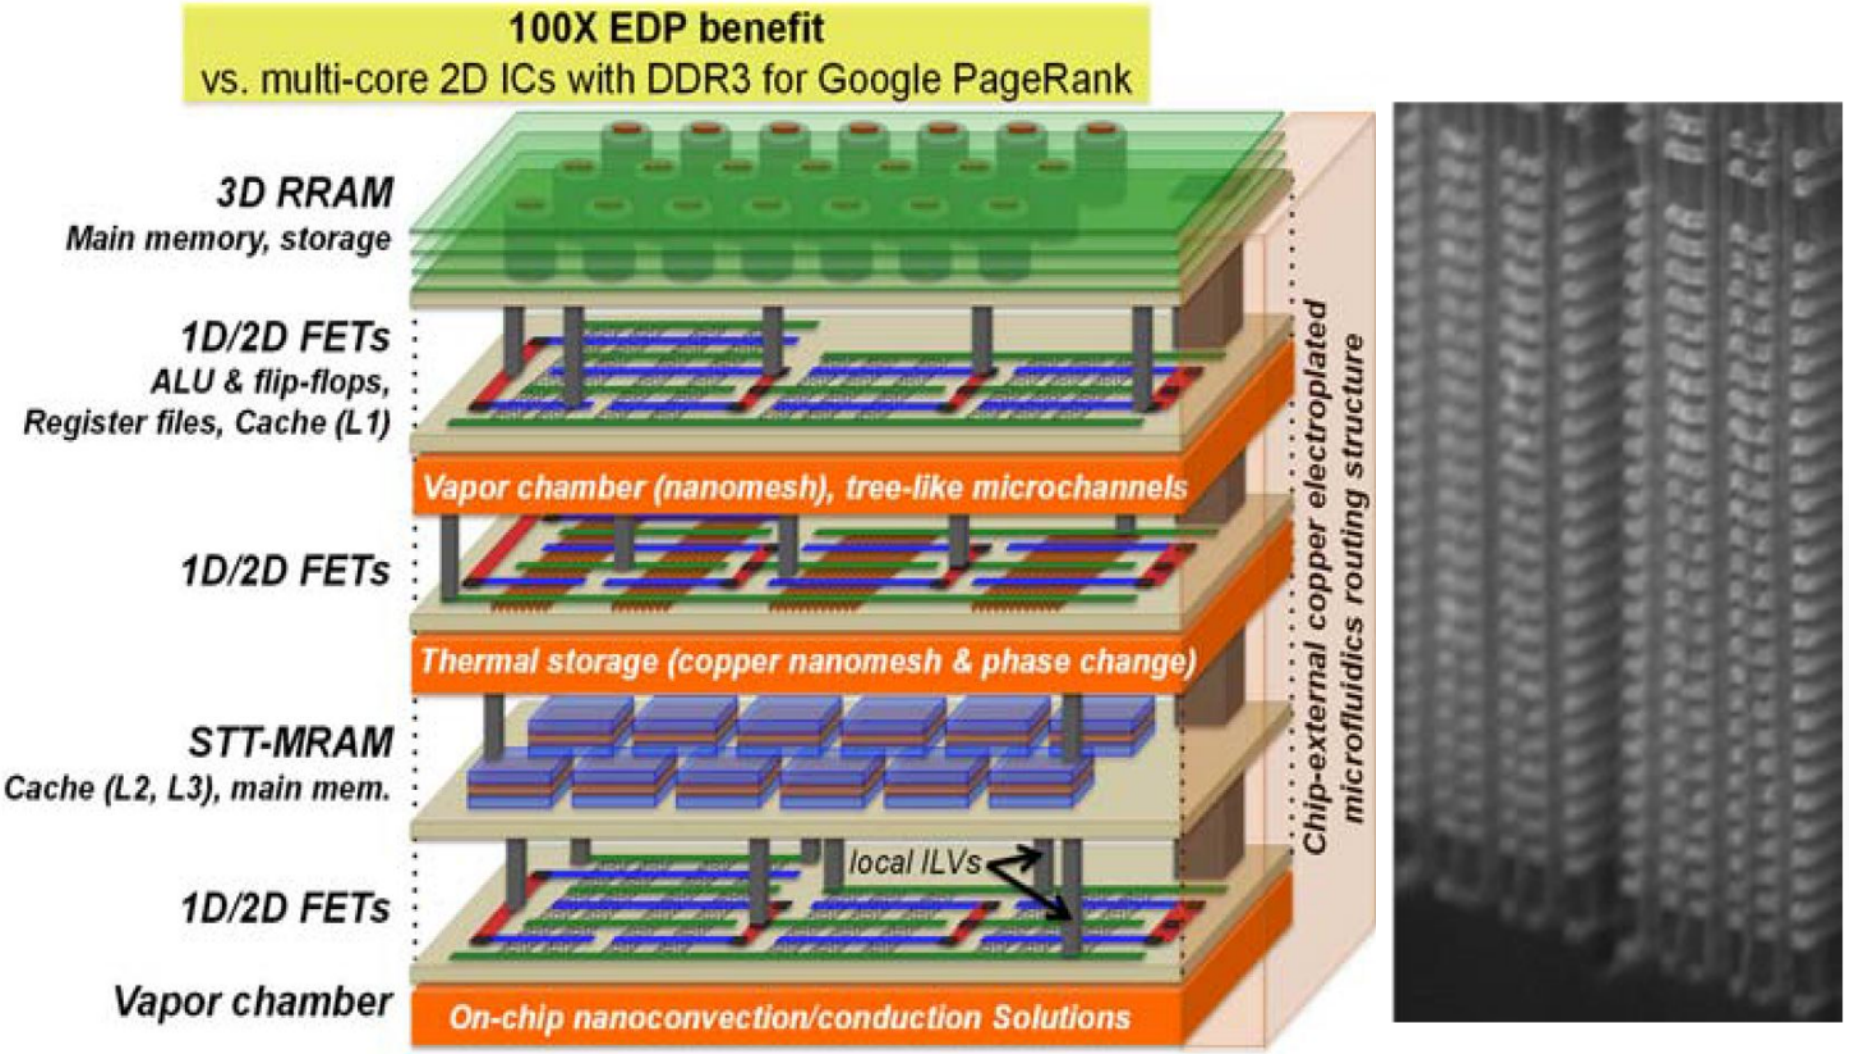

Stanford N3XT

Increase Logic Density and Efficiency using Specialization

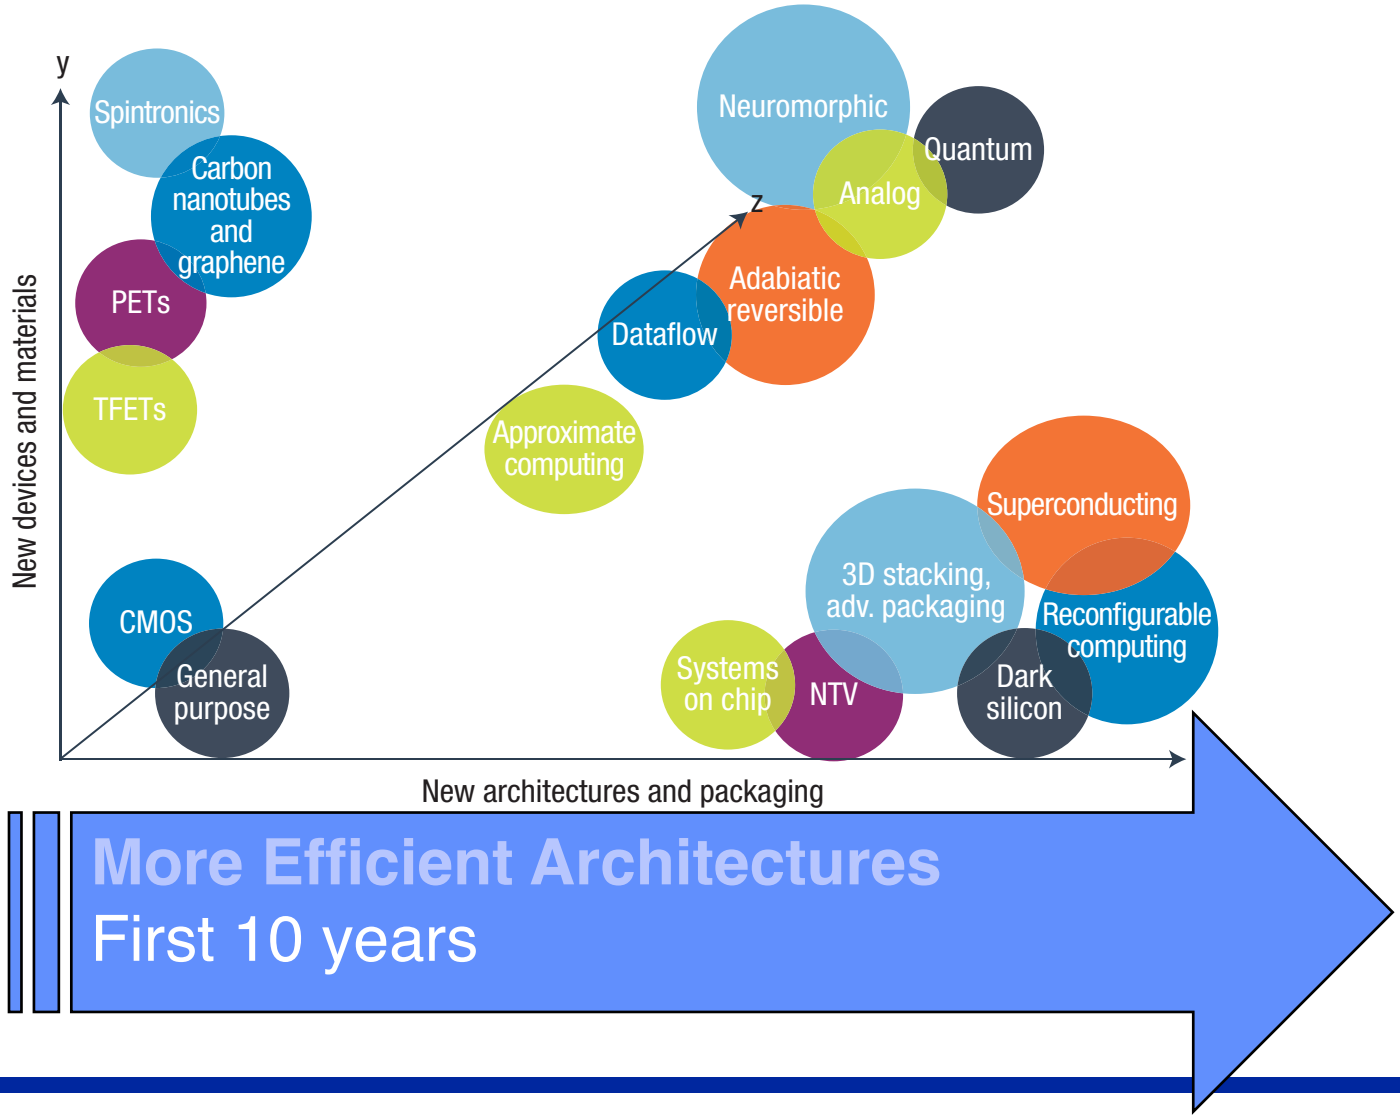

Current Architectures are Wasteful

(how far can we push architecture scaling using specialization?)

Must Dramatically Reduce the cost of creating specialized hardware!

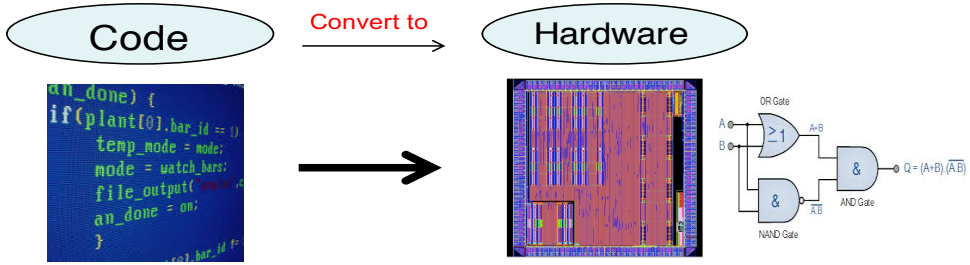

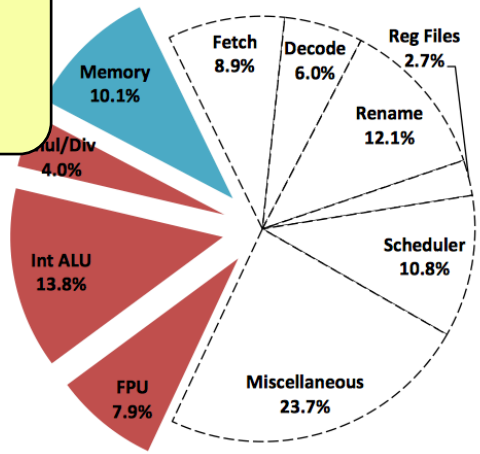

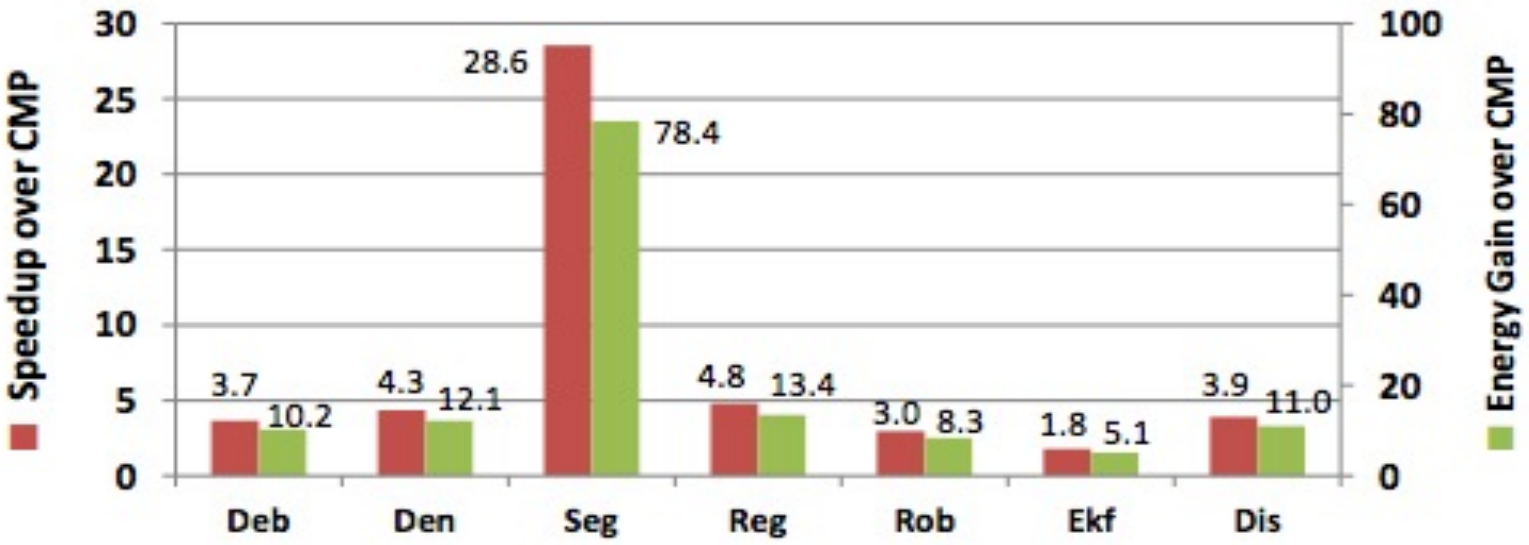

Need to Accelerate Pace of Discovery for Advanced Architectures

Simulate Architectures/Devices
(Chisel, OpenSOC, GEM5)

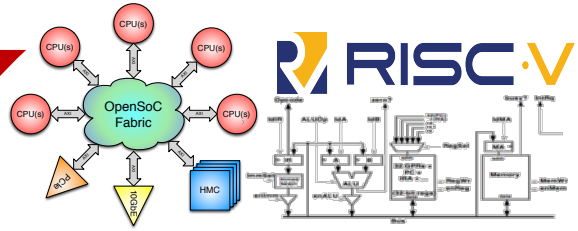

Synthesize/Prototype
(CXRO, Nanofab)

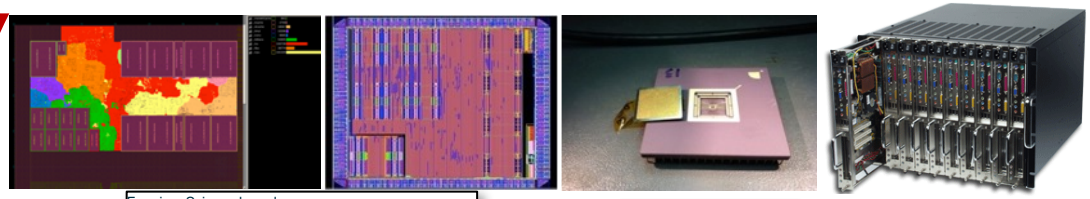

Characterize
(Benchmark, analyze apps.)

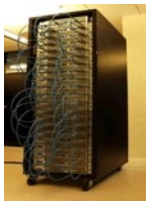

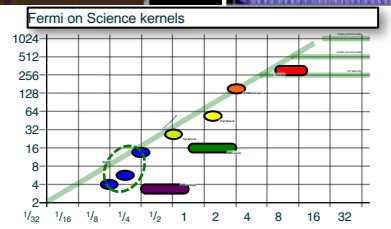

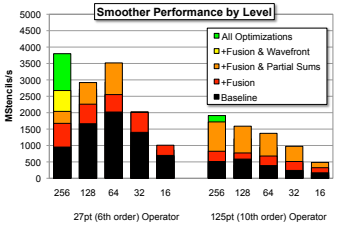

Performance models, e.g., roofline

Integrate
(application)

Feedback

Connection to Next Layer

Open Hardware for Flexible SoCs
(Synthesis & Simulation)

Chisel

DSL for rapid prototyping of circuits, systems, and arch simulator components

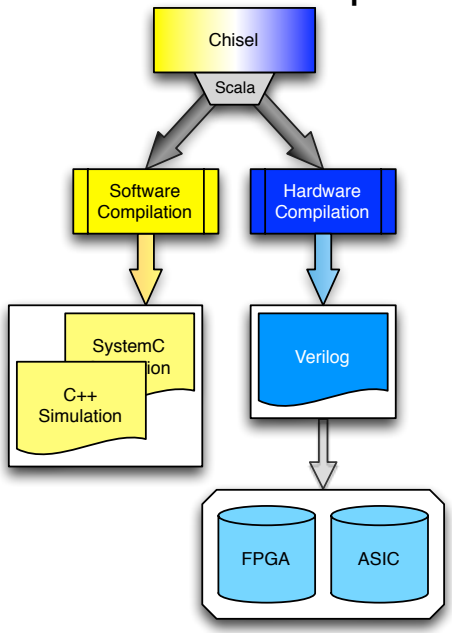

Back-end to synthesize HW with different devices Or new logic families

RISC-V

Open Source Extensible ISA/Cores

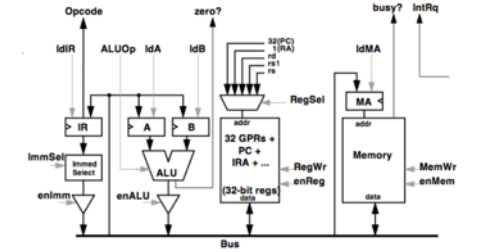

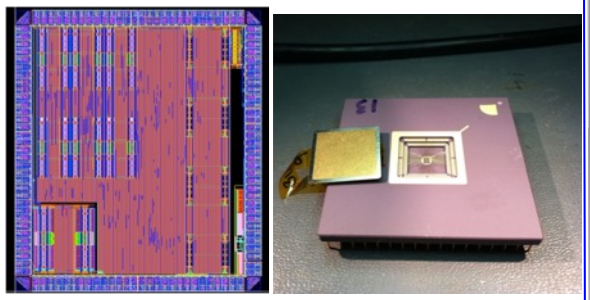

Re-implement processor With different devices or Extend w/accelerators

OpenSOC

Open Source fabric To integrate accelerators And logic into SOC

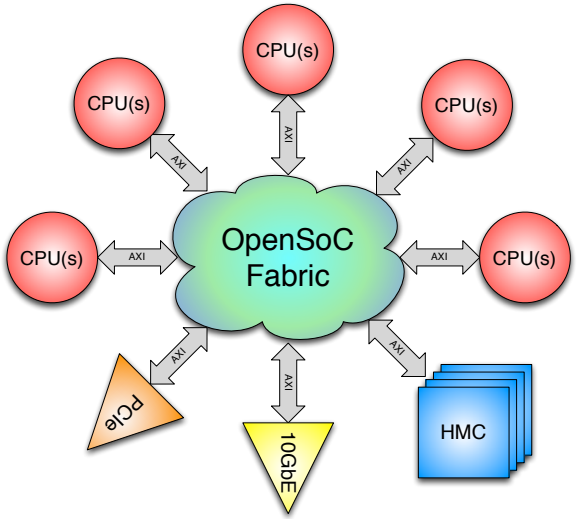

Platform for experimentation with specialization to extend Moore's Law

Combining Compact Device Models with Hardware Architectural Simulators

George Michelogiannakis
Dilip Vasudevan

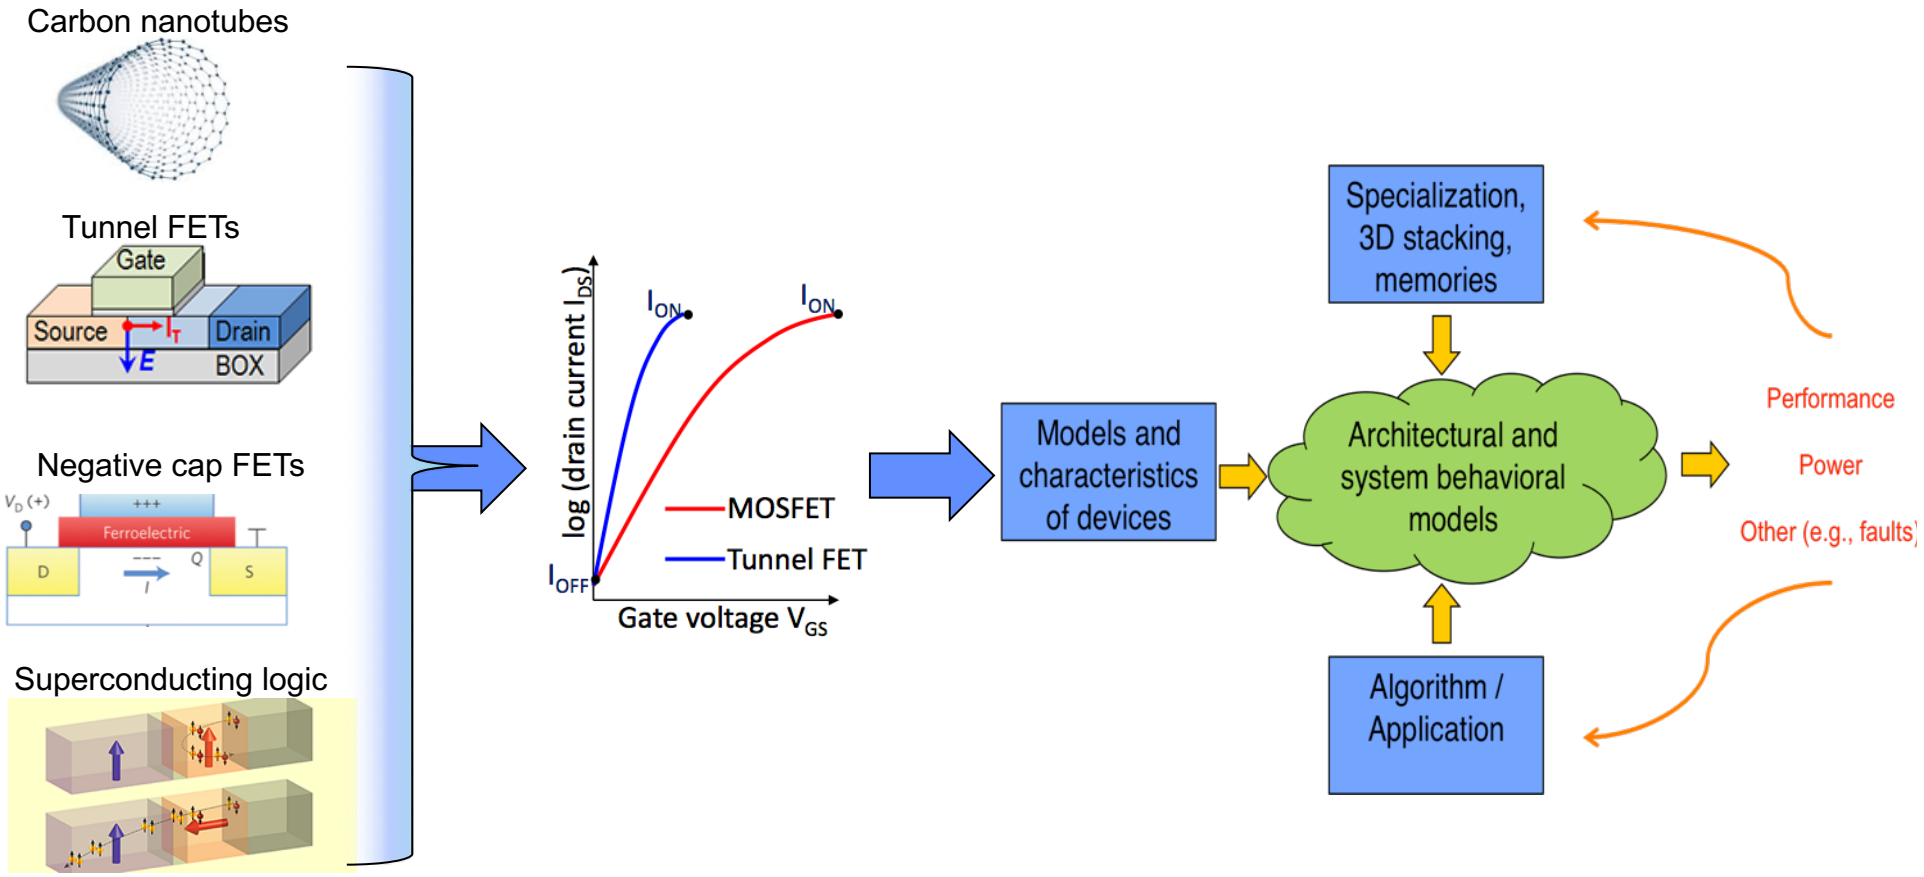

Holistic Electronic Device simulations

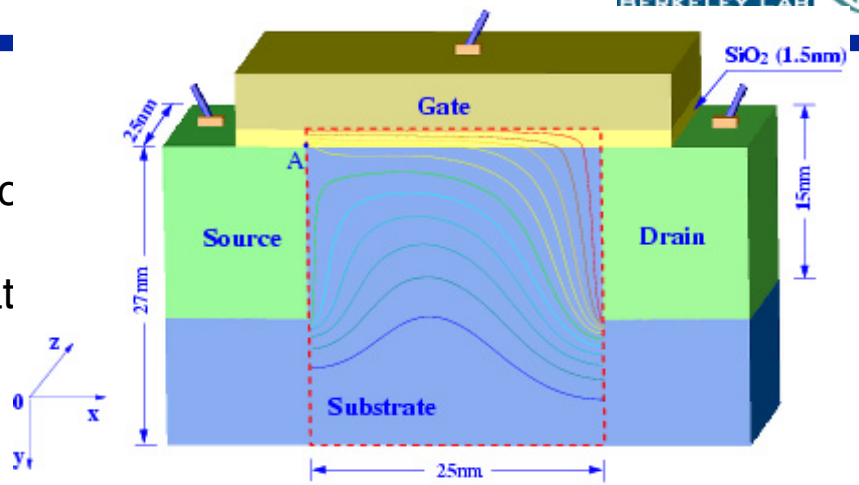

- Combined several techniques for a holistic, ab-initio, atomistic (beyond TCAT) device simulatic
- LS3DF Device-size selfconsistent ab initio calculat get atomistic potential profile, band alignment, based boundary conditioned Poisson solver
- Based on the potential profile, and scattering state calculations to simulate the device transport, and leakage current etc.
- Using electron-phonon coupling to calculate the heat generation and dissipation at atomic scale

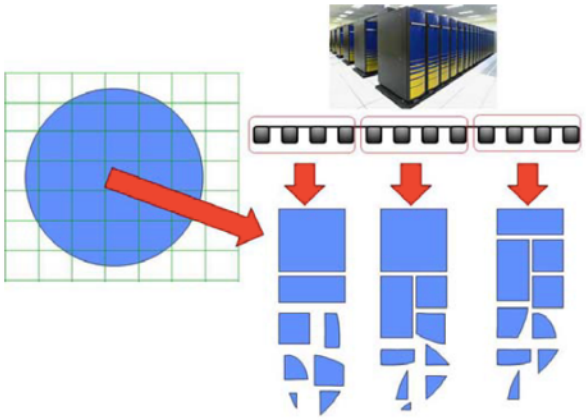

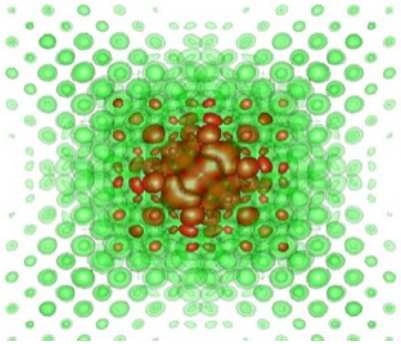

A shallow defect state in Si.

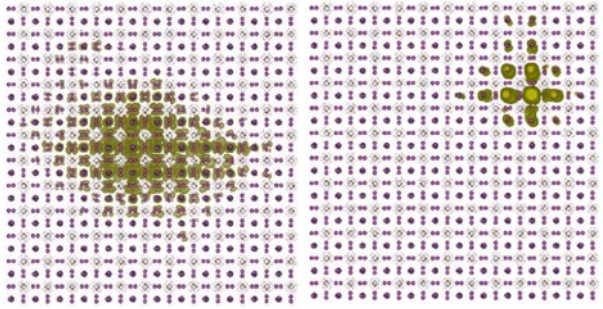

the electron (left) and hole (right) localizations in a bulk $CH_3NH_3PbI_3$ material. The small dots are atoms.

Beyond Moore Modeling Workflow

George Michelogiannakis
Dilip Vasudevan

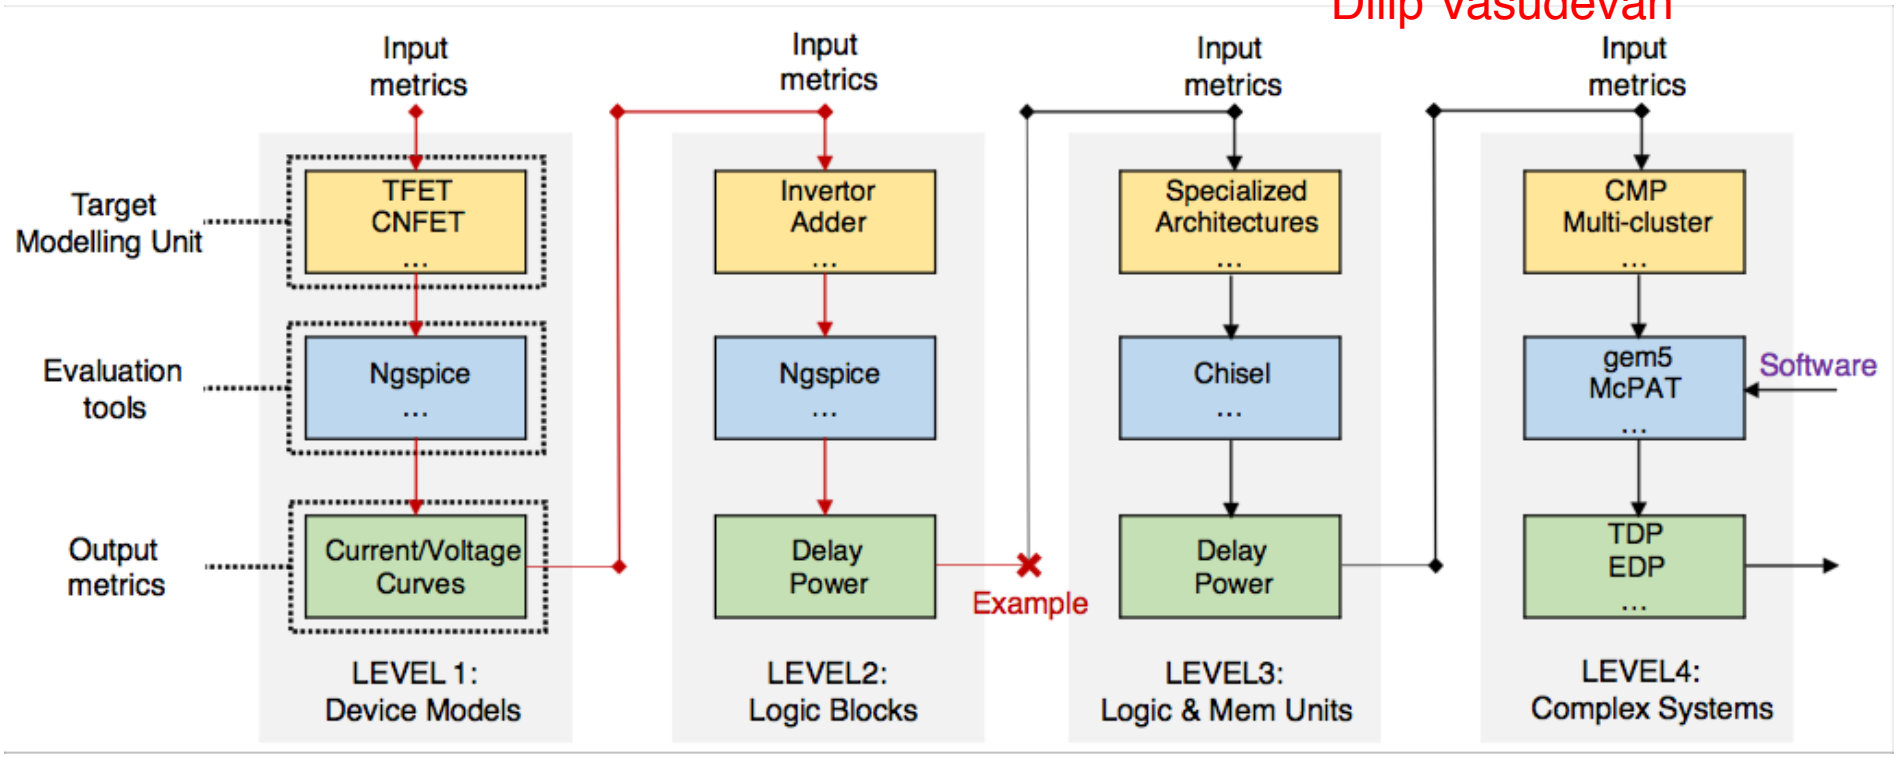

End-to-End Post-Moore Design Space Exploration Tool Flow

Incorporating Emerging Device Models into Architecture Simulation

- We incorporated TFET models into transistor-level netlists and implemented higher-level logic blocks
- Next steps:
  - NCFETs, CNFETs
  - More complicated logic blocks

George Michelogiannakis
Dilip Vasudevan

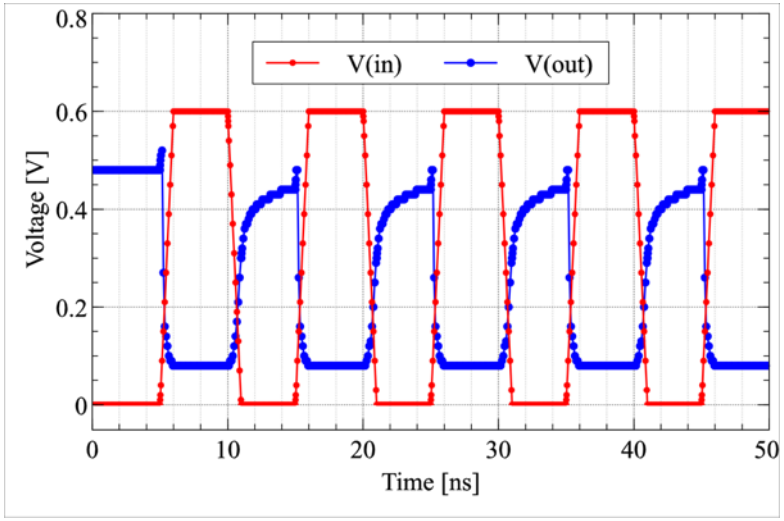

TFET Spice Simulation
- Inverter

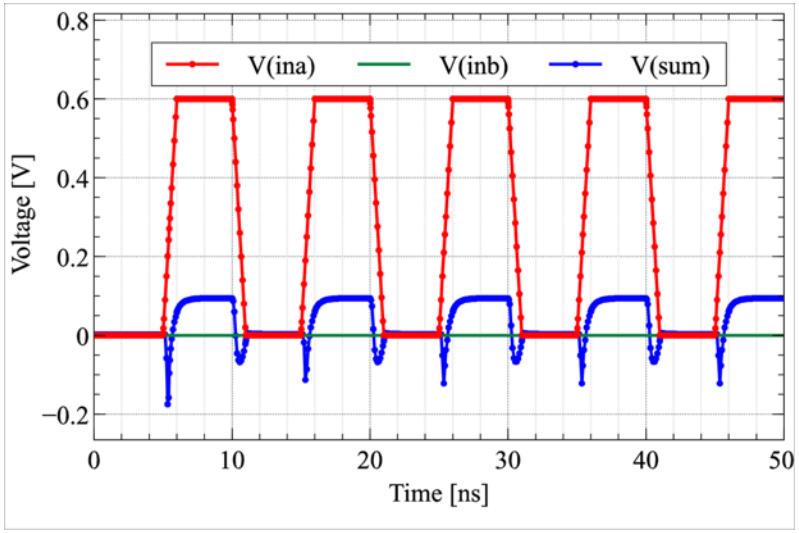

TFET Spice Simulation
- Adder

Design Architectures Around Design Patterns

7 Giants of Data (NRC)	7 Motifs of Simulation
Basic statistics	Monte Carlo methods
Generalized N-Body	Particle methods
Graph-theory	Unstructured meshes
Linear algebra	Dense Linear Algebra
Optimizations	Sparse Linear Algebra
Integrations	Spectral methods
Alignment	Structured Meshes

Identify common computational patterns

Organizing principles for Non-Von "Spatial Computing"

- Data Movement will remain a challenge even with exotic materials, but especially CMOS
- Copper is as good of a conductor as you can expect at room temperature
- With even lower power switches, challenges skews even more to data movement *(NEED Spatial Computing approach)*
- Push towards more parallelism (more tesselation of the memory structures)…. Strong Scaling

Strong Scaling extrapolates to ***limit case*** with **no separation of memory and compute** *(e.g. one PDE cell per processing element)*

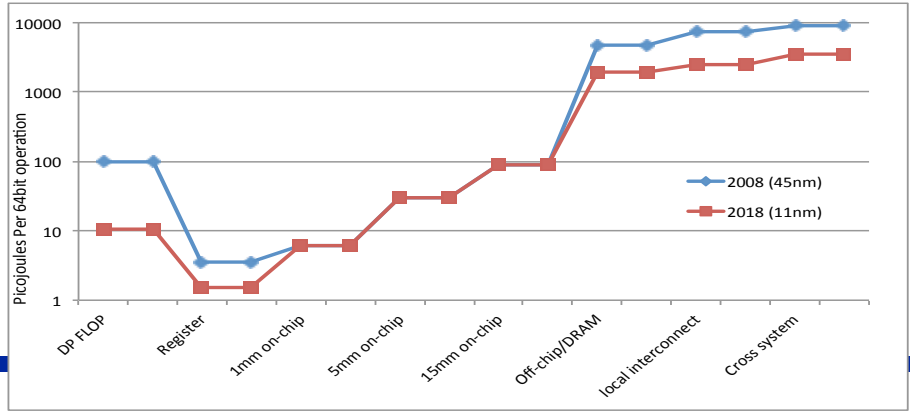

Concept: Solid State Virtual Fluid

Extreme (spatial) Specialization + New Devices + New programming models

PDEcell / PICcell: Ultra-simple compute engine (50k gates) calculates finite-difference updates, and particle forces from neighbors. Microinstructions specify the PDE equation, stencil, and PIC operators. ***Novel features:*** *variable length streaming integer arithmetic and novel PIC particle virtualization scheme.*

Computational Lattice: PDECells are tiles in a lattice/array on each 2D planar chip layer. Target 120x120 tiles per mm^2 @28nm lithography. Novel Features: each tile represents single cell of computational domain (pushes to limit of strong-scaling).

Monolithic 3D Integration: Integrate layers of compute elements using emerging monolithic 3D chip stacking.
Novel Features: *1000 layer stacking (20x more than current practice). Area efficient inter-layer connectivity and new energy efficient transistor logic (ncFET).*
1 Petaflop equivalent performance in 300mm^2 for < 200Watts.

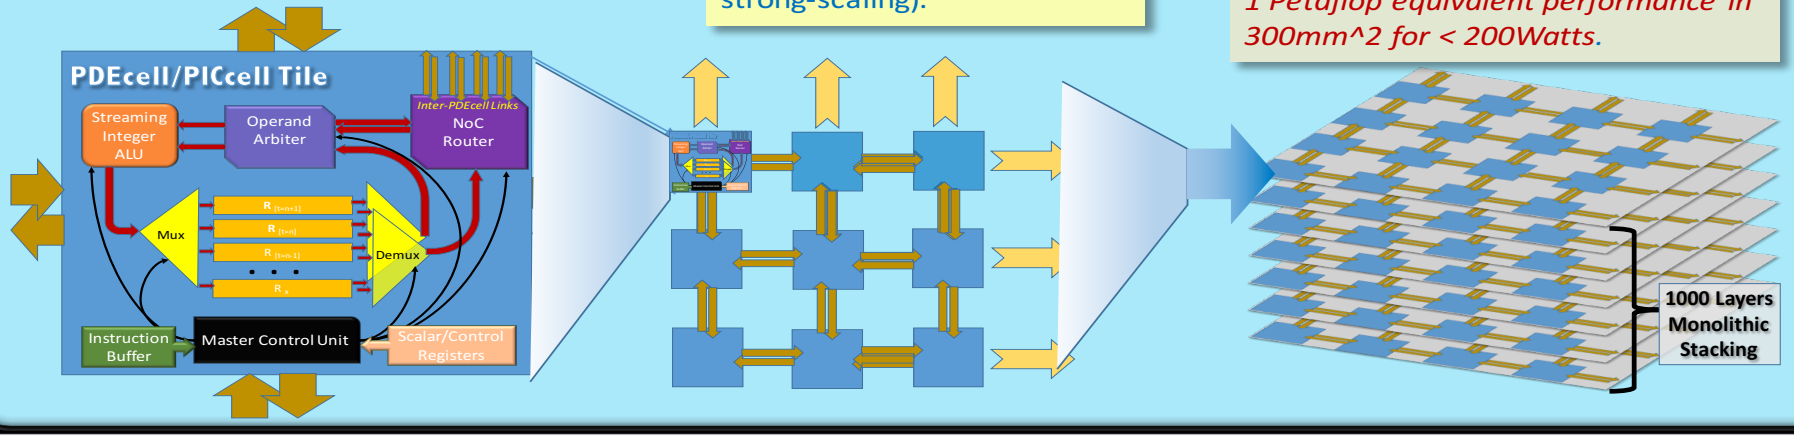

Scalar waves in 3D are solutions of the hyperbolic wave equation: $-\phi_{,tt} + \phi_{,xx} + \phi_{,yy} + \phi_{,zz} = 0$

Initial value problem: given data for ϕ and its first time derivative at initial time, the wave equation says how it evolves with time

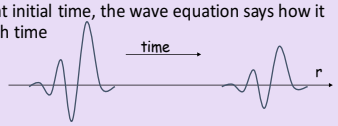

Discretized PDE Representation in DSL

$$\phi^{n+1}_{i,j,k} = 2\phi^{n}_{i,j,k} - \phi^{n-1}_{i,j,k}$$
$$+ \Delta t^2/\Delta x^2(\phi^{n}_{i+1,j,k} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i-1,j,k})$$
$$+ \Delta t^2/\Delta y^2(\phi^{n}_{i,j+1,k} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i,j-1,k})$$
$$+ \Delta t^2/\Delta z^2(\phi^{n}_{i,j,k+1} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i,j,k-1})$$

Compiles to MicroOps

```
R[n+1](0,0,0) = 0
R[n+1](0,0,0) += 2*R[n](0,0,0)
R[n+1](0,0,0) -= R[n-1](0,0,0)
R[n+1](0,0,0) += C * R[n+1](+1,0,0)
R[n+1](0,0,0) -= C * 2 * R[n](0,0,0)
R[n+1](0,0,0) += C * R[n](-1,0,0)
R[n+1](0,0,0) += C * R[n+1](0,+1,0)
    ....
```

Executes in Wavefronts

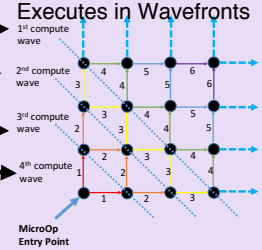

Programming Model Challenges for Non-VonNeumann & Specialized Architectures

Von Neumann ~= Instruction Processor

```
int main()
{
   int n = 0;
   while(n < 100)
   {
      n = n + 5;
      print("n = %d\n", n);
      pause(200);
      if(n == 50) break;
   }
   print("All done!");
}
```

What the heck is this?

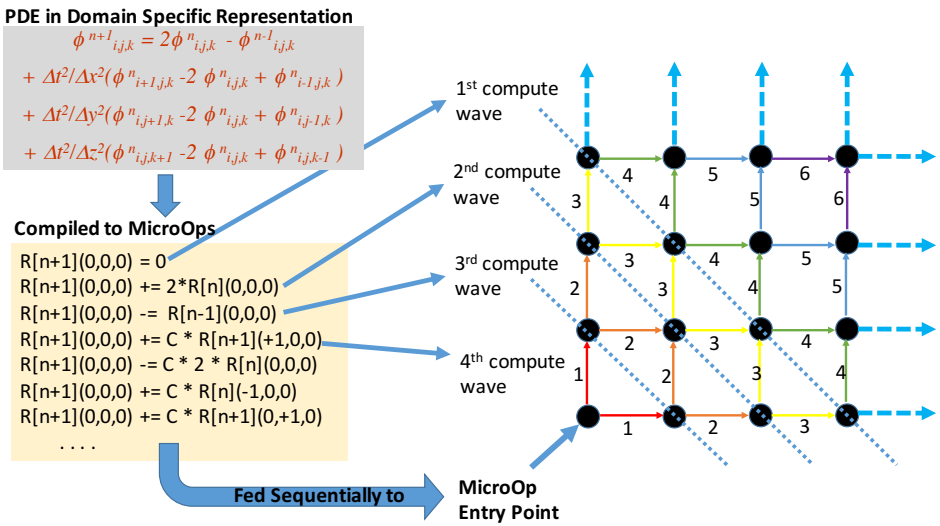

Modern languages (including many classes of DSLs Were designed with instruction processors in mind

A Framework for Accelerated Technology Development Beyond Moore's Law

Drive Focus and Impact via a Multiscale CoDesign Framework

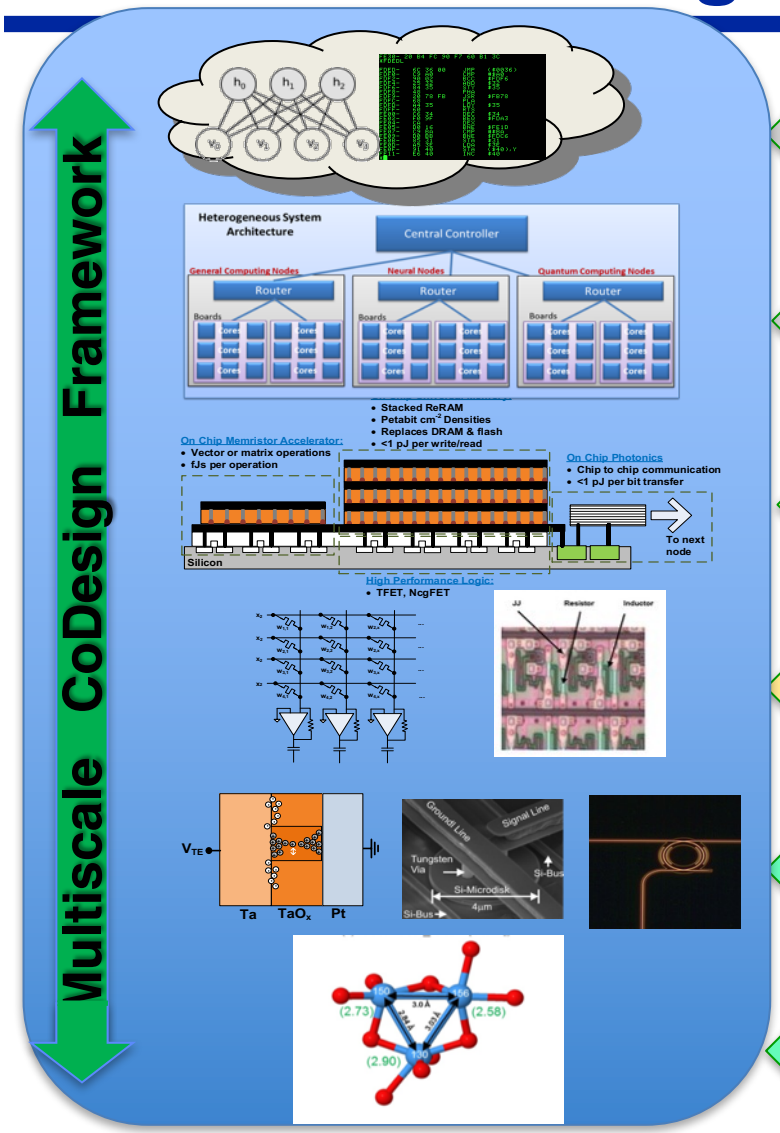

- Programming paradigms & Apps
- System architecture modeling
- Component hetero-integration
- Device demonstration
- Device physics
- Fundamental material science

Berkeley Lab · Argonne National Laboratory · Oak Ridge National Laboratory · Los Alamos National Laboratory · Lawrence Livermore National Laboratory · Pacific Northwest National Laboratory · The Ames Laboratory · Sandia National Laboratories

Reducing Solution Space

TR1 ideas and concepts academia, labs, industry

Input from academic and industrial partners

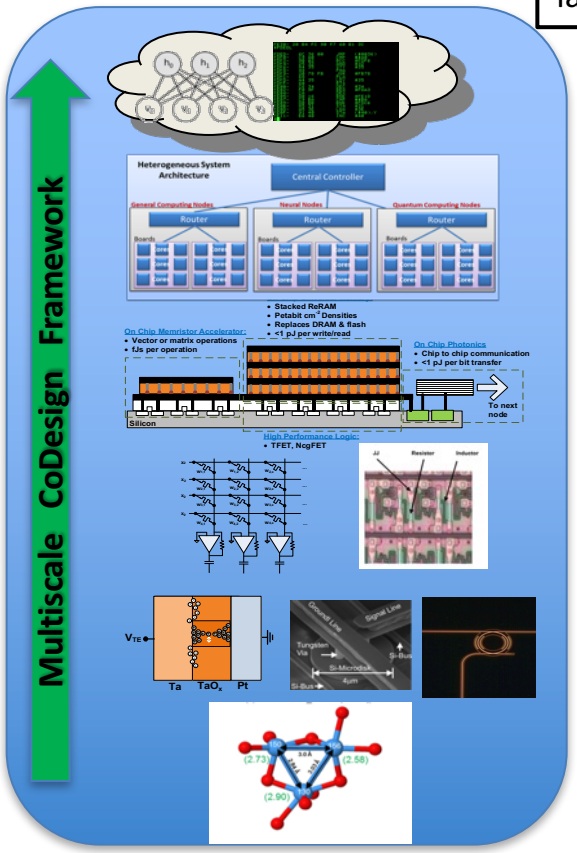

Software and algorithms

Architecture

Devices

Materials

Prototyping and testing

scaling

acceptance

Cross-functional assessments, deployment assessment

Multi-scale modeling and simulations, multi-modal characterization – reduce options

Beyond Moore Codesign Framework

Modeling

Application Performance Modeling

10,000x improvement:
20 fJ per instruction equivalent

Systems

Computer System Architecture Modeling
- Next generation of Structural Simulation Toolkit
- Heterogeneous systems HPC models

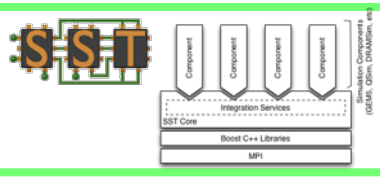

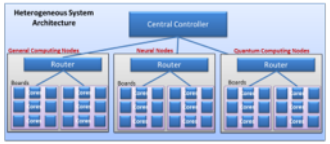

Component Level Models
- Gem5, McPAT, HiHAT, CACTI, NVSIM

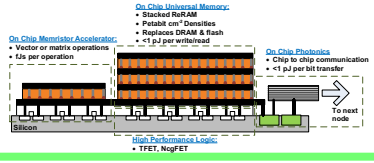

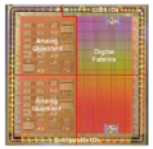

Component Fabrication
- Processors, ASICs
- Photonics
- Memory

Circuit/IP Block Design and Modeling
- SPICE/Xyce model

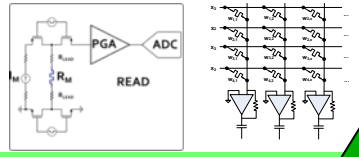

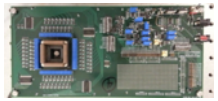

Test Circuit Fab and Measurement
- Subcircuit measurement

Compact Device Models
- Single device electrical models
- Variability and corner models

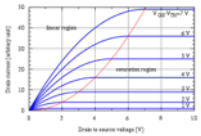

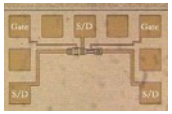

Devices

Device Measurements
- Single device electrical behavior
- Parametric variability

Device Physics Modeling
- Device physics modeling (TCAD)
- Electron transport, ion transport
- Magnetic properties

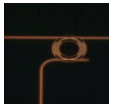

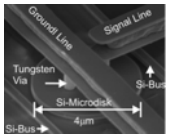

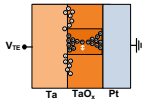

Device Structure Integration and Demonstration
- Novel device structure demonstration

Process Module Modeling
- Diffusion, etch, implant simulation
- EUV and novel lithography models

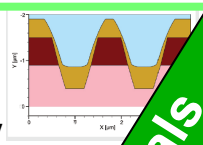

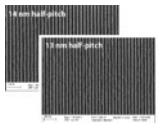

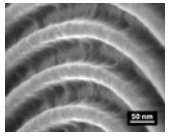

Process Module Demonstrations
- EUV and novel lithography
- Diffusion, etch, implant simulation

Materials

Atomistic and Ab-Initio Modeling
- DFT – VASP
- MD – LAAMPS

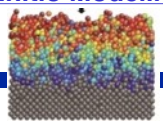

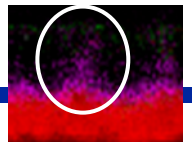

Fundamental Materials Science
- X-ray: XRD, XPS, HAXPES
- Microscopy: TEM, SEM, EELS
- Scanning probe technique

Experimental

Conclusion

- The end of lithography scaling as we know it is coming within a decade (*about when Exascale is done*)
- Neuromorphic and Quantum do not address this challenge
  - They expand computing to exciting new areas!
  - But do not *replace* Digital logic where
  - And ***all*** are affected by lithography challenge!
- But it need not mean the *end of Moore's Law*
  - We believe in ***More Moore!***
  - But it will require ***innovation***!
- Requires a LOT of lead time, so we must start today!

Extra

Scientific Computing on Non-Von Neumann Digital Electronics

Von Neumann Architecture

Von Neumann

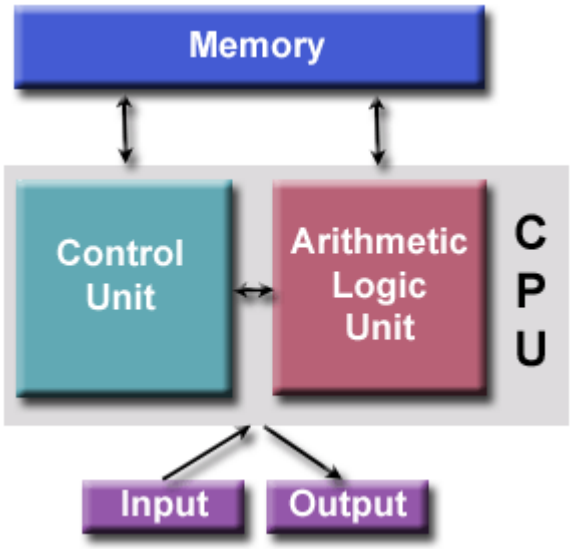

Emulates

Logic Array

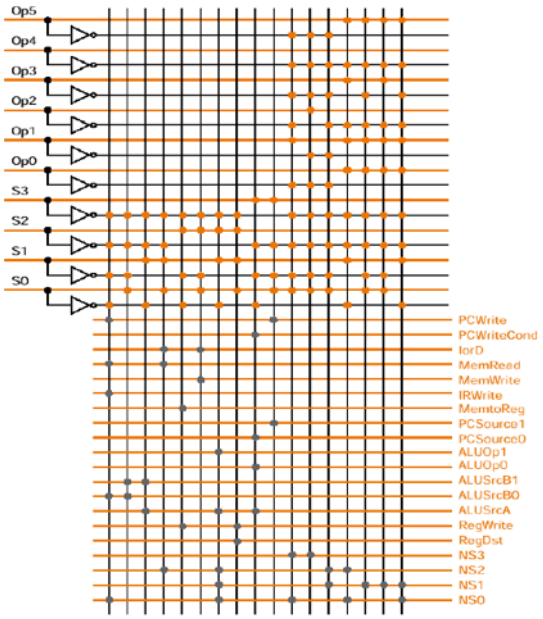

PIM is NOT Non-Von Neumann

Its just better packaging

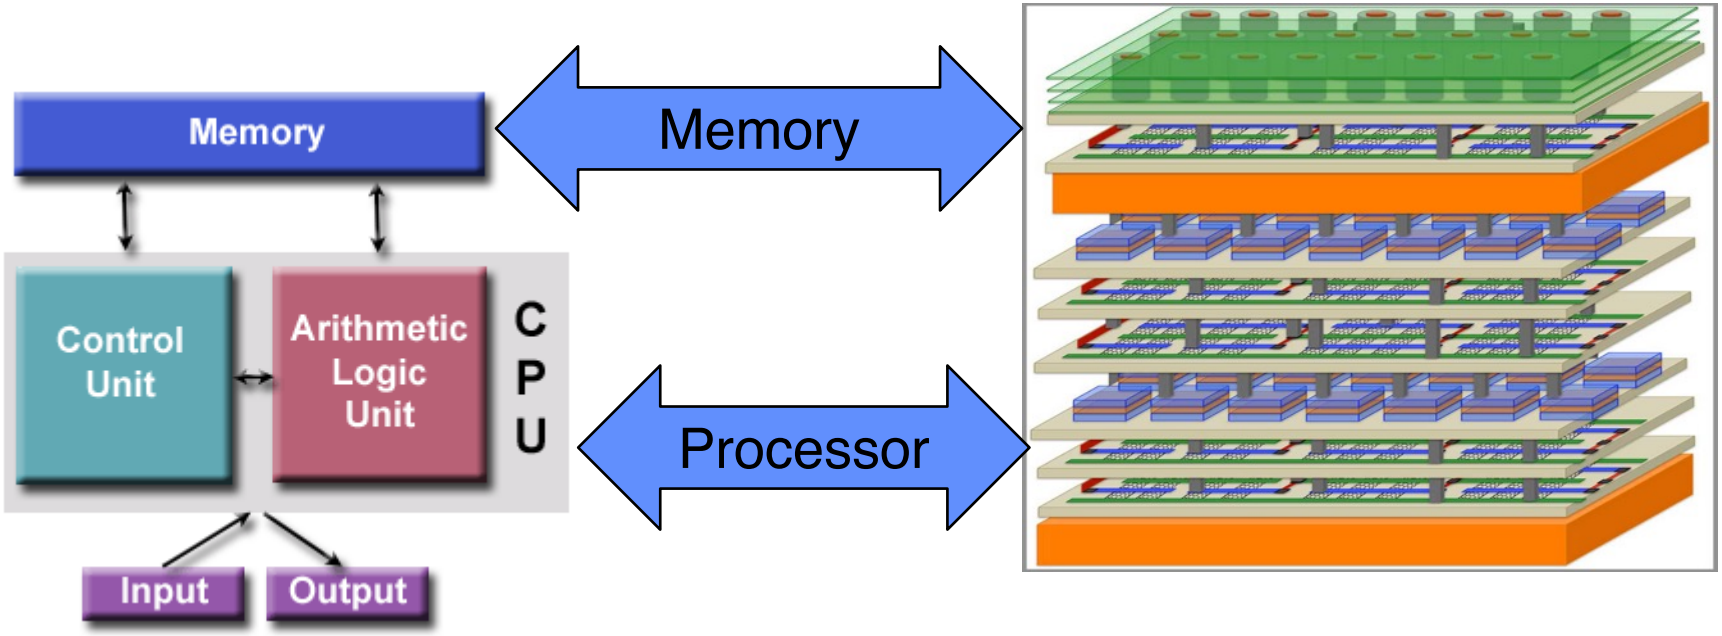

These ARE Non-Von Neumann

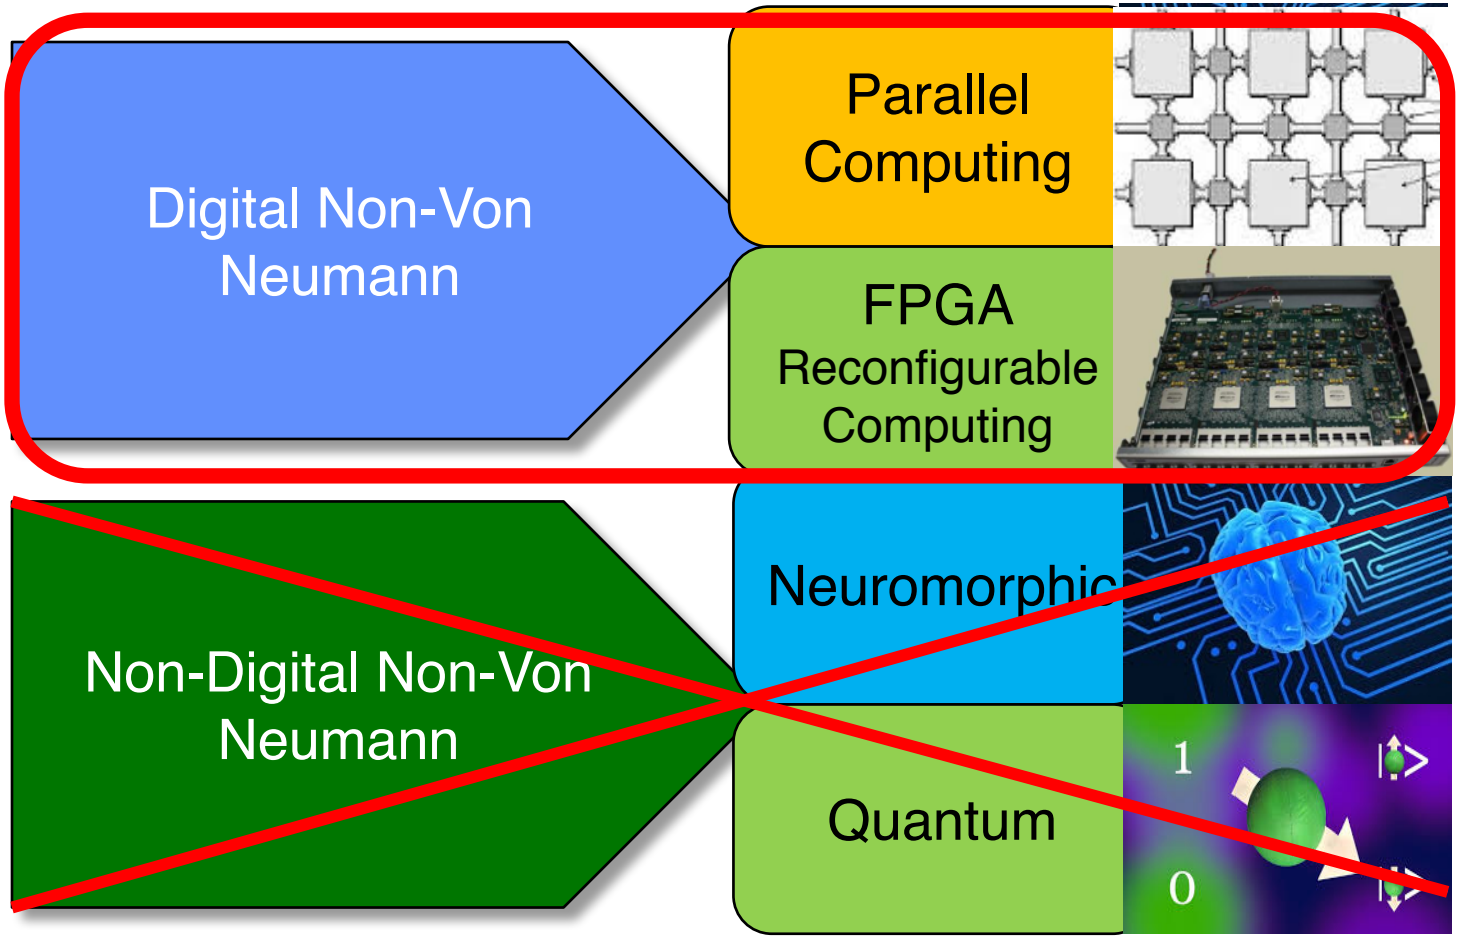

Cost of Data Movement Increasing Relative to Ops

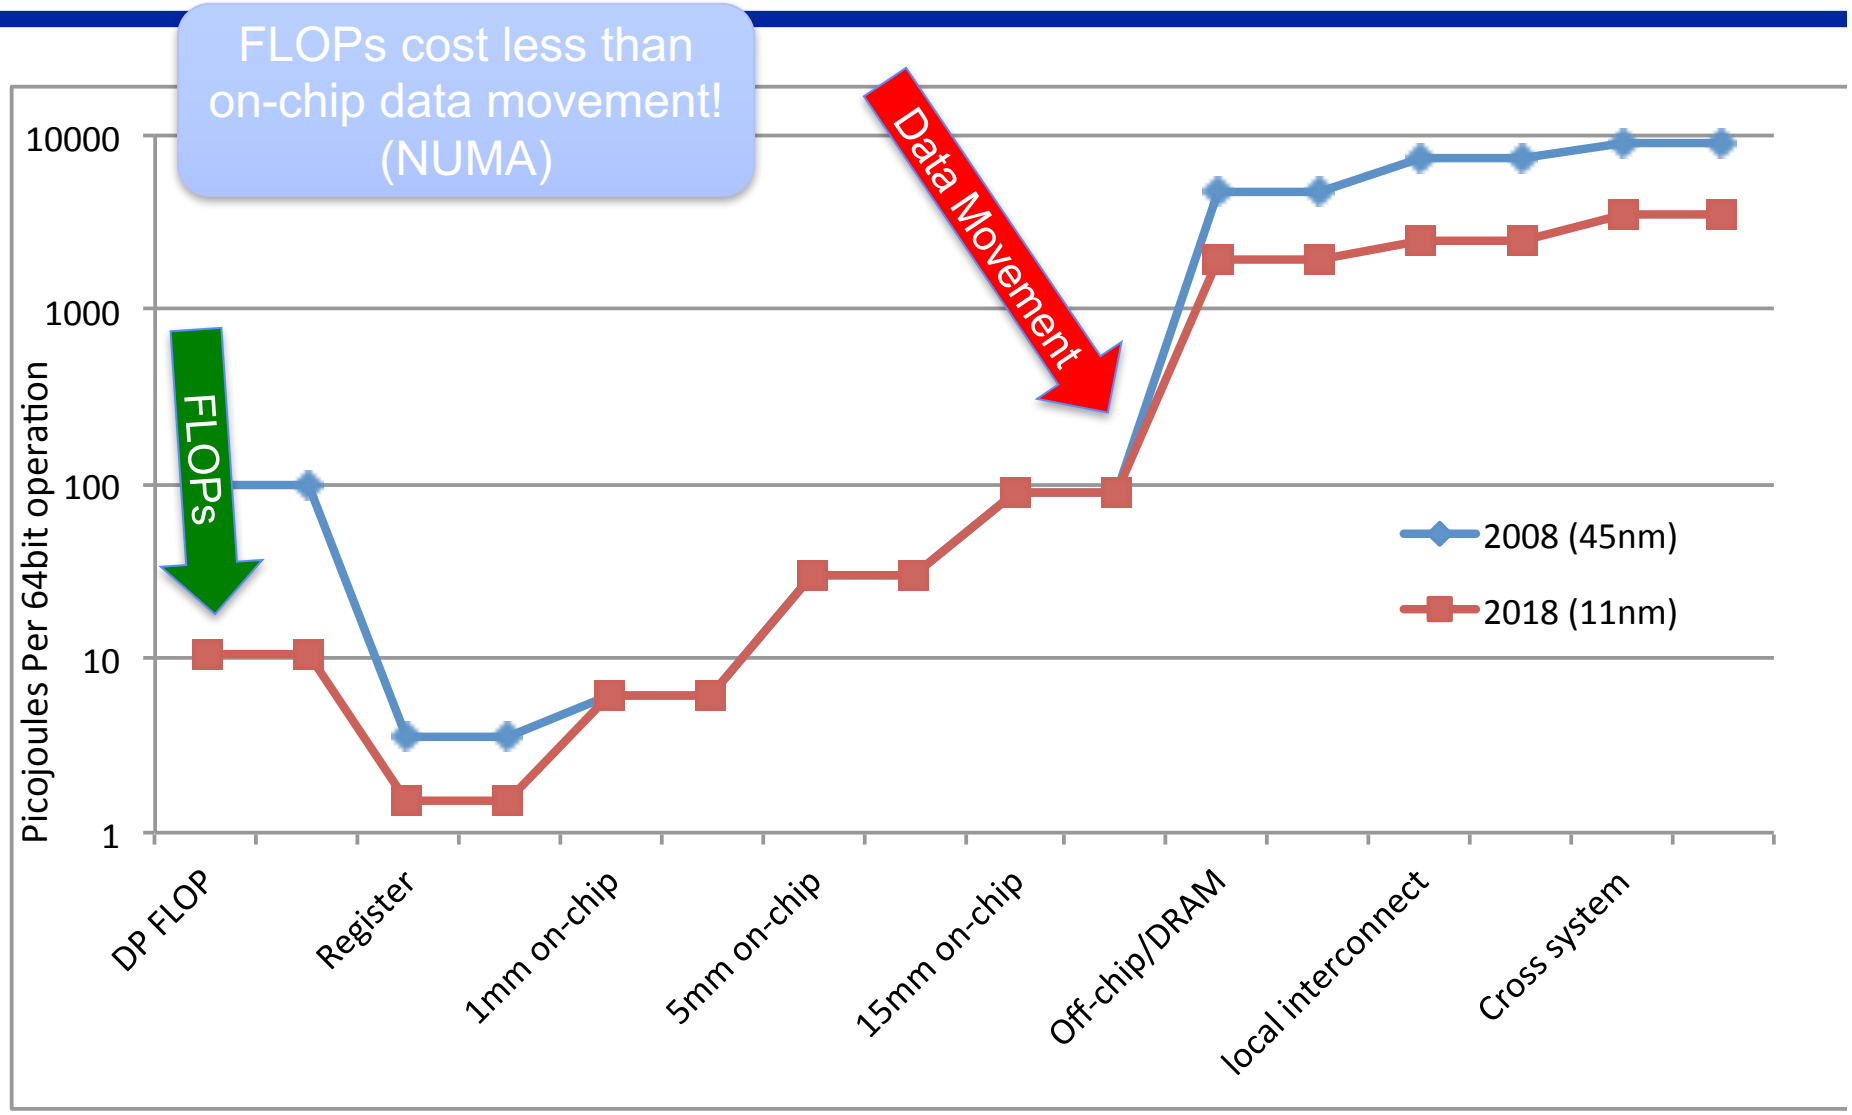

Organizing principles for Non-Von Digital Design

- Data Movement will remain a challenge even with exotic materials, but especially CMOS
- Copper is as good of a conductor as you can expect at room temperature
- With even lower power switches, challenges skews even more to data movement *(NEED Spatial Computing approach)*
- Push towards more parallelism (more tesselation of the memory structures)…. Strong Scaling

Strong Scaling extrapolates to ***limit case*** with **no separation of memory and compute** *(e.g. one PDE cell per processing element)*

Spatial Computing

PDE on a Block Structured Grid Extrapolated to Non-Von Neumann

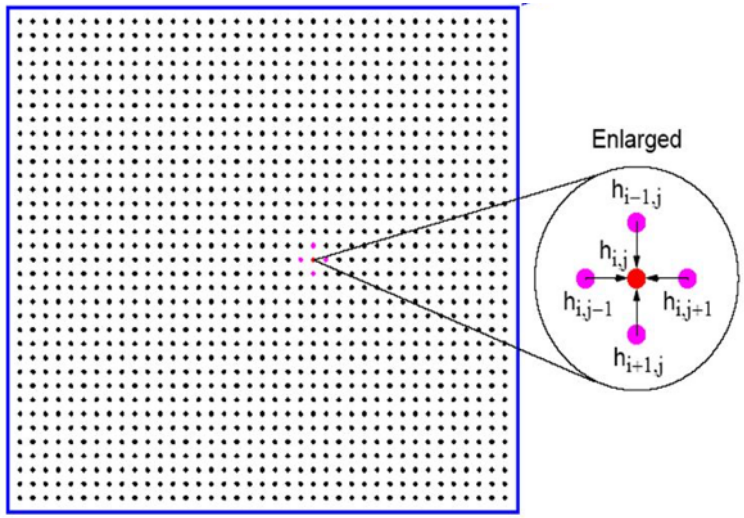

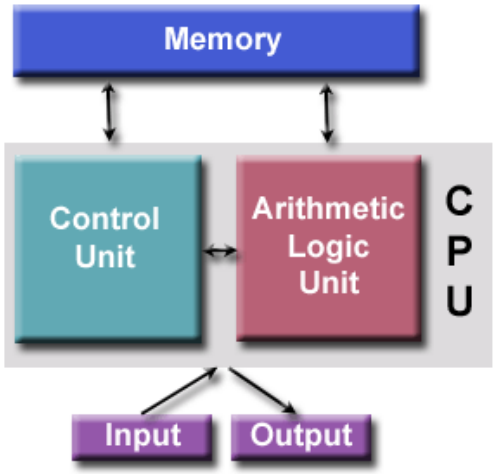

PDE on a Block Structured Grid Extrapolated to Non-Von Neumann

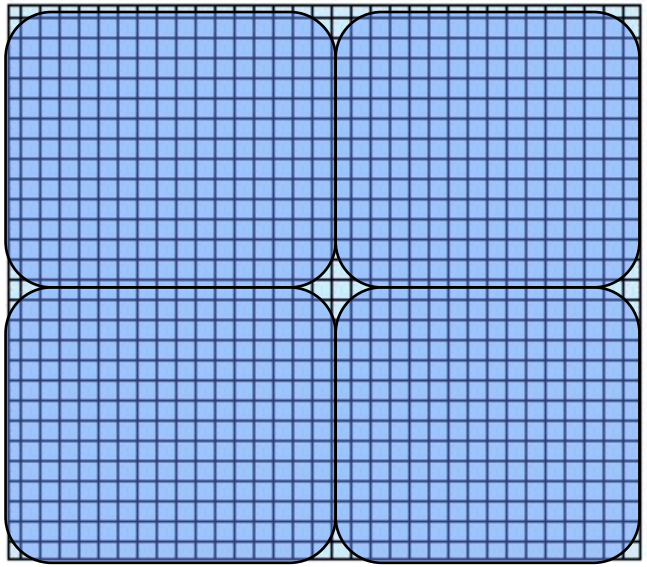

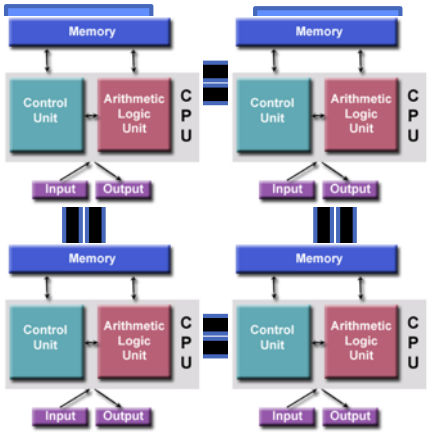

PDE on a Block Structured Grid Extrapolated to Non-Von Neumann

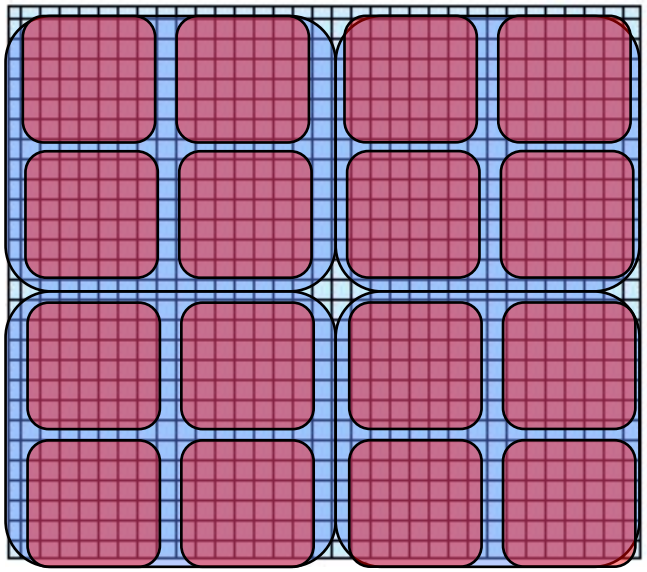

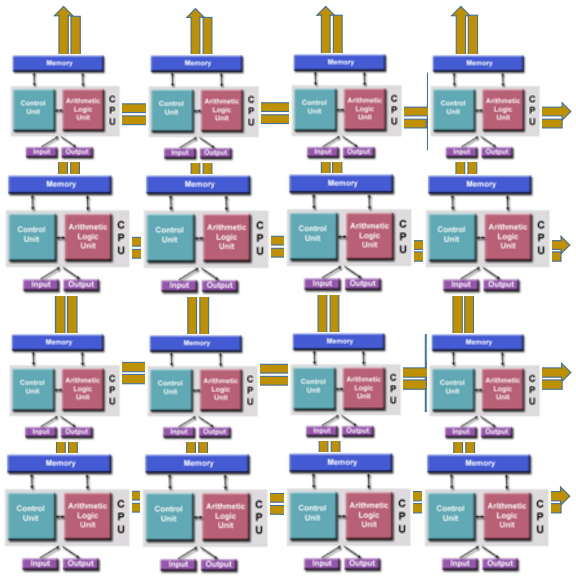

PDEcell / PICcell: Ultra-simple compute engine (50k gates) calculates finite-difference updates, and particle forces from neighbors. Microinstructions specify the PDE equation, stencil, and PIC operators.
Novel features: *variable length streaming integer arithmetic and novel PIC particle virtualization scheme.*

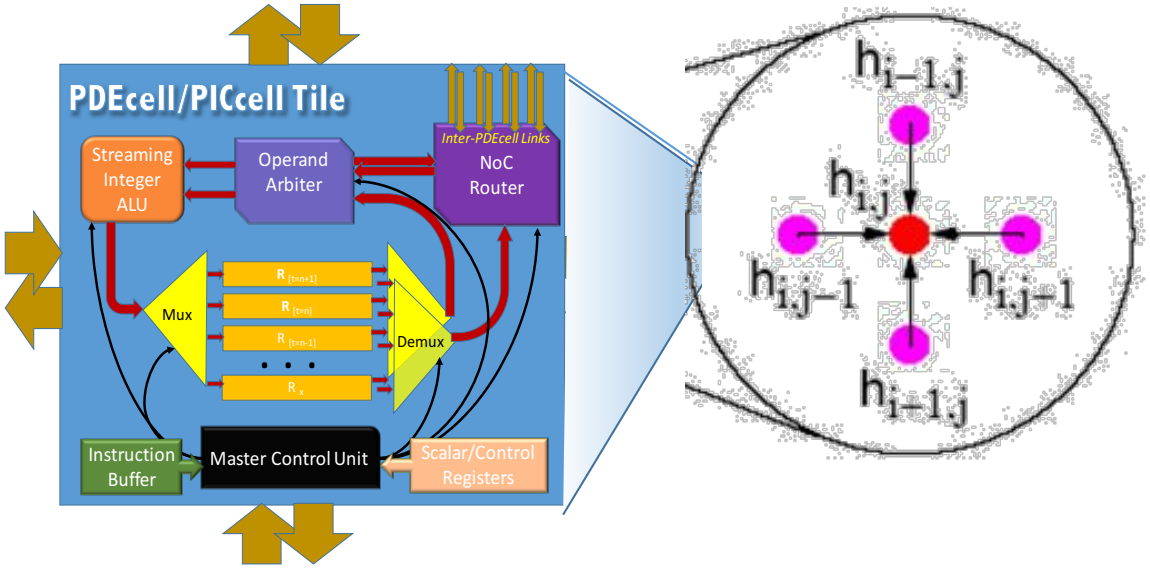

Concept: Solid State Virtual Fluid

Extreme (spatial) Specialization + New Devices + New programming models

PDEcell / PICcell: Ultra-simple compute engine (50k gates) calculates finite-difference updates, and particle forces from neighbors. Microinstructions specify the PDE equation, stencil, and PIC operators. ***Novel features:*** *variable length streaming integer arithmetic and novel PIC particle virtualization scheme.*

Computational Lattice: PDECells are tiles in a lattice/array on each 2D planar chip layer. Target 120x120 tiles per mm^2 @28nm lithography. Novel Features: each tile represents single cell of computational domain (pushes to limit of strong-scaling).

Monolithic 3D Integration: Integrate layers of compute elements using emerging monolithic 3D chip stacking.
Novel Features: *1000 layer stacking (20x more than current practice). Area efficient inter-layer connectivity and new energy efficient transistor logic (ncFET).*
1 Petaflop equivalent performance in 300mm^2 for < 200Watts.

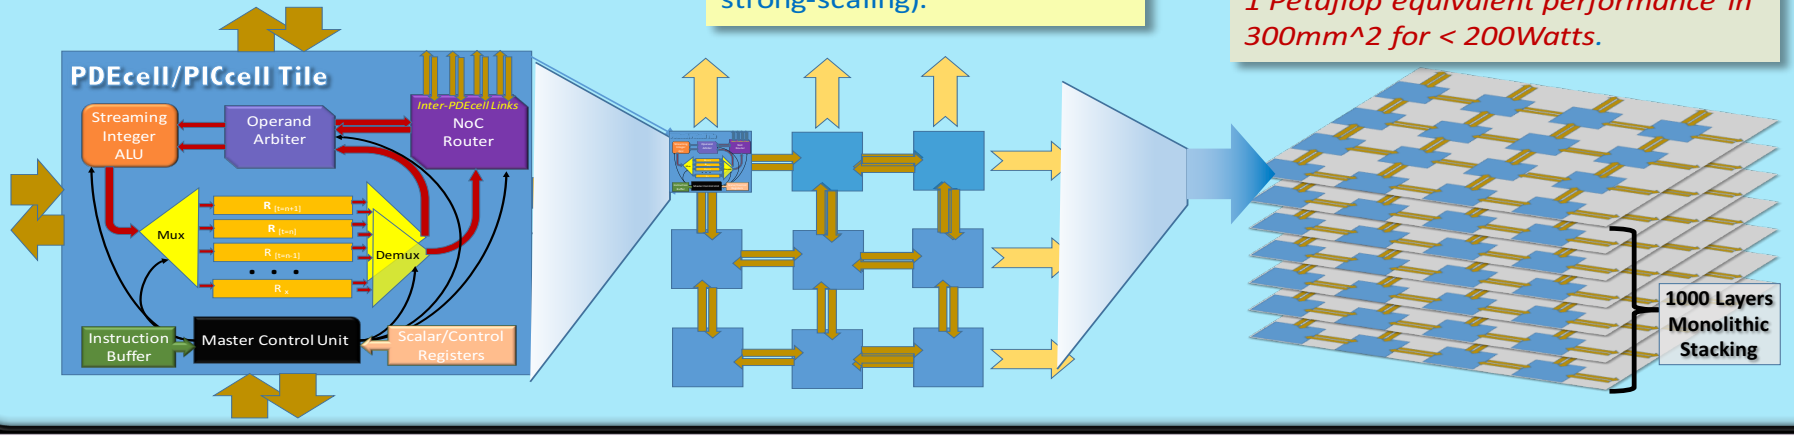

Scalar waves in 3D are solutions of the hyperbolic wave equation: $-\phi_{,tt} + \phi_{,xx} + \phi_{,yy} + \phi_{,zz} = 0$

Initial value problem: given data for ϕ and its first time derivative at initial time, the wave equation says how it evolves with time

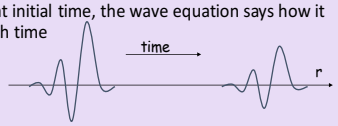

Discretized PDE Representation in DSL

$$\phi^{n+1}_{i,j,k} = 2\phi^{n}_{i,j,k} - \phi^{n-1}_{i,j,k}$$
$$+ \Delta t^2/\Delta x^2(\phi^{n}_{i+1,j,k} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i-1,j,k})$$
$$+ \Delta t^2/\Delta y^2(\phi^{n}_{i,j+1,k} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i,j-1,k})$$
$$+ \Delta t^2/\Delta z^2(\phi^{n}_{i,j,k+1} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i,j,k-1})$$

Compiles to MicroOps

```
R[n+1](0,0,0) = 0
R[n+1](0,0,0) += 2*R[n](0,0,0)
R[n+1](0,0,0) -= R[n-1](0,0,0)
R[n+1](0,0,0) += C * R[n+1](+1,0,0)
R[n+1](0,0,0) -= C * 2 * R[n](0,0,0)
R[n+1](0,0,0) += C * R[n](-1,0,0)
R[n+1](0,0,0) += C * R[n+1](0,+1,0)
    ....
```

Executes in Wavefronts

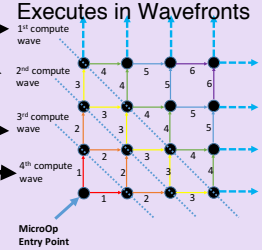

Initial Value Problem (continuous space)

Scalar waves in 3D are solutions of the hyperbolic wave equation: $-\phi_{,tt} + \phi_{,xx} + \phi_{,yy} + \phi_{,zz} = 0$

Initial value problem: given data for ϕ and its first time derivative at initial time, the wave equation says how it evolves with time

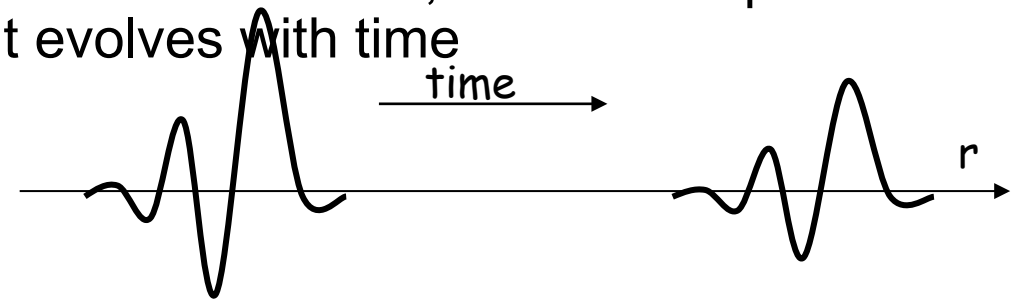

Discretized Representation

Numerical solve by discretising on a grid, using explicit *finite differencing* (centered, second order)

$$\phi^{n+1}_{i,j,k} = 2\phi^{n}_{i,j,k} - \phi^{n-1}_{i,j,k}$$
$$+ \Delta t^2/\Delta x^2(\phi^{n}_{i+1,j,k} - 2\,\phi^{n}_{i,j,k} + \phi^{n}_{i-1,j,k})$$
$$+ \Delta t^2/\Delta y^2(\phi^{n}_{i,j+1,k} - 2\,\phi^{n}_{i,j,k} + \phi^{n}_{i,j-1,k})$$
$$+ \Delta t^2/\Delta z^2(\phi^{n}_{i,j,k+1} - 2\,\phi^{n}_{i,j,k} + \phi^{n}_{i,j,k-1})$$

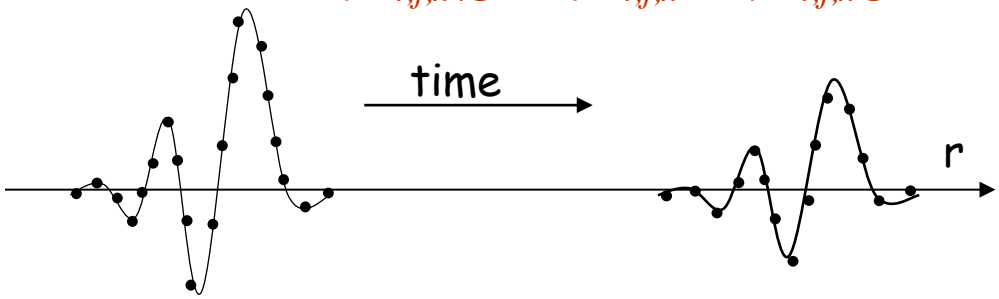

Decomposing Into PDE Weights

$$\phi^{n+1}_{i,j,k} = 2\phi^{n}_{i,j,k} - \phi^{n-1}_{i,j,k}$$
$$+ \Delta t^2/\Delta x^2(\phi^{n}_{i+1,j,k} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i-1,j,k})$$
$$+ \Delta t^2/\Delta y^2(\phi^{n}_{i,j+1,k} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i,j-1,k})$$
$$+ \Delta t^2/\Delta z^2(\phi^{n}_{i,j,k+1} - 2\phi^{n}_{i,j,k} + \phi^{n}_{i,j,k-1})$$

$R_{[t=n+1]}(0,0,0) = 0$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= 2*R_{[t=n]}(0,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{-}= R_{[t=n-1]}(0,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= C * R_{[t=n+1]}(+1,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{-}= C * 2 * R_{[t=n]}(0,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= C * R_{[t=n]}(-1,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= C * R_{[t=n+1]}(0,+1,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{-}= C * 2 * R_{[t=n]}(0,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= C * R_{[t=n]}(0,-1,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= C * R_{[t=n+1]}(0,0,+1)$
$R_{[t=n+1]}(0,0,0) \mathrel{-}= C * 2 * R_{[t=n]}(0,0,0)$
$R_{[t=n+1]}(0,0,0) \mathrel{+}= C * R_{[t=n]}(0,0,-1)$
Rotate Registers

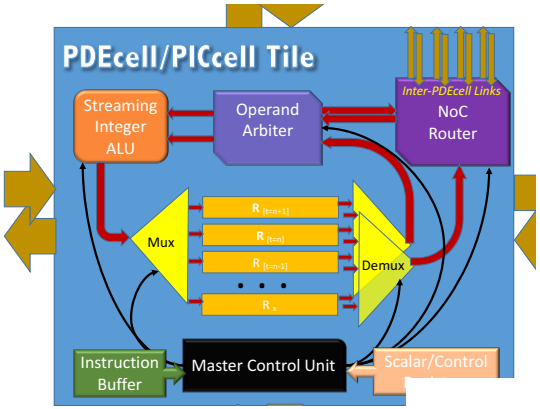

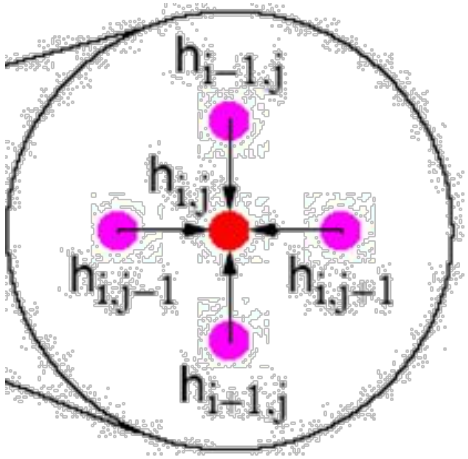

PDE Domain Specific Representation

PDE in Domain Specific Representation

$$\phi^{n+1}_{i,j,k} = 2\phi^{n}_{i,j,k} - \phi^{n-1}_{i,j,k}$$
$$+ \Delta t^2/\Delta x^2(\phi^{n}_{i+1,j,k} - 2\,\phi^{n}_{i,j,k} + \phi^{n}_{i-1,j,k})$$
$$+ \Delta t^2/\Delta y^2(\phi^{n}_{i,j+1,k} - 2\,\phi^{n}_{i,j,k} + \phi^{n}_{i,j-1,k})$$
$$+ \Delta t^2/\Delta z^2(\phi^{n}_{i,j,k+1} - 2\,\phi^{n}_{i,j,k} + \phi^{n}_{i,j,k-1})$$

PDE Weights

```
R[n+1](0,0,0) = 0
R[n+1](0,0,0) += 2*R[n](0,0,0)
R[n+1](0,0,0) -=  R[n-1](0,0,0)
R[n+1](0,0,0) += C * R[n+1](+1,0,0)
R[n+1](0,0,0) -= C * 2 * R[n](0,0,0)
R[n+1](0,0,0) += C * R[n](-1,0,0)
R[n+1](0,0,0) += C * R[n+1](0,+1,0)
    . . . .
```

Fed Sequentially to

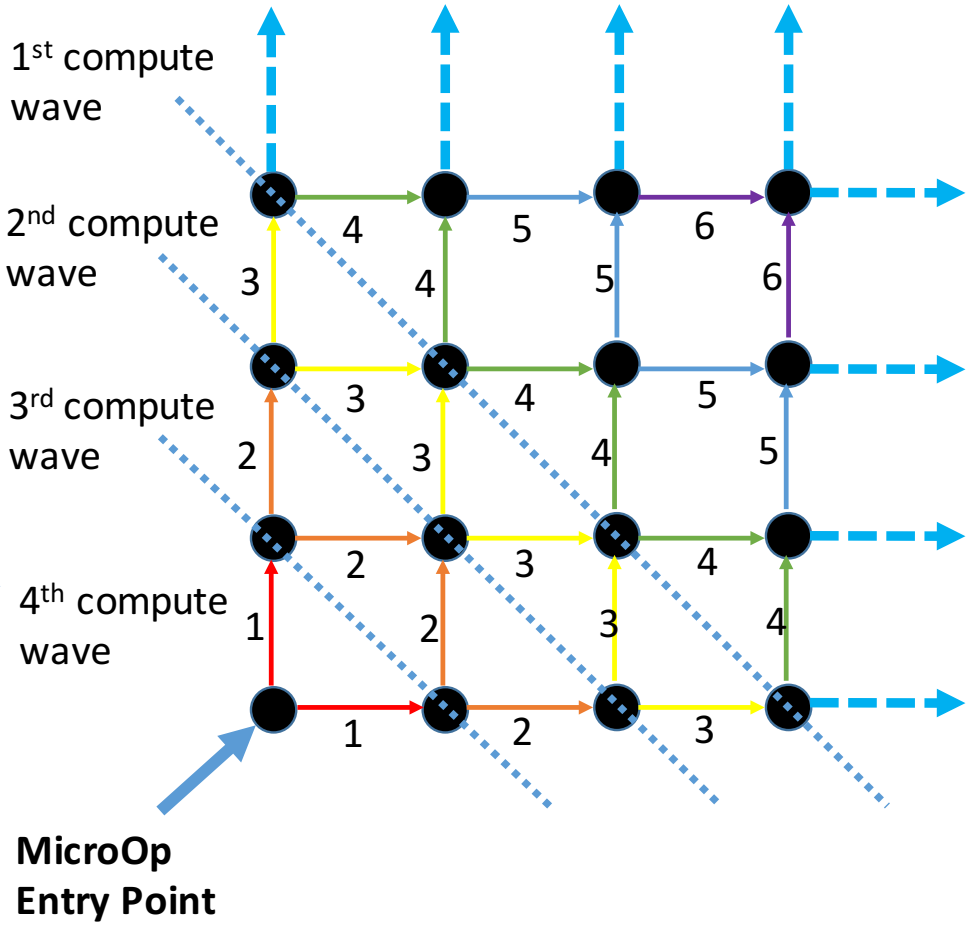

Challenges of Limit-Case Non-Von Neumann Digital

Von Neumann ~= Instruction Processor

```
int main()
{
   int n = 0;
   while(n < 100)
   {
      n = n + 5;
      print("n = %d\n", n);
      pause(200);
      if(n == 50) break;
   }
   print("All done!");
}
```

What the heck is this?

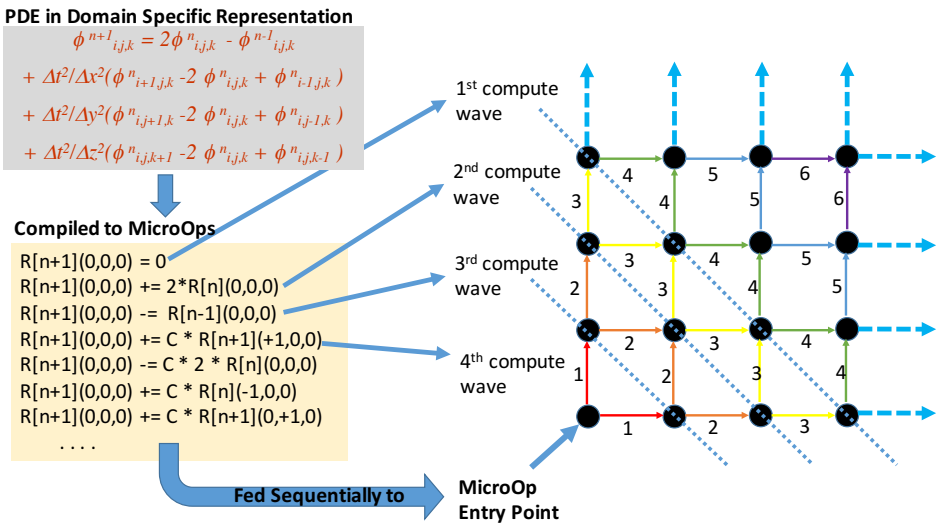

Modern languages (including many classes of DSLs Were designed with instruction processors in mind