
Basic Research Needs for Microelectronics Workshop October 23-25, 2018: *Preliminary findings*

Workshop Chair: Cherry Murray (Harvard Univ.)
Associate Chairs: Supratik Guha (ANL)
Dan Reed (Univ. of Utah)

Technology Liaison: Gil Herrera (SNL)
SC Technical Leads: Andy Schwartz (BES)
Robinson Pino (ASCR)
Eric Colby (HEP)

on behalf of DOE Office of Science
Basic Energy Sciences, Advanced Scientific Computing, High Energy Physics



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Basic Research Needs – Use Inspired Basic Research

- **Transformative, not incremental research directions**
- **Fundamental science challenges to move the technology forward**
- **New techniques and methods**
- **10-30 years out**





BRN for Microelectronics Workshop – Motivation

- Semiconductor-based microelectronics are critical to the U.S. economy, scientific advancement, and national security
 - Semiconductor products are currently the third largest class of U.S. exports (behind aircraft and automobiles)
 - U.S. companies account for more than 50% of the world market by revenue
 - Semiconductor industry directly employs ~250,000 people; ~1 million associated jobs
- The decades long success of Moore’s Law was driven by innovation
 - Materials and chemical sciences
 - Computer science
 - Electrical engineering
 - Fabrication technologies
- Additional innovation needed to keep up with dramatic market growth




Motivation: CMOS scaling slowdown




http://www.swtest.org/swtw_library/2015proc/PDF/SWTW2015_Keynote_McCann_GlobalFoundries.pdf

Uncertainty at 7 nm node
 complexity and cost
 Physics gets in way



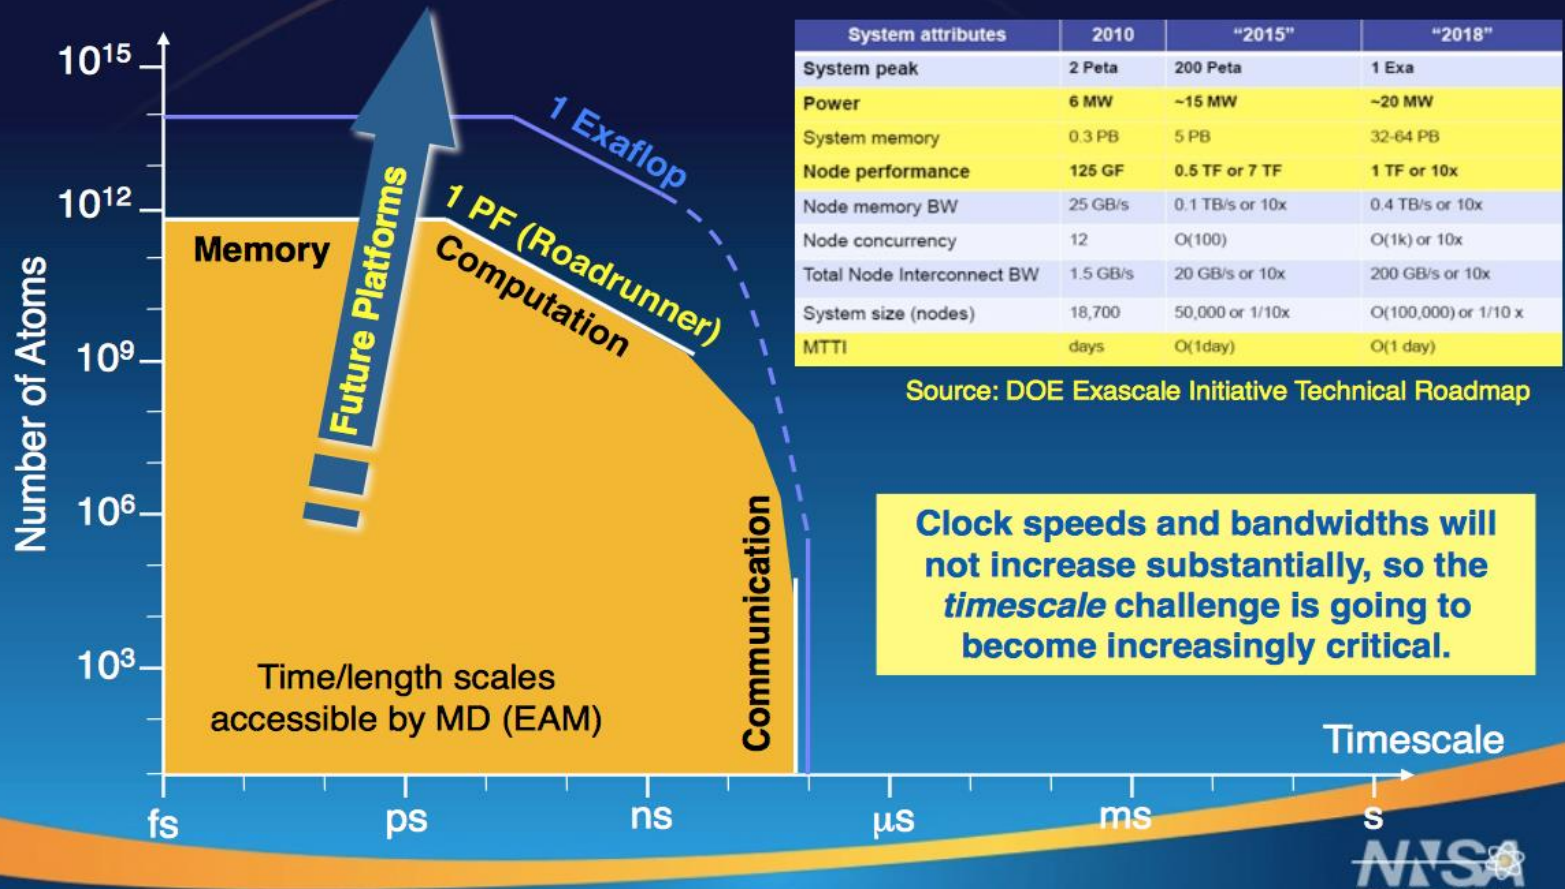
bulk



<https://www.quora.com/What-is-a-FinFET-transistor>

Impact to computational materials science--example

Current trends will increase the *length* scales accessible by large-scale molecular dynamics simulations



System attributes	2010	"2015"	"2018"
System peak	2 Peta	200 Peta	1 Exa
Power	6 MW	~15 MW	~20 MW
System memory	0.3 PB	5 PB	32-64 PB
Node performance	125 GF	0.5 TF or 7 TF	1 TF or 10x
Node memory BW	25 GB/s	0.1 TB/s or 10x	0.4 TB/s or 10x
Node concurrency	12	O(100)	O(1k) or 10x
Total Node Interconnect BW	1.5 GB/s	20 GB/s or 10x	200 GB/s or 10x
System size (nodes)	18,700	50,000 or 1/10x	O(100,000) or 1/10 x
MTTI	days	O(1day)	O(1 day)


Source: DOE Exascale Initiative Technical Roadmap

Rise of data intensive & edge computing: two examples of Office of Science data management needs

Future needs at BES Light Sources (ALS, APS, LCLS/LCLS-II, NSLS-II, SSRL)

Compute needs 10-15X from now



A data storage and movement problem



From: BES User Facilities Data Management and Analysis Resource Needs BES/ASCR Data Call

Future CMS needs at Large Hadron Collider (LHC)

disk and tape storage, 50x w.r.t. now
CPU needs 5M cores, 20x w.r.t. now



(From talk by L Bauerdick, Fermilabs, conveyed by S. Habib, ANL)

Rise of data intensive & edge computing

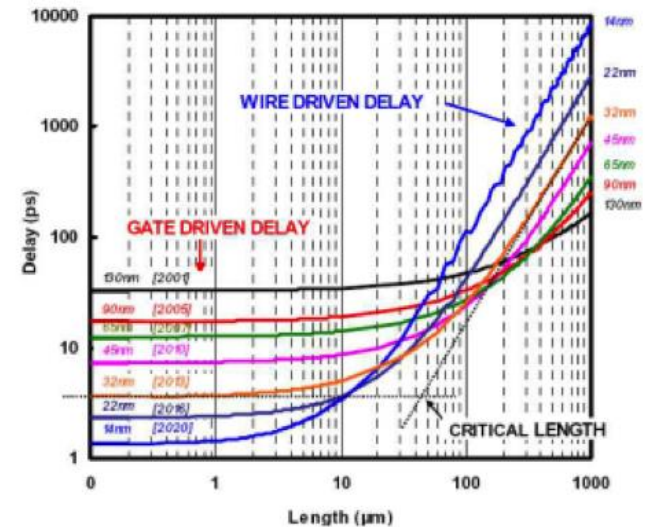
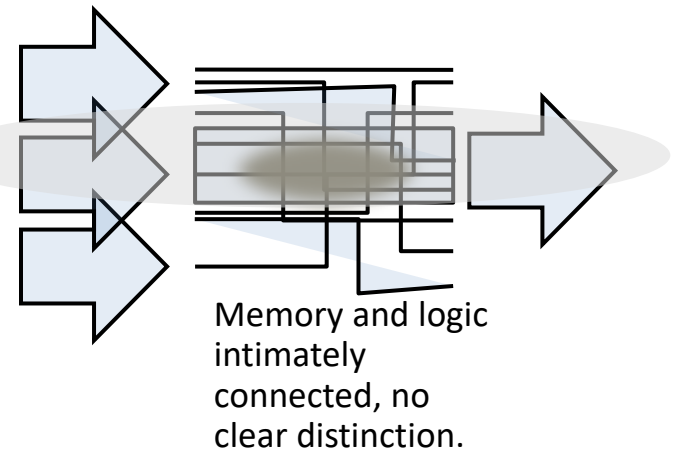



Need for new computing paradigms

Memory bottlenecks

Data transport

Low power computing

One approach



Source: M. Sellier, ISQED 2008

Future Computing Technologies are Important to DOE

as well as many other Federal Agencies

- High-performance computing & simulation underpin DOE missions in energy, environment, and national security
 - Historical role of computing in DOE
 - DOE/vendor synergies in deploying computing technologies
- Future computing technologies (e.g., quantum, neuromorphic, probabilistic, etc.) hold promise for next-generation DOE mission applications
 - DOE research and facilities (e.g. HEP experiments, ASCR HPC, BES light sources) will depend on advanced computing and sensing technologies
 - Likely will augment, not replace, conventional supercomputing
 - Could open new avenues for use of computing in science (data analytics, machine learning, artificial intelligence, ...)
- New directions for applied mathematics and computer science are likely to emerge that could enable new science across DOE-SC

Call To Action

- Significant challenges as CMOS extends below 5nm
- The end to Moore's Law will impact U.S. industry and competitiveness
- The importance of this issue and its technical complication will require *innovative approaches* to keep the U.S. in a leadership position
- Solving a problem of this scale will require “*whole of government*” approach and a robust *public/private partnership* to apply the best research from *industry, academia and government research facilities* to allow the U.S. to successfully make this technology transition
- **DOE, and particularly the Office of Science, will play a significant role in this effort**
- **DOE-SC was charged with organizing a *Basic Research Needs Workshop* to define the highest priority research directions**

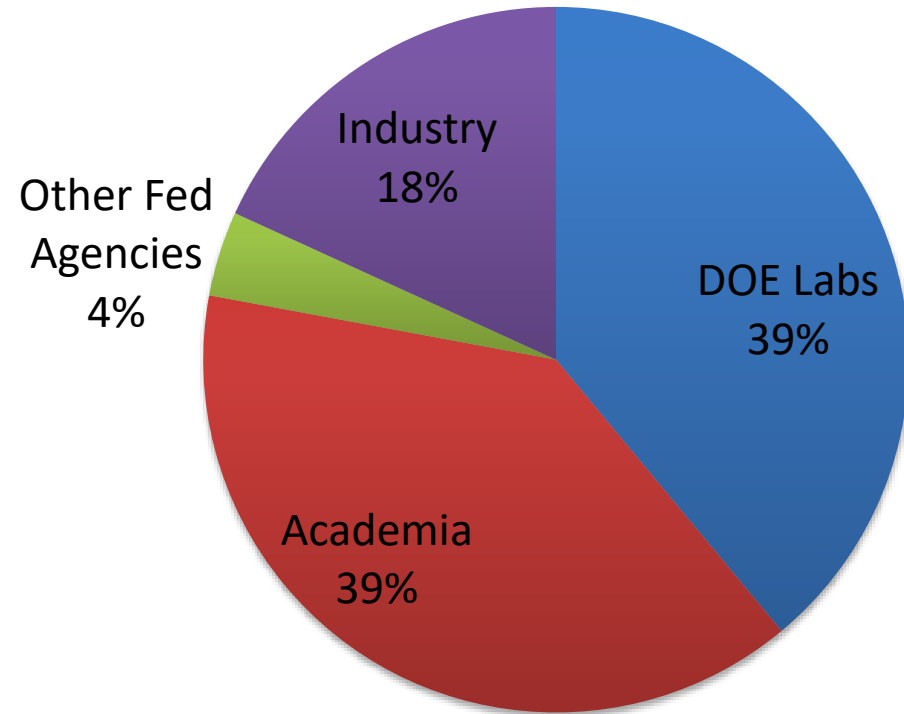


Basic Research Needs for Microelectronics – Charge


- **A thorough assessment of the scientific issues** associated with advanced microelectronics technologies for applications relevant to the DOE mission.
- **Identify critical scientific challenges, fundamental research opportunities,** and priority research directions that require further study as a foundation for advances in microelectronics over the next decade and beyond.
- Particular emphasis on energy-relevant applications, and areas that are aligned with the missions and needs of ASCR, BES, HEP including data management and processing, power electronics, and high performance computing.
- Examine extension of CMOS and beyond CMOS technologies, **beyond exascale technologies. however Quantum Information Science is outside the scope of this workshop.**
- focus on a co-design innovation ecosystem in which materials, chemistries, devices, systems, architectures, and algorithms are researched and developed in a closely integrated fashion.

Basic Research Needs for Microelectronics Workshop participation

77 panelists, ~70 observers



By affiliation



By expertise

Systems: circuits, micro-architecture, architecture, algorithms, software

Hardware: devices, materials, physics, chemistry

Panels

Panelists were invited for their expertise – and are assigned to a particular panel that will determine priority research directions in the breakout sessions

1) Big data collection, analytics, processing for SC facilities

Leads: Kirsten Kleese van Dam (BNL) and Sayeef Salahuddin (UC Berkeley)

2) Co-design for high performance computing beyond exascale

Leads: James Ang (PNNL) and Thomas Conte (Georgia Tech)

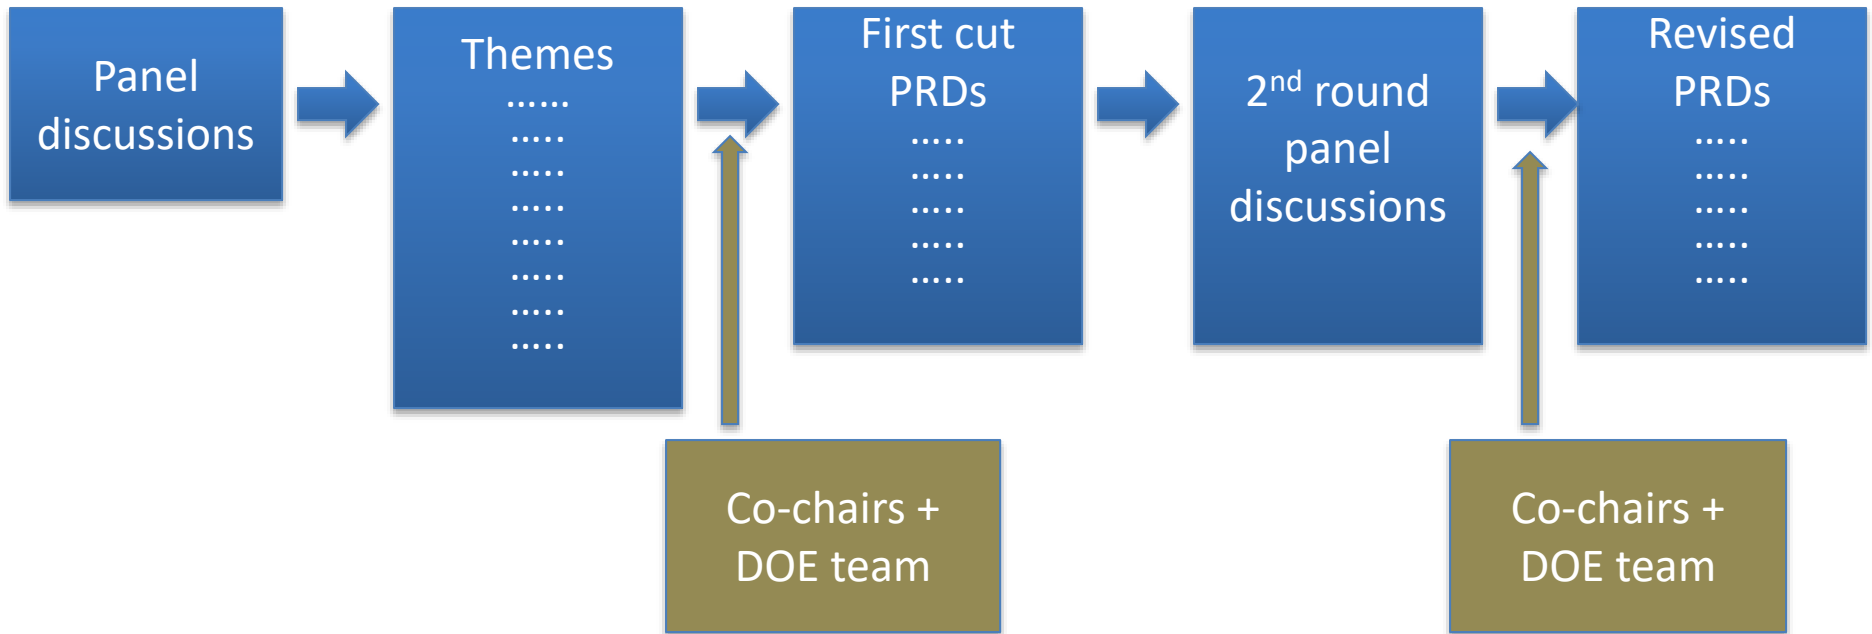
3) Power control, conversion and detection

Leads: Debdeep Jena (Cornell U) and Robert Kaplar (SNL)


4) Crosscutting themes – *may roam and join other panels*

Leads: Harry Atwater (Caltech) and Rick Stevens (ANL)

Oct 22-25



Summary Brochure Published on 7 December 2018



Five Priority Research Directions (PRDs) Identified

Priority Research Directions

- Flip the current paradigm: Define innovative material, device, and architecture requirements driven by applications, algorithms, and software**

Key Questions: How can we optimize and integrate across physical, logical, and communication and control hierarchies? How will system-level optimization enable directed materials/device discovery and innovation?

Materials properties, microelectronic devices, architectures, and algorithms must be understood and designed from the atomistic to the systems level to address the critical technical challenges facing DOE in its missions of science, energy, and national security. The outcome of an "end-to-end co-design framework" will reshape high performance computing, data analytics, the electricity grid, and other computing intensive and high power applications.
- Revolutionize memory and data storage**

Key Questions: How do we link physics, materials, architectures, and algorithms to overcome current physical limits on access and retention times for memory and storage? What innovations will minimize data movement and reduce energy consumption by orders of magnitude?

Memory technologies are critically important in all aspects of data acquisition, analysis, and storage, and have the potential to perform efficient computations within, or proximally close to, the memory element. We face fundamental tradeoffs between fast memory access, capacity, and data retention time, as well as key challenges in energy usage and heat dissipation. Meeting these challenges will require coordinated breakthroughs in materials, device design, computer architecture, and algorithms.
- Reimagine information flow unconstrained by interconnects**

Key Questions: How can we minimize data movement while maximizing information transfer? What novel electronic/optical states of matter can be discovered and manipulated to design non-traditional interconnects at the atomic, micro, and macro scales?

A co-design approach to developing novel interconnect architectures will enable seamless integration of large-scale, real-time computation with communications and sensing to dramatically improve data transfer rates, connectivity, and reconfigurability.
- Redefine computing by leveraging unexploited physical phenomena**

Key Questions: What unexplored materials, phenomena, or alternative computing models could perform computation for more efficiently than today's technology? How will these new systems be modeled and programmed?


The capabilities of the prevailing model of computation, the von Neumann model, are increasingly constrained by the energy inefficiency of established hardware and architecture. Understanding and using new computing models based on unexploited phenomena require a co-design approach spanning architectures and algorithms to physics, materials science, and new devices.
- Reinvent the electricity grid through new materials, devices, and architectures**

Key Question: Using a co-design approach, how do we create novel devices based on new materials to enable revolutionary breakthroughs in the performance, reliability, and security of power conversion systems?

Revolutionary advances in power electronics for the electricity grid will require the design, synthesis, understanding, processing, and integration of advanced semiconductors and magnetic and dielectric materials. Novel device, circuit, and thermal transport concepts will be developed to exploit the unique physical properties of these materials. Such energy-efficient power conversion systems are necessary to replace the century-old electricity grid with one appropriate for the 21st century. They could also be applicable to electric transportation and use in extreme environments such as accelerators and power generation facilities.

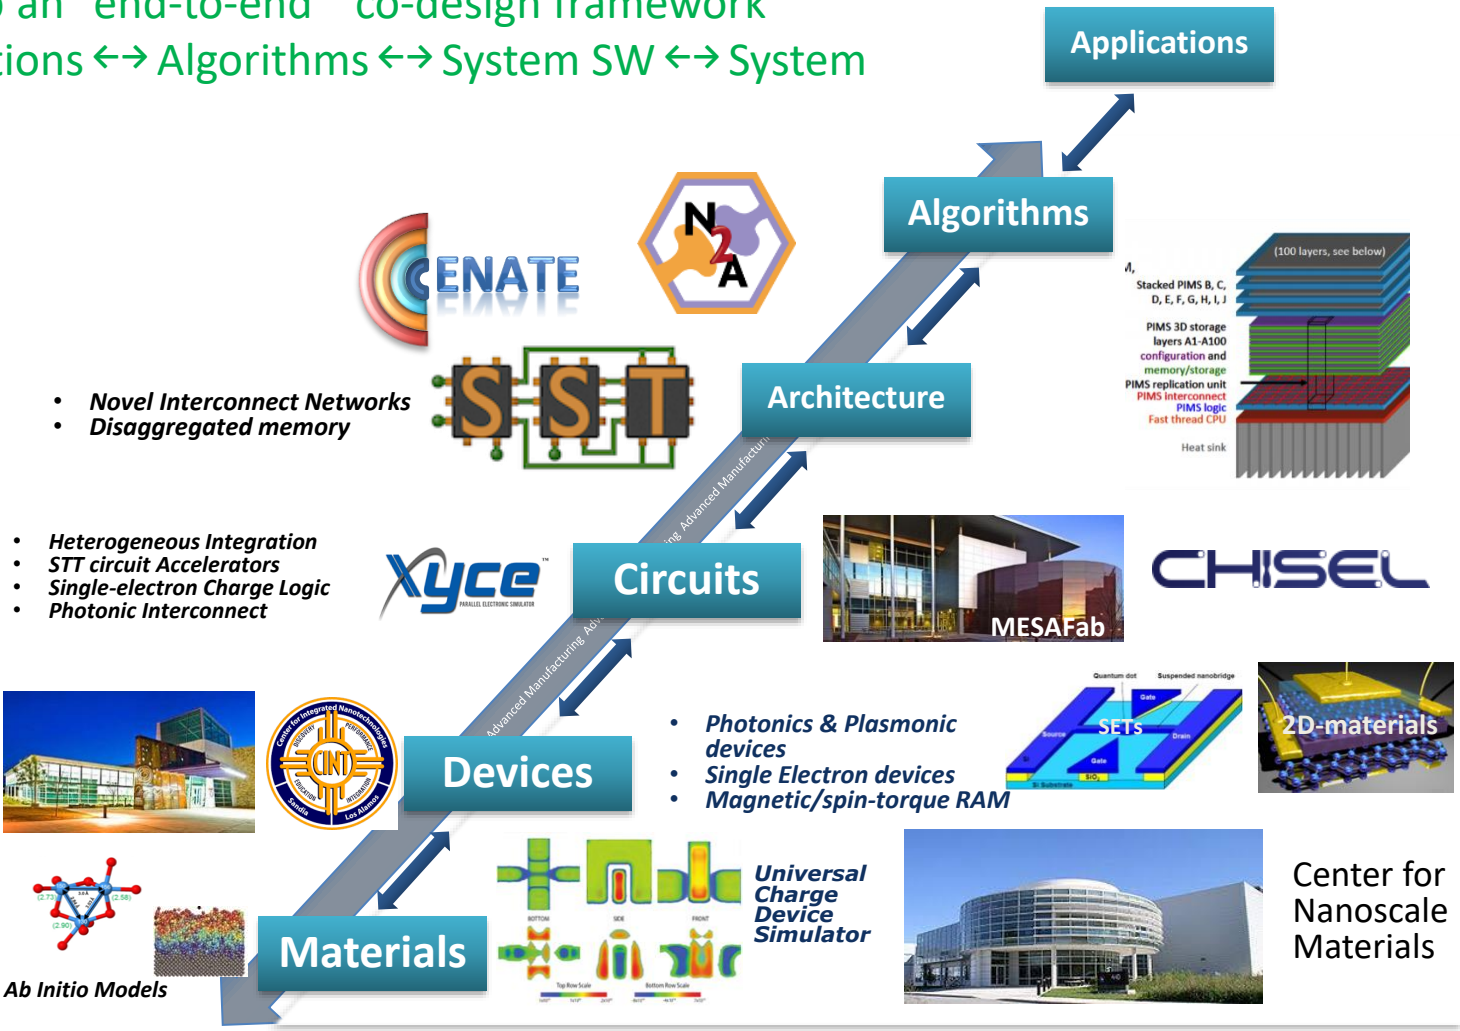
Full report target date: Apr 2019

Principles of co-design underpin all five priority research directions (PRDs)



PRD 1: Flip the current paradigm: Define innovative materials, device, and architecture requirements driven by applications, algorithms, and software

- Develop an “end-to-end” co-design framework
- Applications ↔ Algorithms ↔ System SW ↔ System HW




- Novel Interconnect Networks
- Disaggregated memory
- Heterogeneous Integration
- STT circuit Accelerators
- Single-electron Charge Logic
- Photonic Interconnect

PRD 2: Revolutionize memory and data storage

Storage & Memory

SRAM



>\$300/GB
10 nano-sec
~ 1Gb/sq. in. equiv

DRAM



\$7/GB
50 nano-sec
~100 Gb/sq. in.

Flash



\$1/GB
micro-sec
>1Tb/sq. in.

Hard Disk



\$0.03/GB
milli-sec
>1Tb/sq. in.

Tape



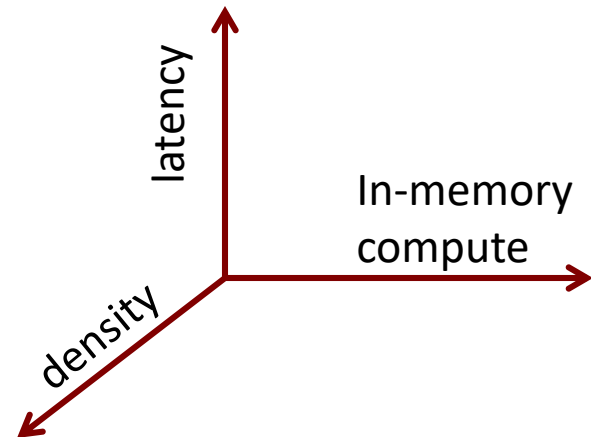

\$0.01/GB
Seconds
> 1Tb/sq. in.

DNA



>\$3.5K/MB
Hours
~80 Exa-b/in³

- Density: 1-100s of atom memory
- Latency
- Bandwidth
- In-memory computing

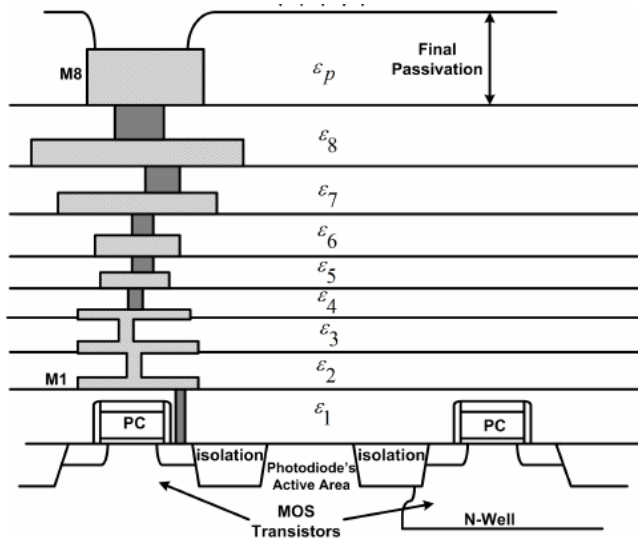


PRD 3: Reimagine information flow unconstrained by interconnects


Data movement is growing exponentially

pJ/bit not ramping down significantly

Worthy Goal: \gg Tbyte/sec-mm channel capacity for <100 fJ/bit




https://www.researchgate.net/figure/General-structure-of-130-nm-technology-with-Back-end-of-line-metallization-and-dielectric_fig11_224918168



<https://images.anandtech.com/doci/8367/14nmInterconnect.jpg>

Microelectronics moving to three dimensions



PRD 4: Redefine computing by leveraging unexploited physical phenomena



Finding and understanding physical phenomena that can express computation

New ways of reasoning about computation


Leveraging physical processes to compute (“analogous computing”)

NvN Optimizers, both continuous and integer

Artificial Neural Networks



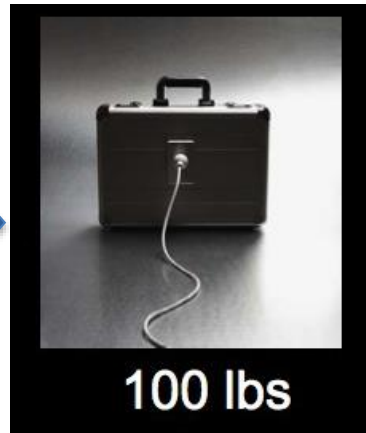
LEGEND: No Disruption



Substation in a Suitcase



8000 lbs, 60 Hz Distribution Transformer



*Silicon Carbide IGBT;
15 kV, 100 A;
50 kHz from Cree Inc.*

