

# Specialized Architectures: Enabling Performance Scaling in a Post Moore Era

David Donofrio

Berkeley National Labs

March 17, 2017

Supercomputing Frontiers



U.S. DEPARTMENT OF  
**ENERGY**

Office of  
Science

# Post Moore Technology Curve

Great *opportunities* exist for innovation through the end of Moore's Law

## End of Moore's Law

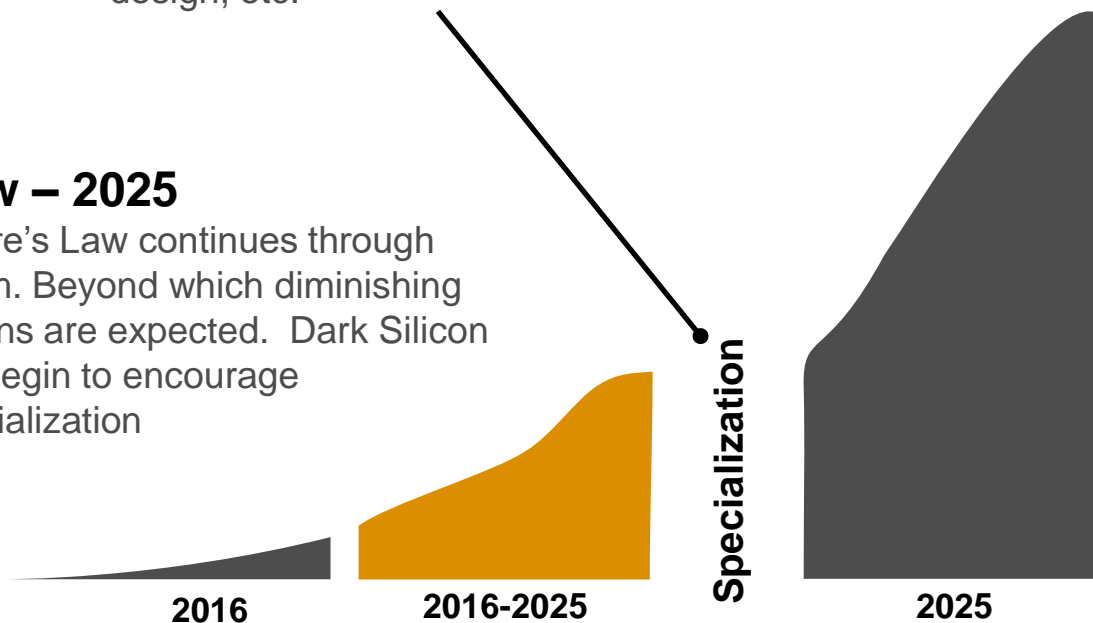
End of Moore's Law requires a different set of optimizations to continue performance scaling. Opportunities for additional specialization, reconfigurable computing, hardware / software co-design, etc.

## Throughout...

Continued increases in parallelism and heterogeneity will require advanced runtimes, programming environments and compiler optimizations in order to take full advantage of these new architectures

## Now – 2025

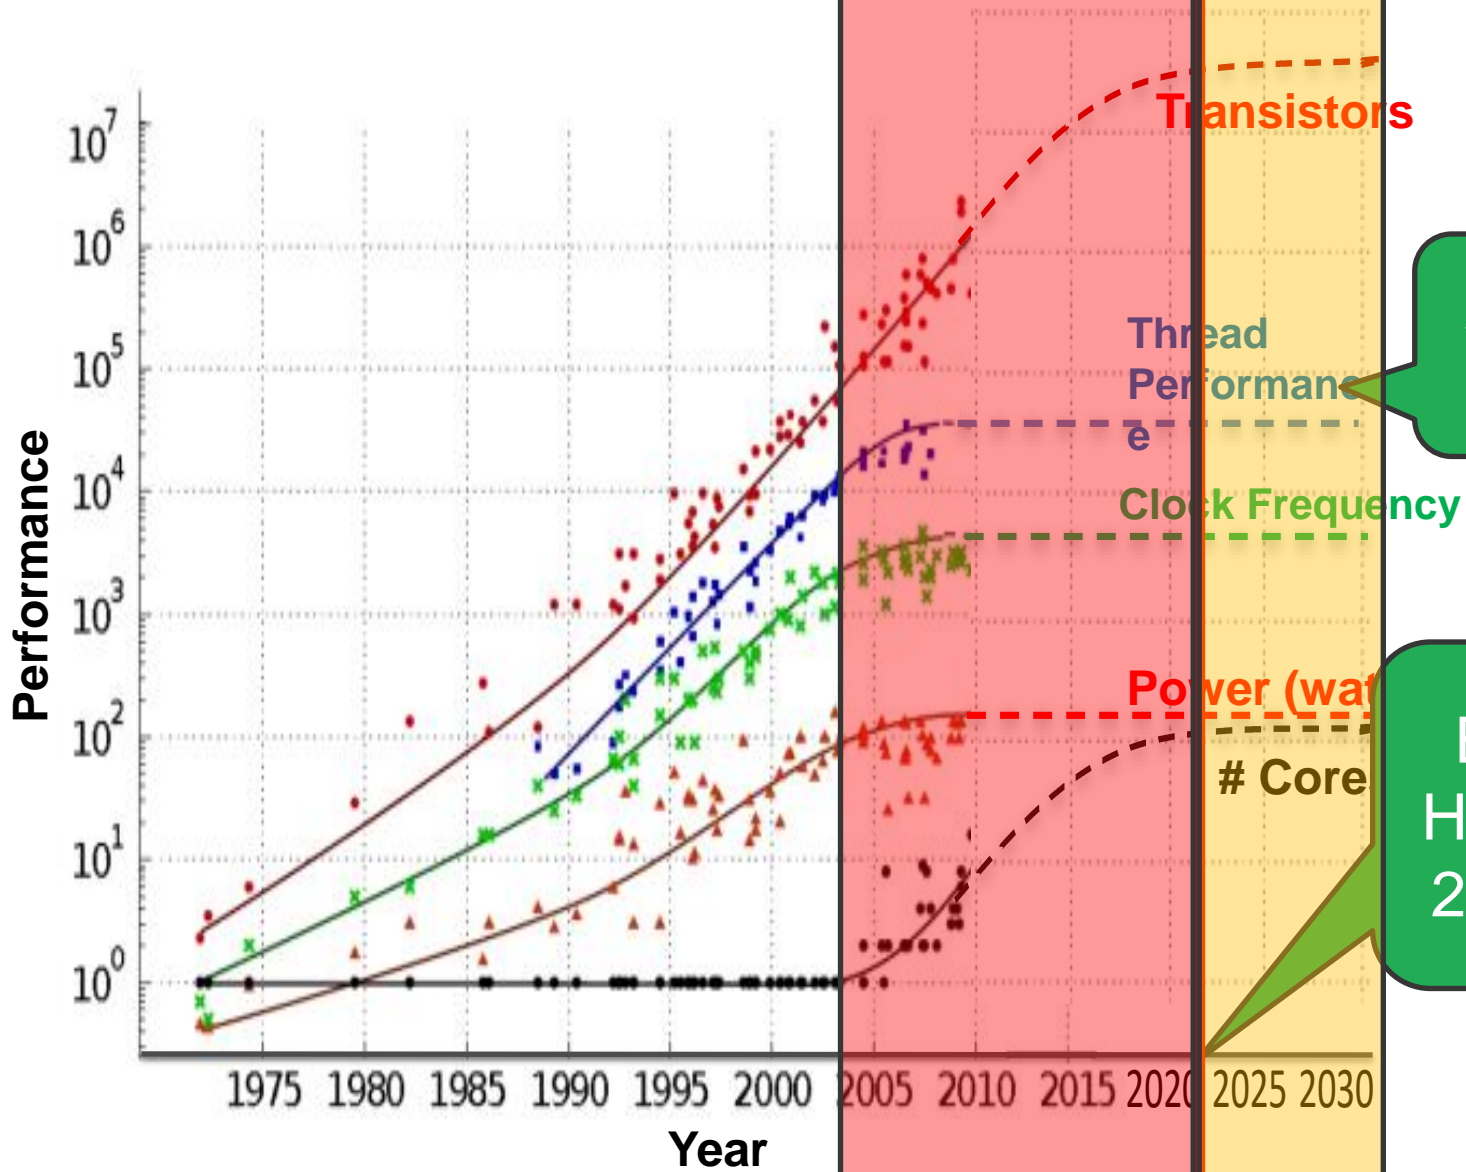
Moore's Law continues through ~5nm. Beyond which diminishing returns are expected. Dark Silicon will begin to encourage specialization



## Post Moore Scaling

New materials possibly introduced to allow continued process and performance scaling.

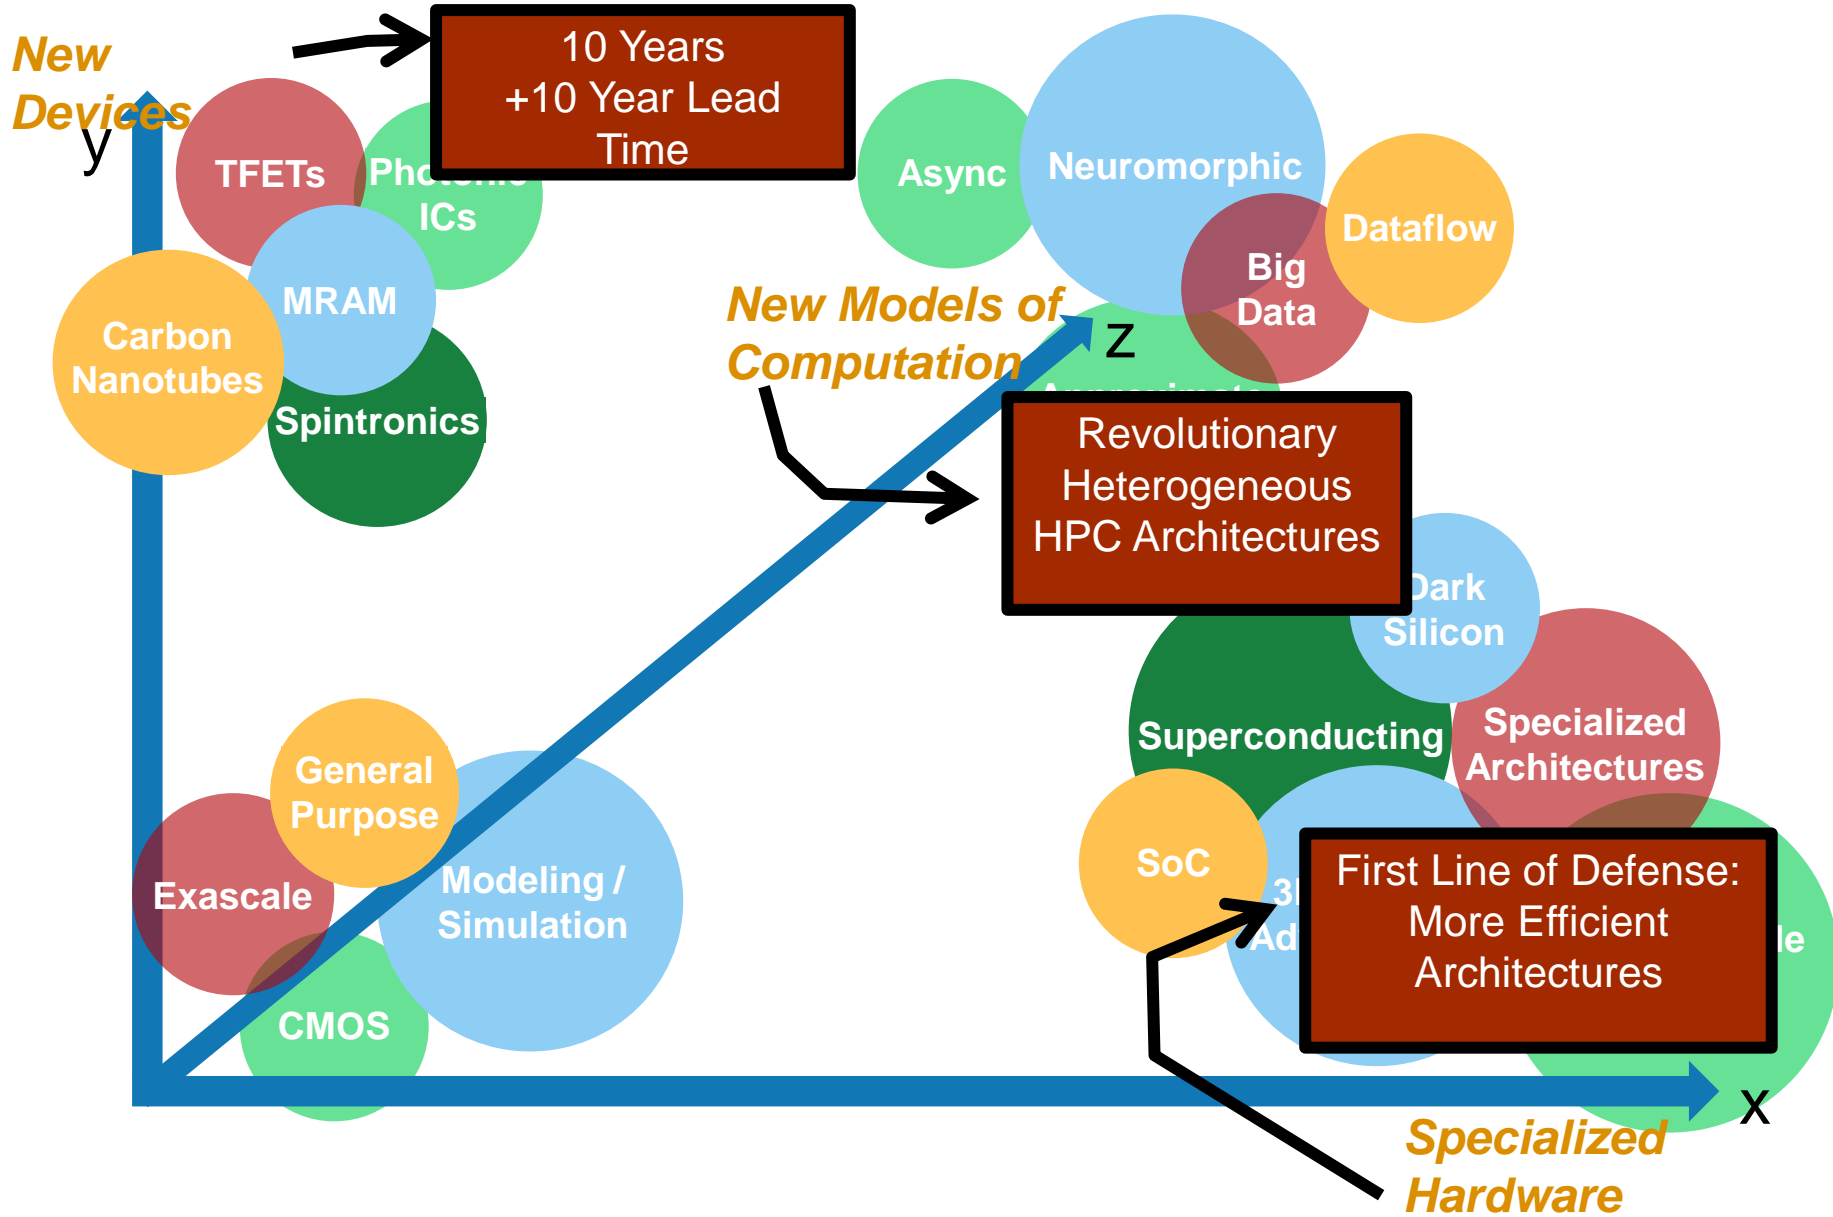
# Technology Scaling Trends



And Then  
What?

Exascale  
Happens in  
2021-2023

# Paths Forward in Post-Moore



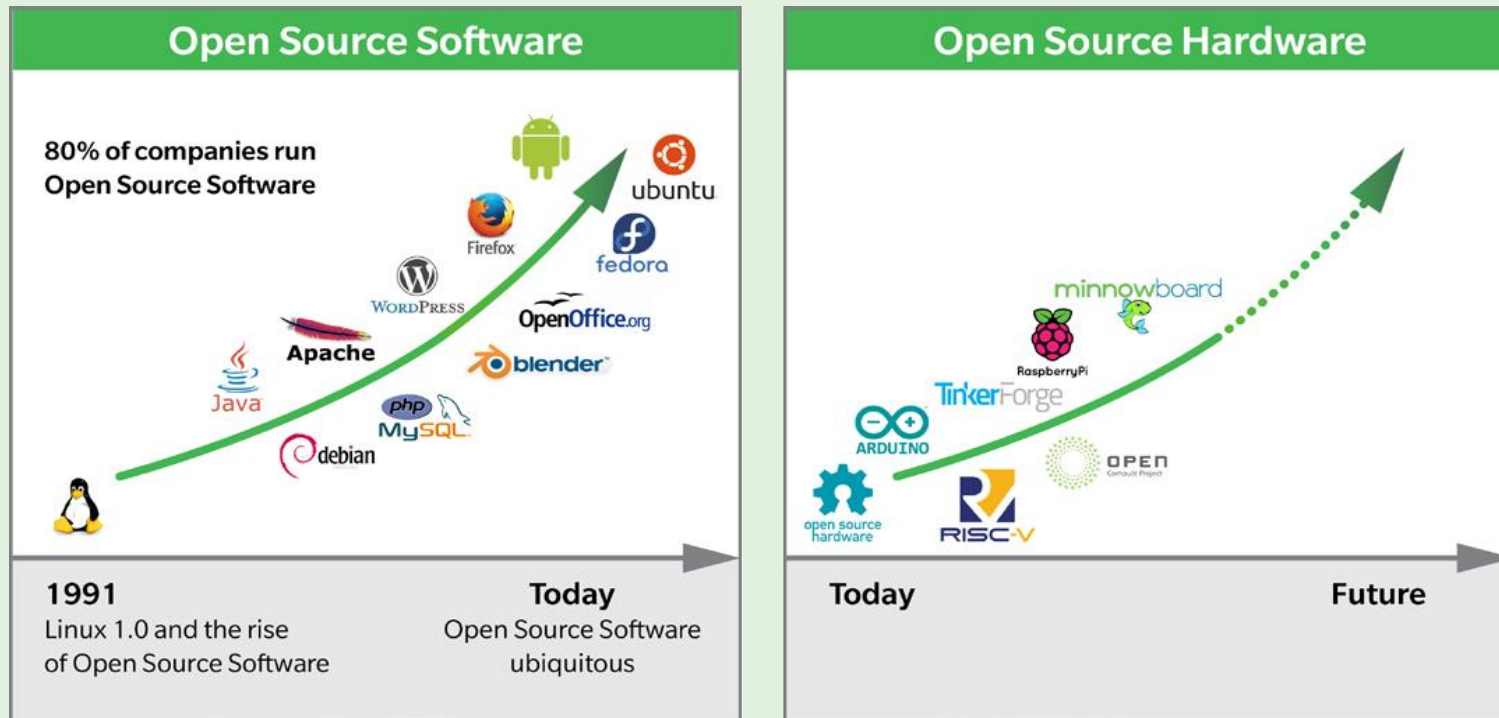
# How do we best encourage architecture diversity?

Open ecosystems allow optimization at every level

# Open Source Hardware

Driving the next wave of innovation

The Rise of Open Source Software: Will Hardware Follow Suit?



- Rapid growth in the adoption and number of open source software projects
- More than 95% of web servers run Linux variants, approximately 85% of smartphones run Android variants
- Will open source hardware ignite the semiconductor industry?  
Is RISC-V the hardware industry's Linux?

# Why Open Source Hardware?

- ▶ **Reducing the cost of development**

- By creating and sharing open hardware (RISCV, OpenSoC)

- ▶ **Closed source IP is a major drag on innovation in all technology spaces**

- Don't need to be a big company to play – lower barrier of entry
- Open nature enables customization at *all* levels of the design – not just at the periphery
- Final product can still be closed

- ▶ **Lower Cost / Complexity for Development**

- Share hardware AND software stack
  - Compilers, debug, Linux ports
- Focus NRE and license on new/innovative IP blocks
- Stop squeezing license costs out of items that students can implement in a summer (license \*hard\* stuff instead)

# Chisel: A DSL for Hardware Design

A productive, flexible language for hardware design and simulation

- ▶ **Increase the productivity of hardware designers**
  - Chisel raises the abstraction level of hardware design
  - Introduces OOP techniques to hardware development
  - Encourages code reuse
- ▶ ***Hardware Generators* are a more efficient technique for generating designs**

- Reduce waste in design
- Reduce design time
- Reduce risk
- Reduce cost






# Chisel Overview

How does Chisel work?

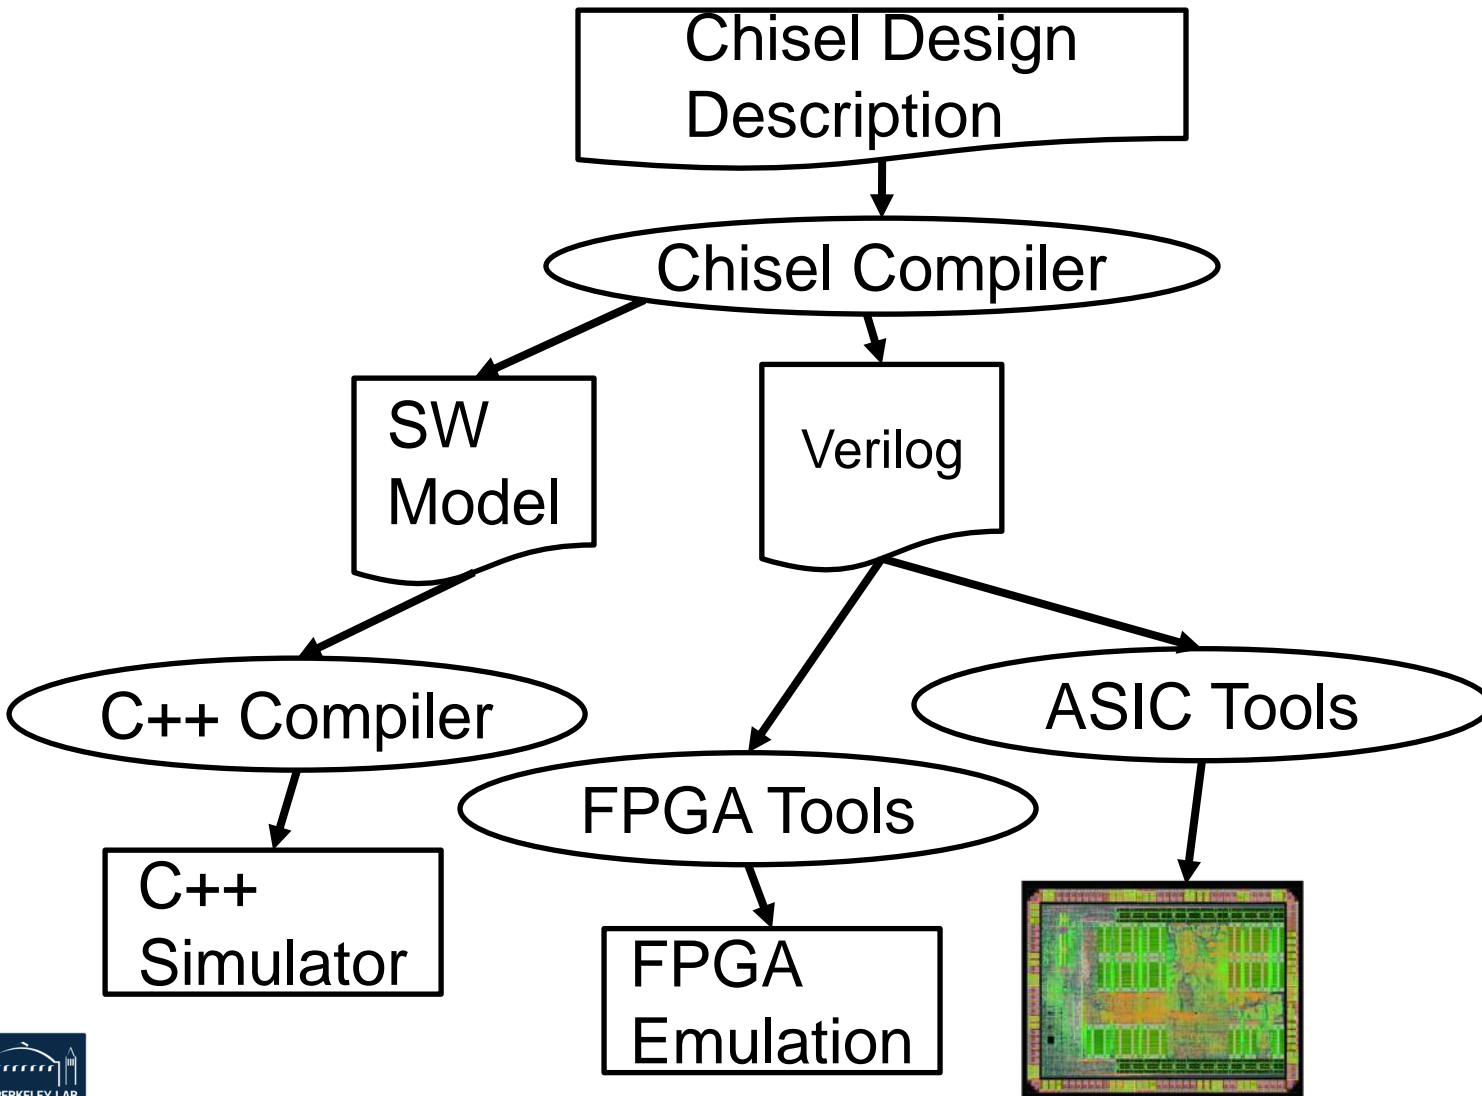
- ▶ **Not “Scala to Gates”**
- ▶ **Describe hardware functionality**
- ▶ **All the abstractions available of a modern language**
  - Now applied to hardware design

```
Mux(x > y, x, y)
```



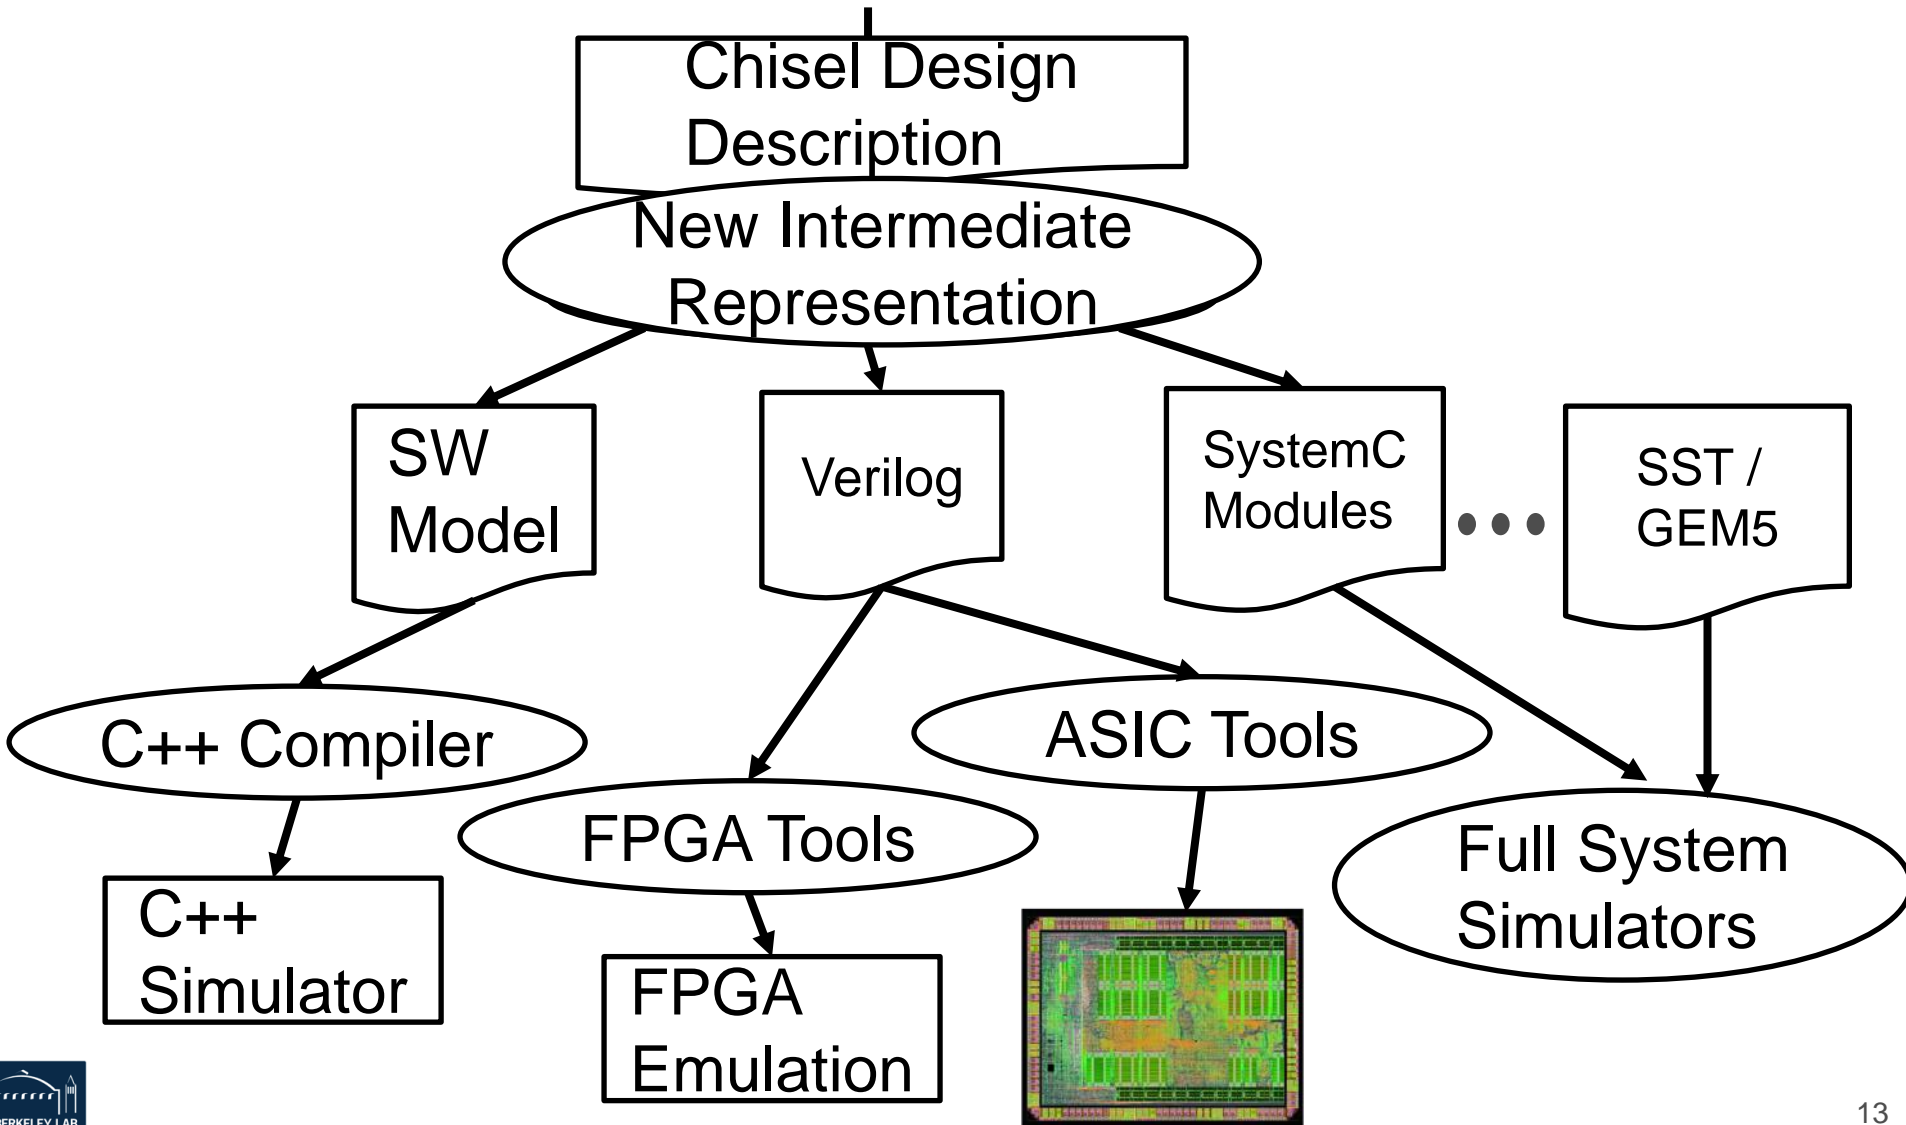
# Chisel: A DSL for Hardware Design

Hardware Generation Infrastructure with Integrated Simulation



# Chisel: A DSL for Hardware Design


Hardware Generation Infrastructure with Integrated Simulation



# RISC V Based Processors

Open source, chisel based processors based on a new ISA

- ▶ **Multiple flavors**
  - Out-of-order (BOOM)
  - In-Order (Rocket)
  - IoT (Z-Scale)
- ▶ **Powerful features**
  - 32, 64, 128-bit addressing
  - Double precision floating point
  - Vector accelerators
- ▶ **Complete SW stack available**
- ▶ **Highly configurable**
  - Only add the features you want!



# “Open Source” *doesn't* mean “Low Performance”

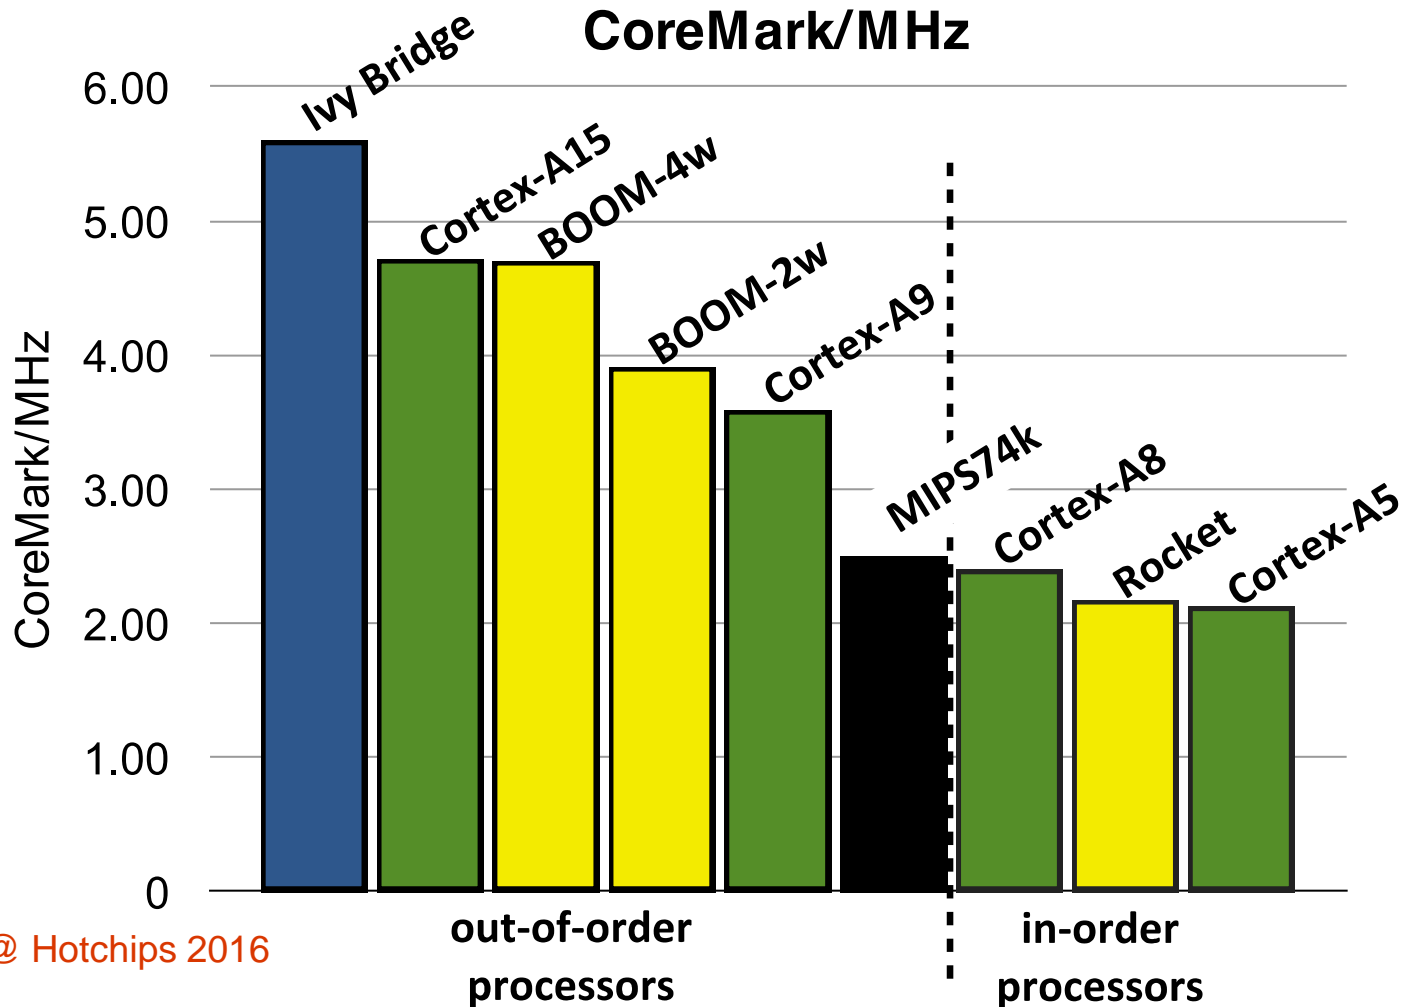
Sometimes you get more than paid for...

Y. Lee UCB @ Hotchips 2016

Category	ARM Cortex A5	RISC-V Rocket
ISA	32-bit ARM v7	64-bit RISC-V v2
Architecture	Single-Issue In-Order 8-stage	Single-Issue In-Order 5-stage
Performance	1.57 DMIPS/MHz	1.72 DMIPS/MHz
Process	TSMC 40GPLUS	TSMC 40GPLUS
Area w/o Caches	0.27 mm <sup>2</sup>	0.14 mm <sup>2</sup>
Area with 16K Caches	0.53 mm <sup>2</sup>	0.39 mm <sup>2</sup>
Area Efficiency	2.96 DMIPS/MHz/mm <sup>2</sup>	4.41 DMIPS/MHz/mm <sup>2</sup>
Frequency	>1GHz	>1GHz
Dynamic Power	<0.080 mW/MHz	0.034 mW/MHz

# “Open Source” *doesn't* mean “Low Performance”

Sometimes you get more than paid for...



Y. Lee UCB @ Hotchips 2016

# Building a HPC System out of RISC-V?

Is it crazier than using ARM?

- ▶ **RISC-V Gaining significant momentum**
  - Large community of SW and HW developers
  - Official ISA of India
- ▶ **Many RISC-V based implementations in the wild**
  - Most not customer facing... yet
- ▶ **Long term investment**

**EE|Times** Connecting the Global Electronics Community

Home | News | Opinion | Messages | Authors | Video | Slideshows | Teardown | Education | EEL

designlines | Planet Analog | Power Management | Programmable Logic | Prototyping

**BREAKING NEWS**

**BLOG: Is Mobileye-Intel New 'Wintel' of Auto?**

designlines OPEN SOURCE

**News & Analysis**

## RISC-V Expands its Audience

**Workshop attracts chip architects, execs**

Rick Merritt

12/1/2016 07:20 AM EST

[Post a comment](#)

NO RATINGS  
1 saves  
[LOGIN TO RATE](#)

[Like](#) 31 [Tweet](#) [in](#) [Share](#) 169 [G+](#) 2

MOUNTAIN VIEW, Calif. — The RISC-V movement is grabbing the attention of a growing set of chip architects and semiconductor executives. Several came to [the group's fifth workshop](#) here to gauge whether the seeds planted by a handful of academics could grow into a disruptive, commercial reality.

The group aims to spread support for its free instruction set architecture across a broad range of products. Talks at the event made clear it will take several years for the ambitious efforts to bear fruit.

Many attendees said they felt exhilarated by the prospects of free, flexible cores unencumbered by patents with an ecosystem of innovations around them. Some feared the efforts could undermine existing markets in an industry already tightening its belt in a cold winter of consolidation.

"RISC-V is the Linux of processor architectures," said Ron Minnich, a Google software developer who [presented his research](#) on creating a hybrid of threads and virtual machines.



More than 330 people registered for the event. (Image: Krste Asanović)

# OpenSoC Fabric

An Open-Source, Flexible, Parameterized, NoC Generator

- ▶ **Greater number of cores per chip driving the evolution of sophisticated on-chip networks**

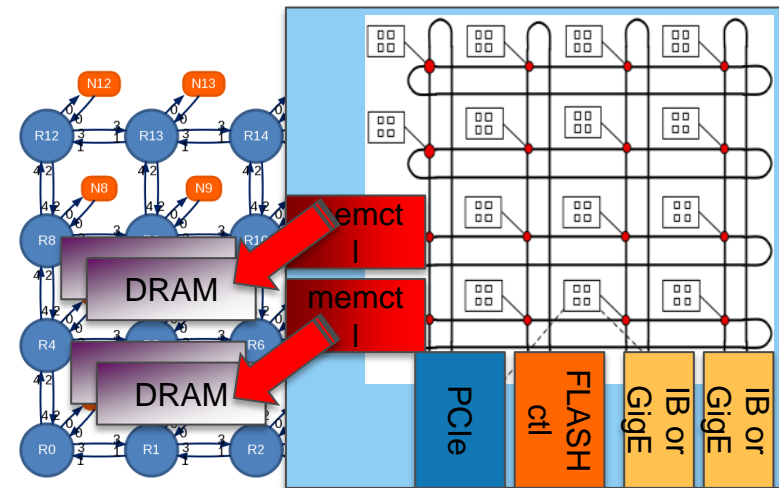

- Needed new tools and techniques to evaluate tradeoffs

- ▶ **Chisel-based**

- Allows high level of parameterization
  - Dimensions, topology, VCs, etc. all configurable
- Fast, functional SW model for integration with other simulators
- Verilog model for FPGA and ASIC flows


- ▶ **Multiple Network Interfaces**

- Integrate with wide variety of existing processors
  - Tensilica, RISC-V, ARM, etc.





# SC15 Demo: 96 Core SoC Design for HPC

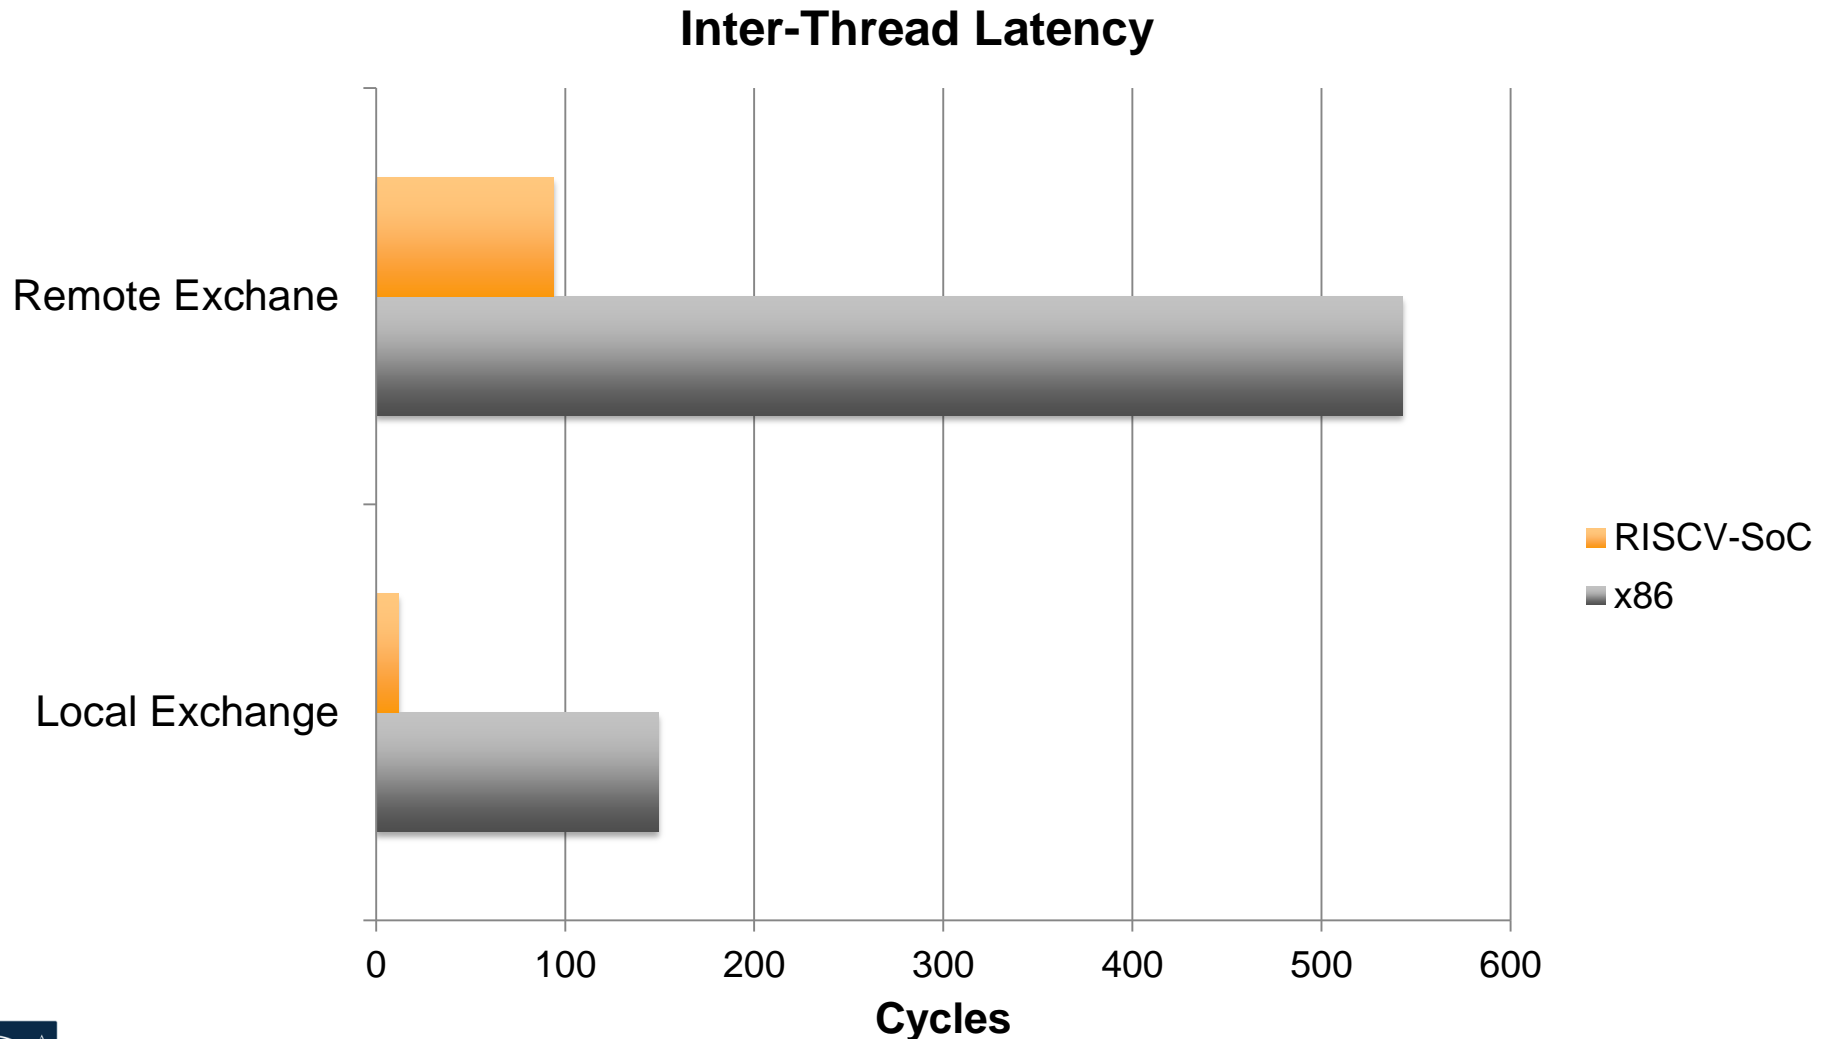


*2 people spent 2 months to create*

- ▶ **Z-Scale processors connected in a Concentrated Mesh**
- ▶ **4 Z-scale processors**
- ▶ **2x2 Concentrated mesh with 2 virtual channels**
- ▶ **Micron HMC Memory**

# 96 Core System: Results

Comparing conventional cache coherence protocol  
To direct hardware support for global address space for halo exch.



# If you thought 96 cores was cool...

GRVI Phalanx Accelerator (Jan Gray)

- ▶ **1680 32-bit RISC-V Cores**
  - 30 rows x 7 Col
- ▶ **26MB SRAM**
- ▶ **300bit NoC**








# Accelerating the Design Process

Creating a complete suite of tools for rapid processor and compiler generation




## OpenSoC Fabric




# Accelerating the Design Process


Creating a complete suite of tools for rapid processor and compiler generation



**OpenSoC**




**OpenSoC Fabric**




# Accelerating the Design Process


Creating a complete suite of tools for rapid processor and compiler generation



**OpenSoC**



**OpenSoC Fabric**




**OpenSoC Compiler**




# Accelerating the Design Process


Creating a complete suite of tools for rapid processor and compiler generation



**OpenSoC**



**OpenSoC Fabric**



**OpenSoC Compiler**




**OpenSoC Cores**






# Accelerating the Design Process


Creating a complete suite of tools for rapid processor and compiler generation



**OpenSoC**



**OpenSoC Fabric**



**OpenSoC Compiler**



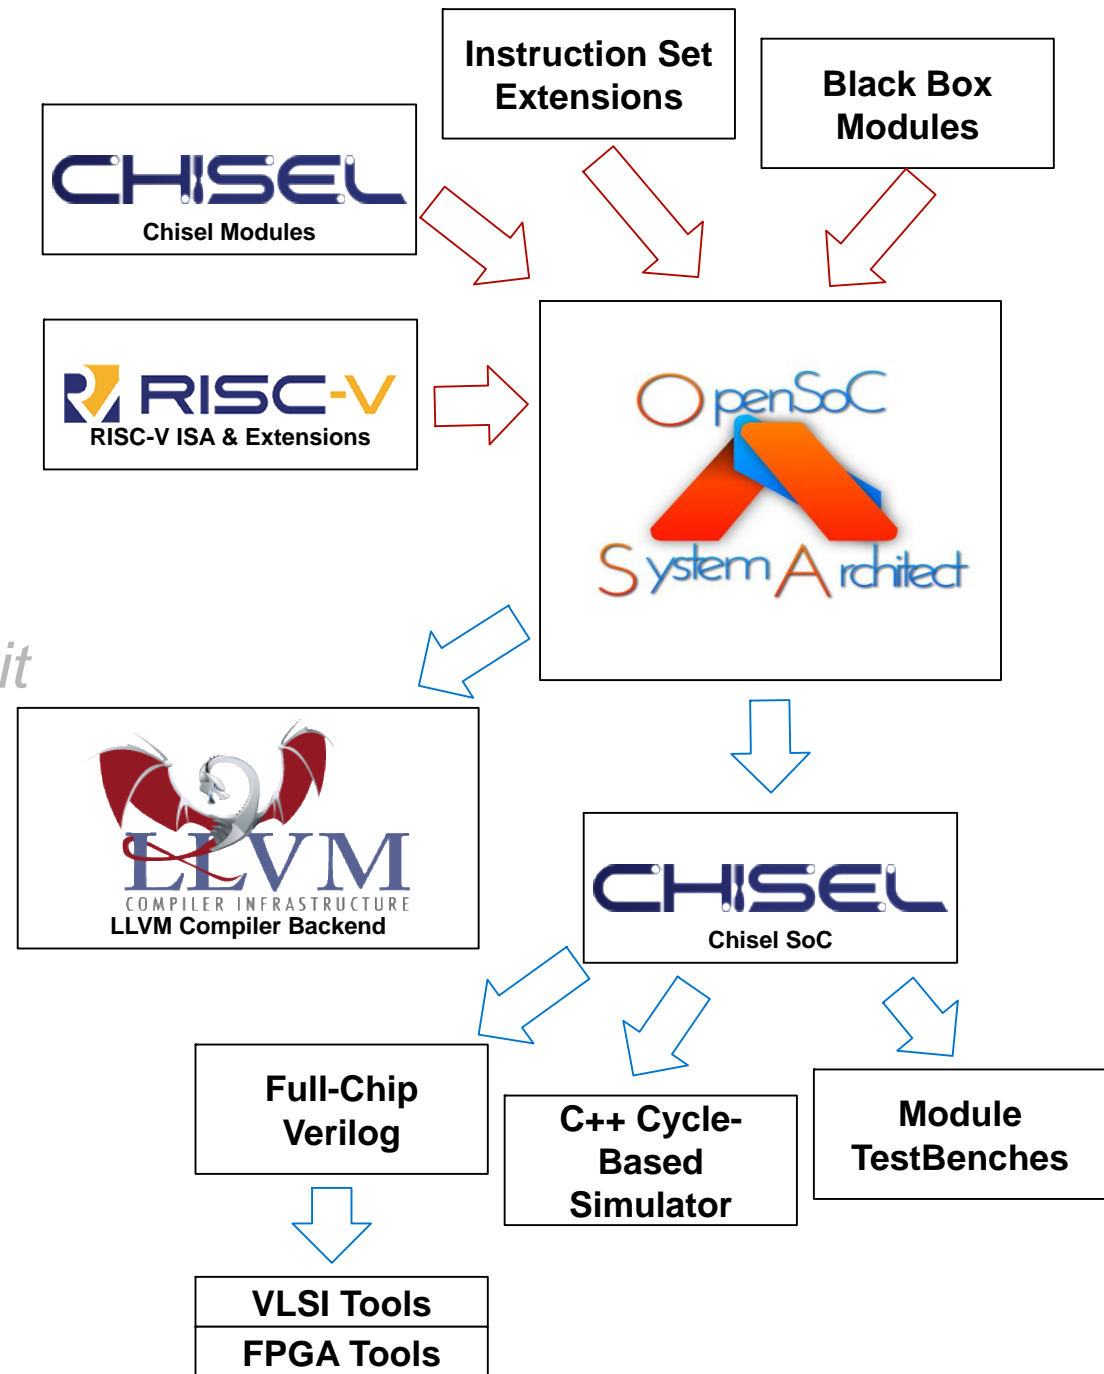
**OpenSoC Cores**



**OpenSoC System Architect**

# OpenSoC System Architect

*A complete hardware and software development toolkit*

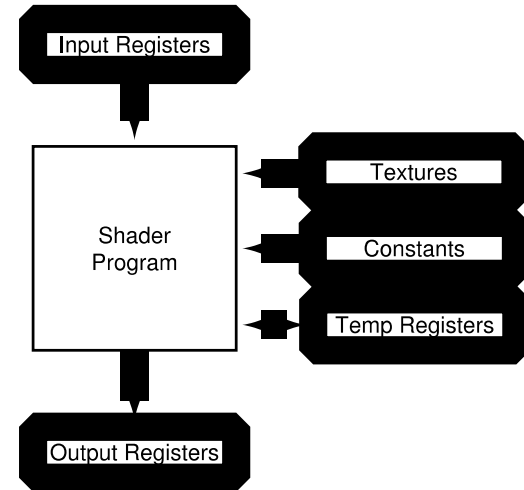



# Specialization and FPGAs Look Compelling...



What do we need to make these concepts mainstream?

# Programmability

Remember 2004?



```
typedef OctreeGpu<vec3f, vec4ub> OctreeType; // 3D float addresses, RGBA values  
OctreeType octree( vec3i(2048, 2048, 2048) ); // effective size 20483
```



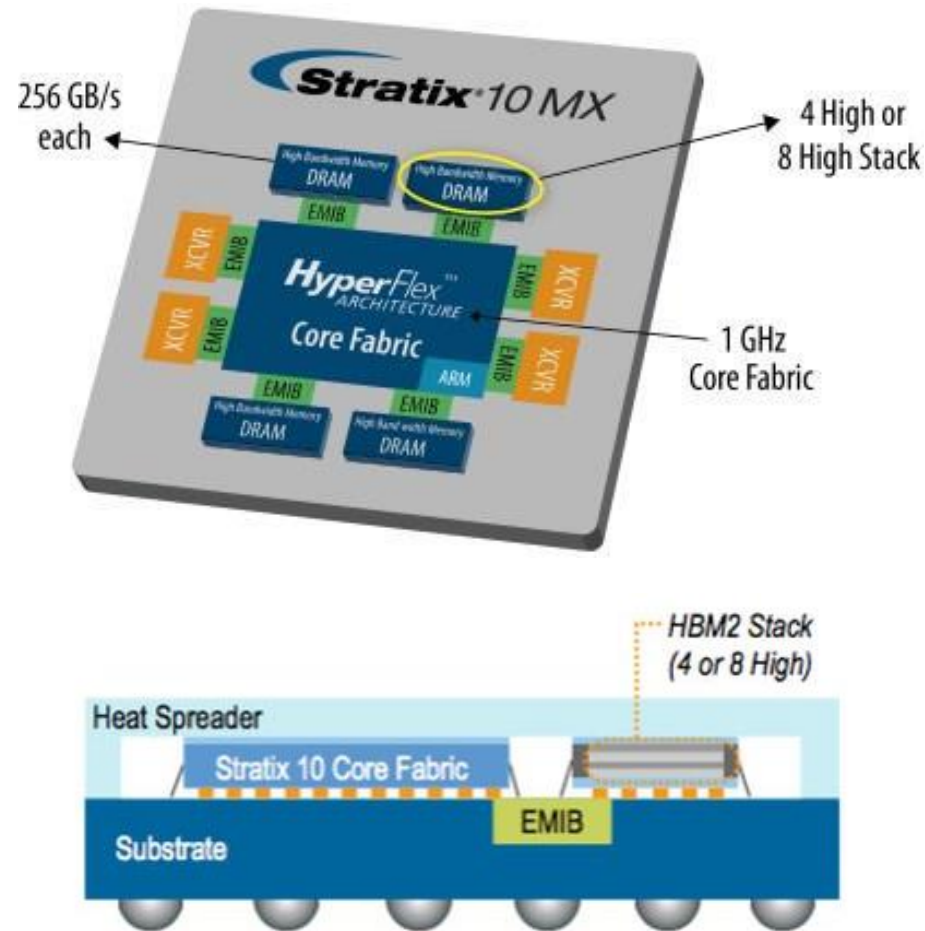
Buck, Owens, Riffel, Lefhon , et al.

# New FPGAs are very capable

Riding Moore's Law to the end...

## ▶ FPGA + Stacked Memory SIP

- 1GHz Fabric
- Quad core A53
- >6 TFlops Single Precision
- 16GB HBM on package, 1TB/s BW
- Tb/s IO



# Programmability

FPGA ecosystem promoting development of new tools and abstractions to accelerate development




- ▶ **Parallels with early days of GPGPU computing**

- VERY capable hardware

- ▶ **New languages raising abstraction levels**

- OpenCL
- OpenACC

- ▶ **But the tools are still terrible**



**Altera® SDK**  
for  
**OpenCL™**

**XILINX**  
ALL PROGRAMMABLE™

# Programmability

No need to translate your algorithm directly to hardware

## ▶ **Previous model:**

- Translate your algorithm from Fortran/C/etc. to a hardware description
- Highest performance gain, but most significant development effort

## ▶ **New model:**

- Instantiate an array of specialized “soft” cores that can be targeted much like a GPU
- Greater flexibility, simpler hardware development
- Enables acceleration of applications not typically associated with FPGA computing

# Creating an architecture per motif

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

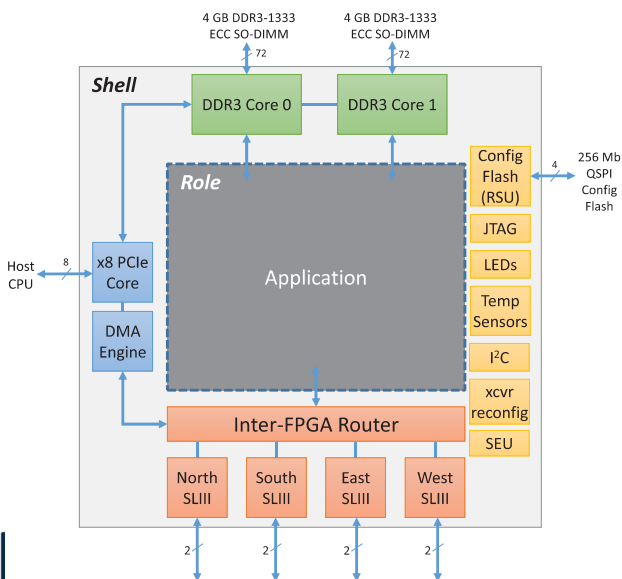



# Catapult


Real world deployment of FPGA accelerators

## ▶ Microsoft Catapult (ISCA 2014)

- Bing search acceleration
- Lower Cost and Power
- Programmable
- Composed of customized soft cores

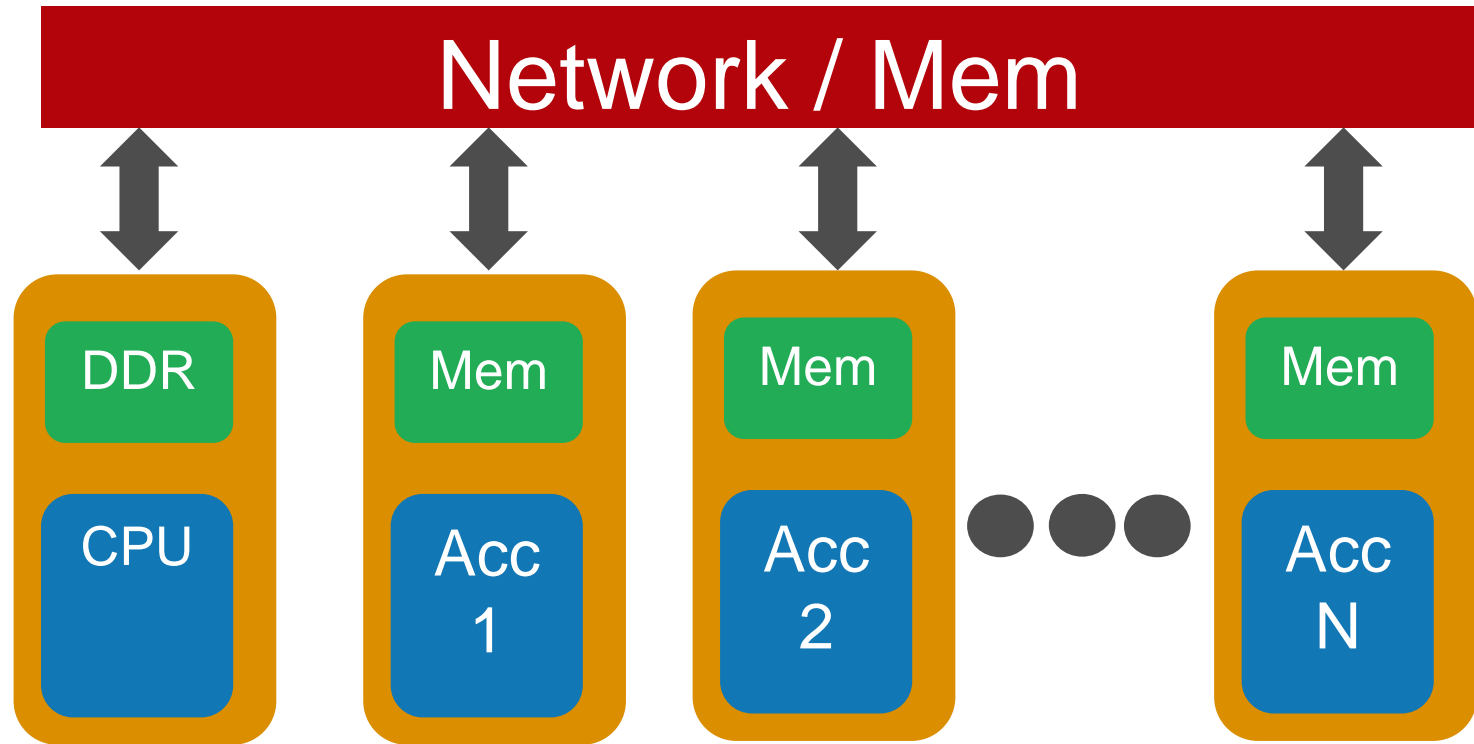


1,632 Servers with FPGAs Running Bing Page Ranking Service (~30,000 lines of C++)



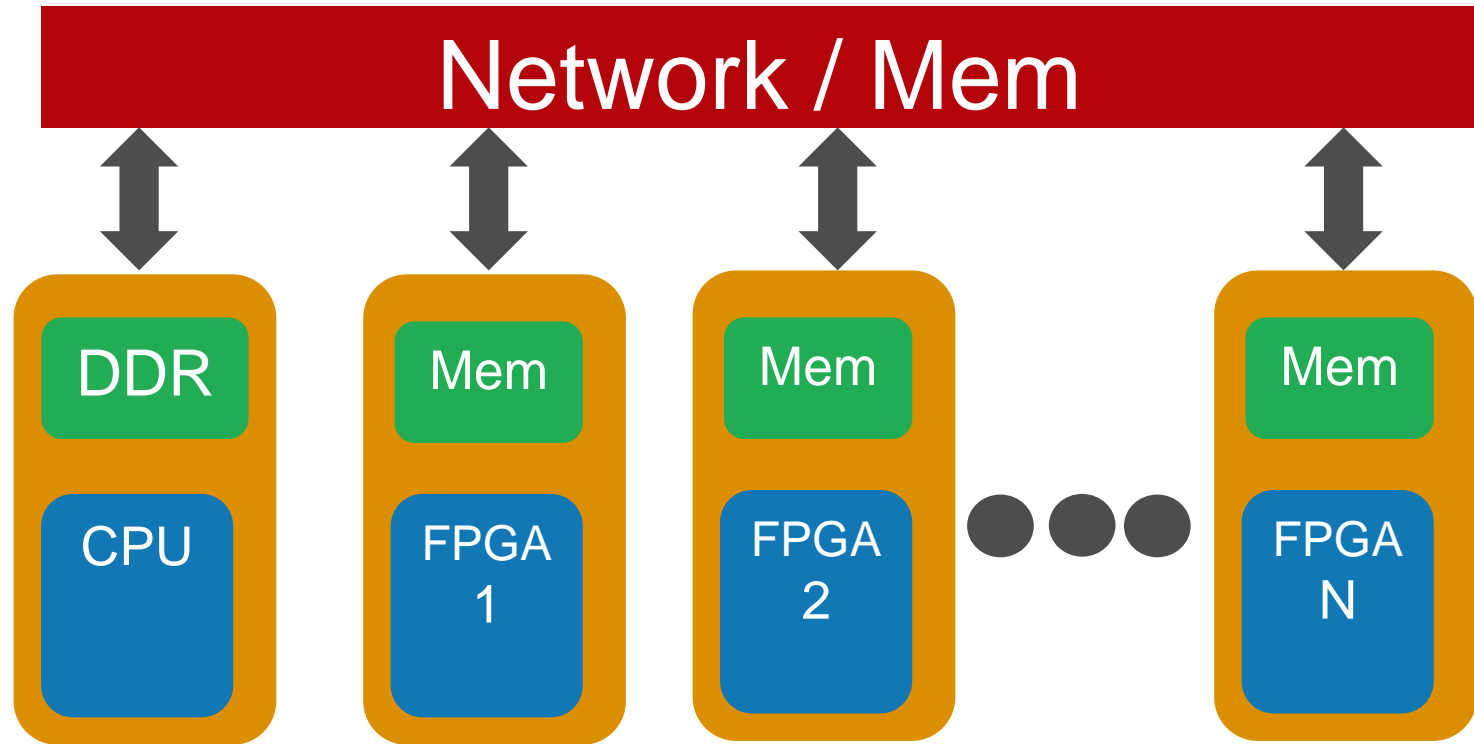
# A Potential Exascale Node

A notional representation using specialization



# A Potential Exascale Node

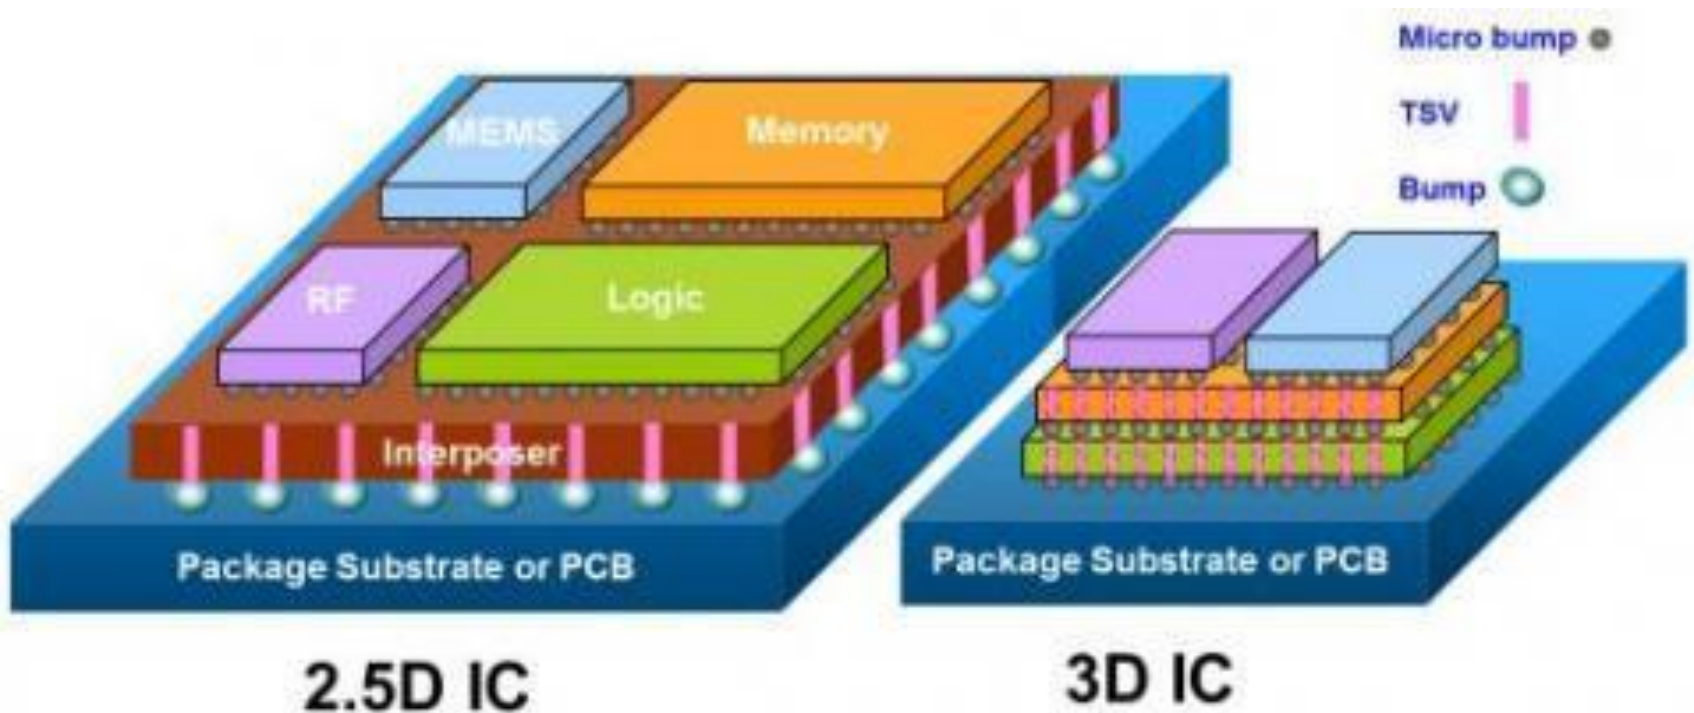
A notional representation using specialization



# Advanced Manufacturing

## 3D Stacking

- ▶ Can combine processing and memory in single package with efficient interconnect



# Looking out to the edge...

Applying these tools and techniques to other scientific / computational areas

# On-detector processing



Putting our hardware design tool suite to work to augment existing HPC resources

## ▶ The Problem:

- Future detectors threaten to overwhelm data transfer and computing capabilities w/ data rates exceeding 1 Tb/s
- Data processing experiment driven

## ▶ Proposed solution:

- Process the data *before it leaves the sensor*
- Application-tailored, programmable processing allows data reduction to occur on-sensor
- Programmability allows data reduction techniques to be tailored to the experiment – even *after* the sensor is built!



# What's Next?

- ▶ **Need to continue to deliver increased performance scaling**
- ▶ **FPGAs emerging as the next computational element**
  - Moving away from monolithic implementations to specialized soft cores
- ▶ **Can we get some better tools?**
  - HW design getting better, but not enough
  - Tune Chisel compiler for FPGAs
- ▶ **Already being deployed at scale**
  - Catapult
  - AWS
    - Bitfile marketplace

***None of this matters without advanced software – programming models, runtimes, etc***

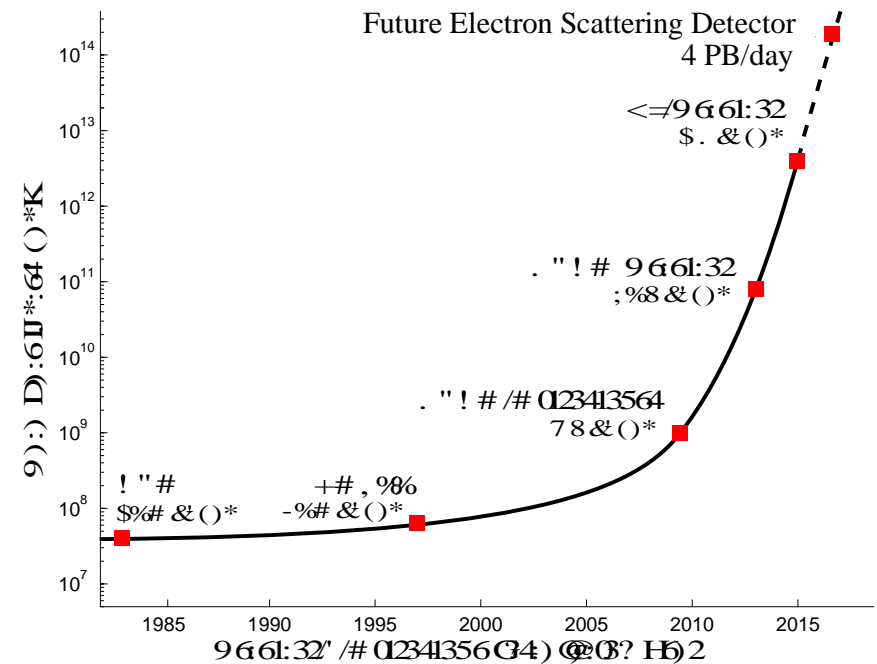

# Thank you!

ddonofrio@lbl.gov



## ▶ Backup


# Future Electron Scattering Detector




- 100,000 fps pixel detector
  - 576 x 576 x 10 #m
- Segmented silicon HAADF
- Fabricate detectors Q4CY2016

- Dedicated (donated) 400 Gbs link to NERSC
  - Link testing underway
- Stream events to processors on Cori
- Future goal: firmware processing (reduce data rate)





# P5\$5(R5\$'3(NA+1%B'3;



o' (' 1>' (/ 1[W-\	
'	JK <sup>i</sup> JK <sup>h</sup> JK <sup>JK</sup> JK <sup>JJ</sup>
z	JK <sup>i</sup> JK <sup>U</sup> JK <sup>l</sup> JK <sup>JJ</sup>
2	JK <sup>i</sup> JK <sup>h</sup> JK <sup>JJ</sup>

Q? Y. (P' \$' # \$+ &

dT

OP(AQ? .

e/- &6  
!'X' 3

n53\$! | P

: ' & n53\$! | P

' #8"4

P5\$5(3' \$3(5& (3' #+263(\$+(B +2\$83

P5\$5(85\$' 3(5& ;(Q+65/ ;(KN(Q N7 TKJU;(Q N7 TKTK;(JK(Q N7

Y2(56B '\$\$' 6/(45&+ #8'51(S" G4+'2\$