SDC

STORAGE DEVELOPER CONFERENCE

SNIA SANTA CLARA, 2017

Memories of Tomorrow

Coughlin Associates

Tom Coughlin, Coughlin Associates

&

Jim Handy, Objective Analysis

How Have New Memories Gotten Funding?

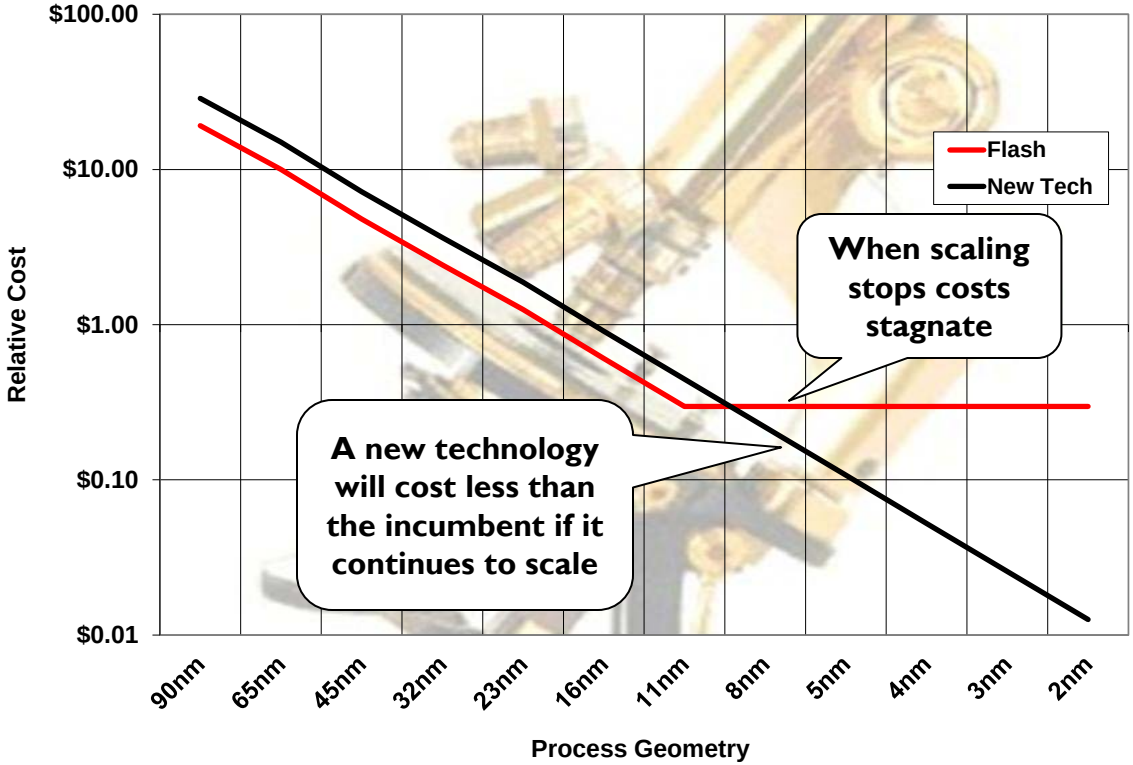

The Scaling Limit Has Been Looming For Years

New Memories Extend Scaling

Limited Scaling

- Flash
- DRAM

Extended Scaling

- ReRAM
- MRAM
- PCM
- FRAM
- CBRAM
- NRAM

But Wait! There's More!

- New memories are all much better than Flash & DRAM:
  - Nonvolatile (persistent)
  - Bit writes
  - No pre-erase requirement
  - Much faster than flash
  - Lower power (no refresh, low write energy)
  - Cheaper (through scaling)

Storage Class Memory!

What is Storage Class Memory?

Storage-class memory (SCM) combines the benefits of a solid-state memory, such as high performance and robustness, with the archival capabilities and low cost of conventional hard-disk magnetic storage.

IBM Almaden Research Labs

What is Storage Class Memory?

Storage-class memory (SCM) combines the benefits of a solid-state memory, such as high performance ~~and robustness~~, with the archival capabilities ~~and low cost~~ of conventional hard-disk magnetic storage.

IBM Almaden Research Labs

Well, That's Just Core!

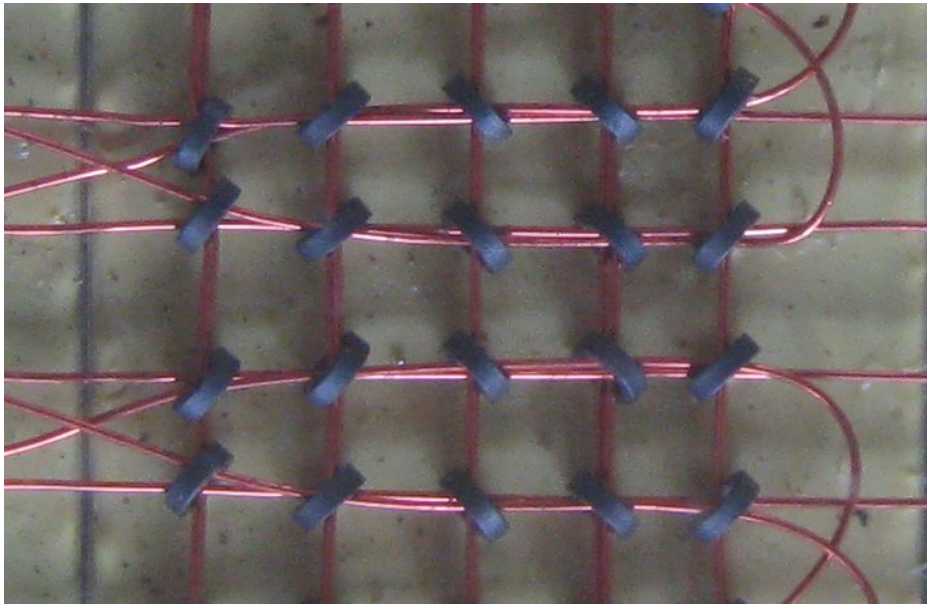

- Goal: Rethink Everything
  - From Turing & von Neuman to now

- Why IEEE?
  - Encompasses the whole computing stack

Tom Conte, Co-chair IEEE Rebooting Initiative, Professor, Georgia Tech

Potential Approaches vs. Disruption in Computing Stack

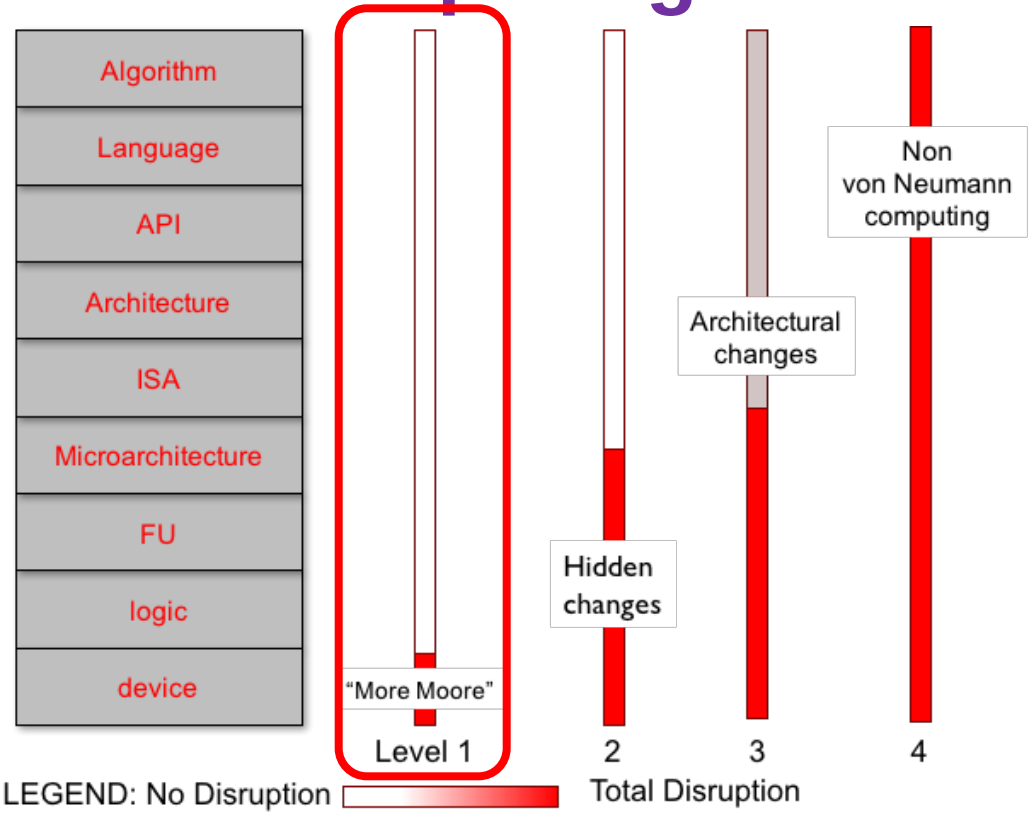

More Moore

- Further scaling
  - New device structures:
    - FinFET, 3D NAND, new switch, unreliable switch, etc.
- Non-silicon device types:
  - MRAM
    - Spintronic systems
  - PCM
  - Memristors
  - FRAM

Potential Approaches vs. Disruption in Computing Stack

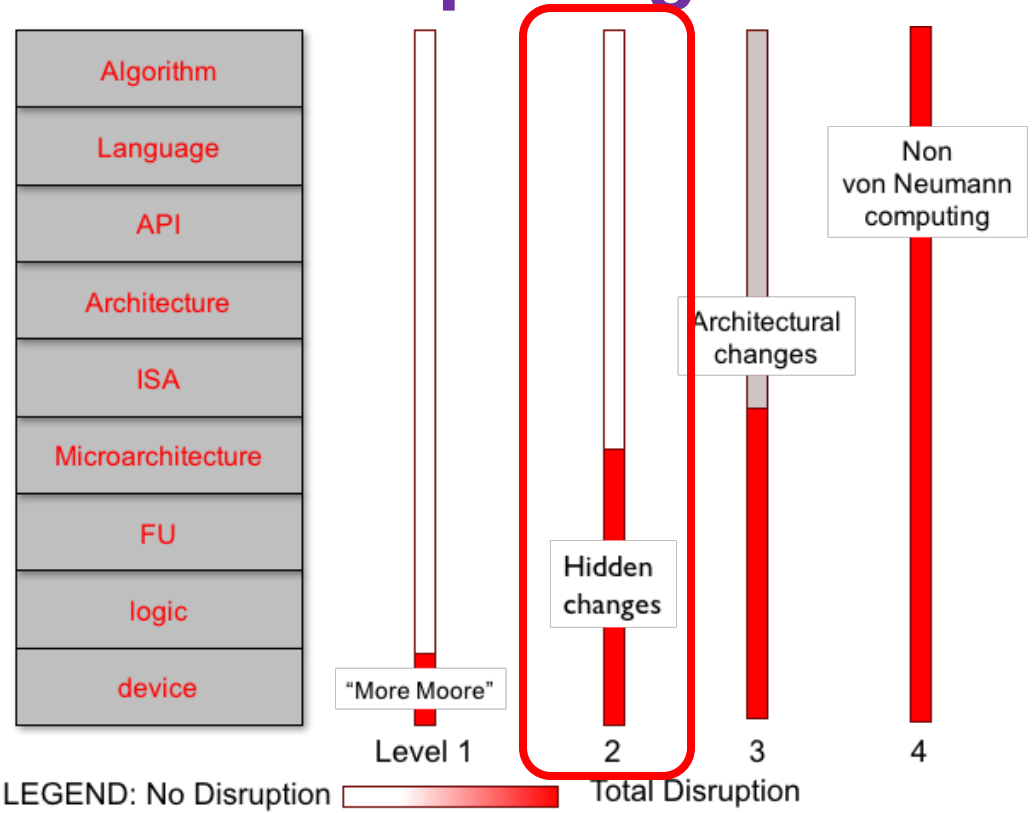

Hidden Changes

- Changes that don't impact the Instruction Set
  - Devices
  - Logic
  - Functional Units
  - Microarchitecture
- Context switch issue fits here (described shortly)
- Microarchitectural support for non-deterministic logic devices
- No significant impact to memory or storage

Potential Approaches vs. Disruption in Computing Stack

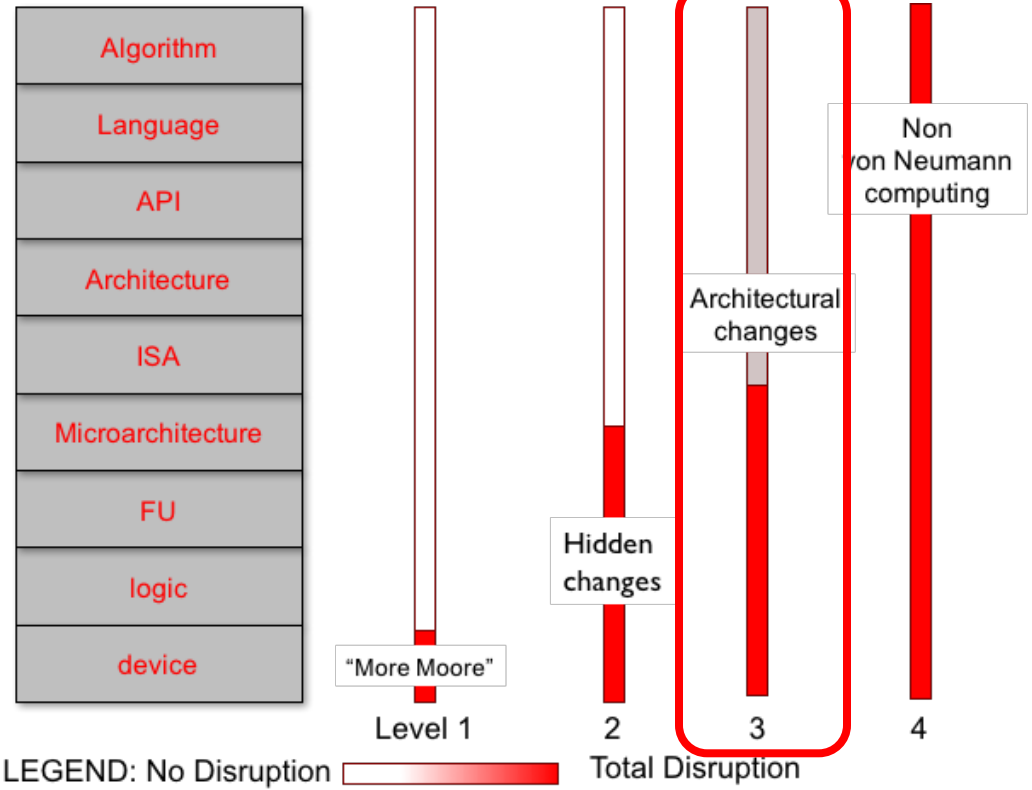

Architectural Changes

- Beyond x86
  - GPU
  - FPGA
- Moving computation to the data
  - In-memory compute
  - *In-situ* processing
- Special purpose computing
- Requires new programing

Potential Approaches vs. Disruption in Computing Stack

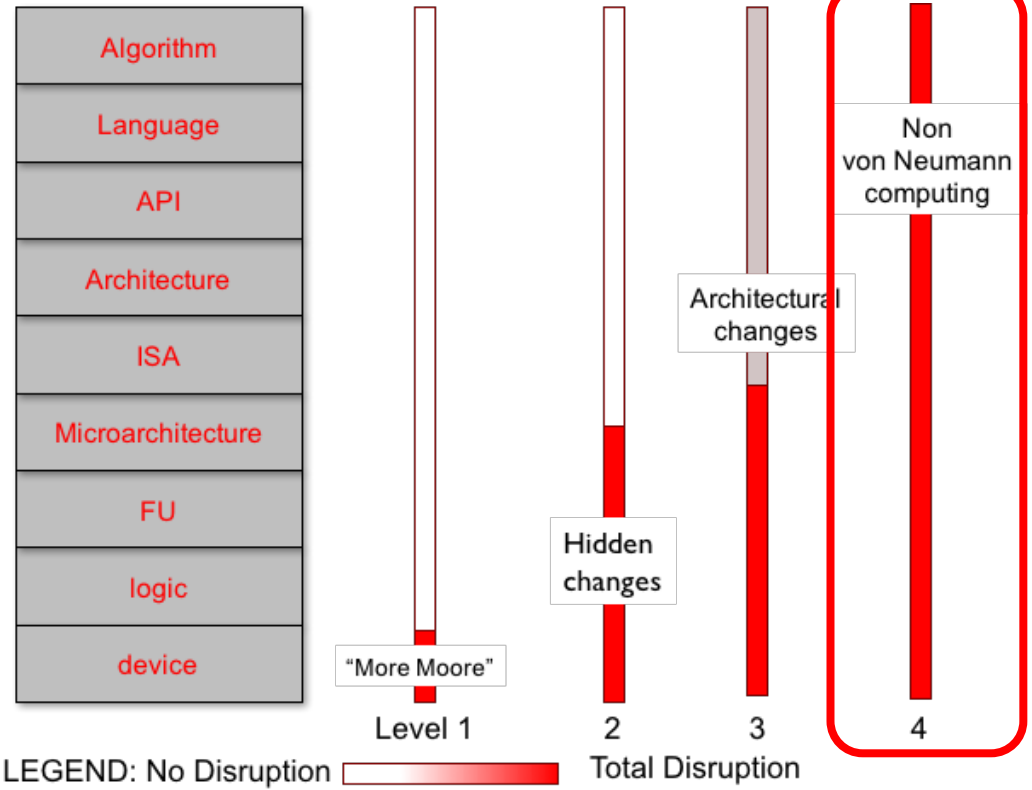

Non von Neumann Architectures

- Quantum computers
- Neuromorphic computing
- Analog Computing
- Very immature, lots of investment

Artificial Intelligence useful in:

- Internet of Things
- Voice and Image recognition
- Autonomous Vehicles
- Cybersecurity
- Industrial automation
- And many other uses

All of These Need Storage & Memory

Two Major Problem Areas:

- Memory loses its data (volatility)
- Storage is slow (latency)

Curing Latency

- Move compute close to the data
- Use lower-latency storage

More on these later

Curing Volatility

- Old Way: Frequent verified copies to storage
  - Faster storage helps this a lot!
- New Way: Make the memory persist
  - Batteries
  - Flash backups
  - New persistent memory types

That's Storage Class Memory (SCM)!

Solution 1: Persistent Memory

- NVDIMM
- Optane
- Other emerging memories
- Issues: Context switch, cost, maturity, S/W support…

NVDIMM

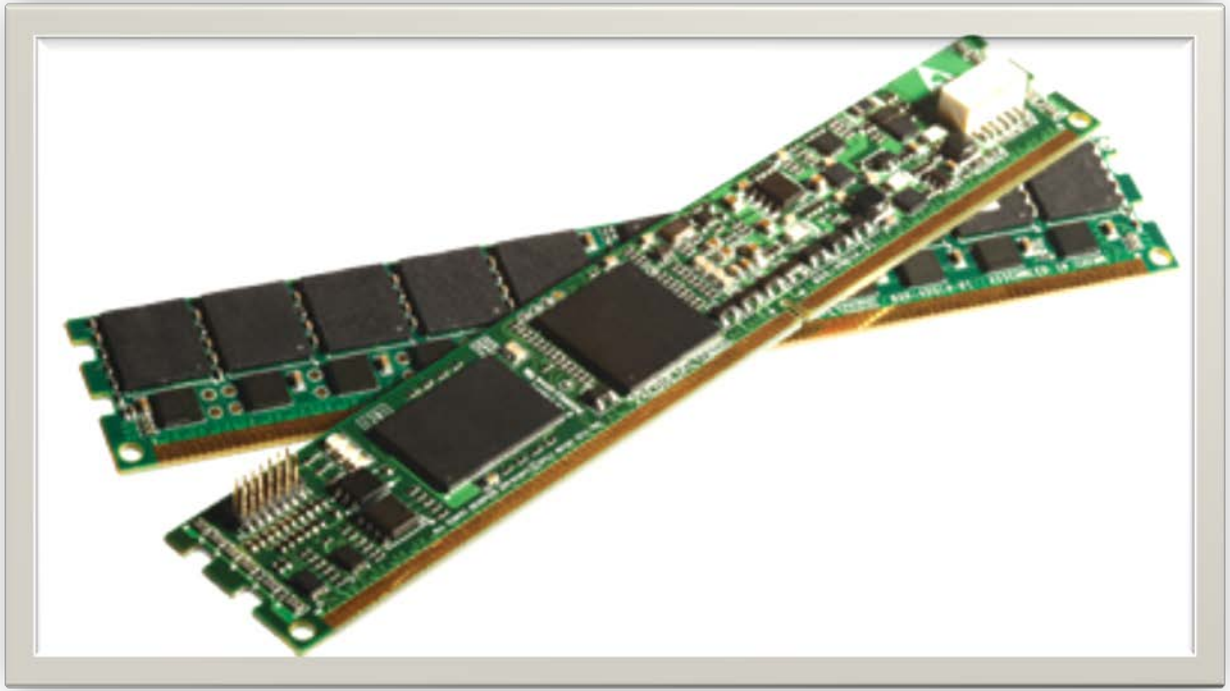

Today's Memory Types

	SRAM	DRAM	ROM	EEPROM	NOR	NAND
Nonvolatile	No	No	Yes	Yes	Yes	Yes
Erasable	Yes	Yes	No	Yes	Yes	Yes
Programmable	Yes	Yes	Factory	Yes	Yes	Yes
Smallest Write	Byte	Byte	N/A	Byte	Byte	Page
Smallest Read	Byte	Page	Byte	Byte	Byte	Page
Read Speed	V Fast	Fast	Fast	Fast	Fast	Slow
Write Speed	V Fast	Fast	N/A	Slow	Slow	Slow
Active Power	High	Med	Med	Med	Med	Med
Sleep Power	V Low	High	Zero	Zero	Zero	Zero
Price/GB	High	Low	Low	High	Med	V Low
Applications	Small Fast	Main Memory	Stable Code	Serial #, Trim	Code	Data

New Memories Perform Better

	MRAM	ReRAM	FRAM	PCM	XPoint
Nonvolatile	Yes	Yes	Yes	Yes	Yes
Erasable	Yes	Yes	Yes	Yes	Yes
Programmable	Yes	Yes	Yes	Yes	Yes
Smallest Write	Byte	Byte	Byte	Byte	Byte
Smallest Read	Byte	Byte	Byte	Byte	Byte
Read Speed	Fast	Fast	Fast	Fast	Fast
Write Speed	Fast	Fast	Fast	Fast	Fast
Active Power	Low	Med	Low	High	High?
Sleep Power	Low	Low	Low	Low	Low
Price/GB	High	High	High	High	High?
Applications	Niche	TBD	Low Power	Obsolete	Main Memory

Solution 2: Bring Compute to Storage

- *In Situ* Processing: Kinetic, WDLabs Microserver, NGD, NVXL, ScaleFlux
- Compute in Memory: TOMI, Automata
  - Why it's failed
- Old failed storage/processing models: Schooner, Virident, Violin

Compute In SSD: *In Situ* Processing

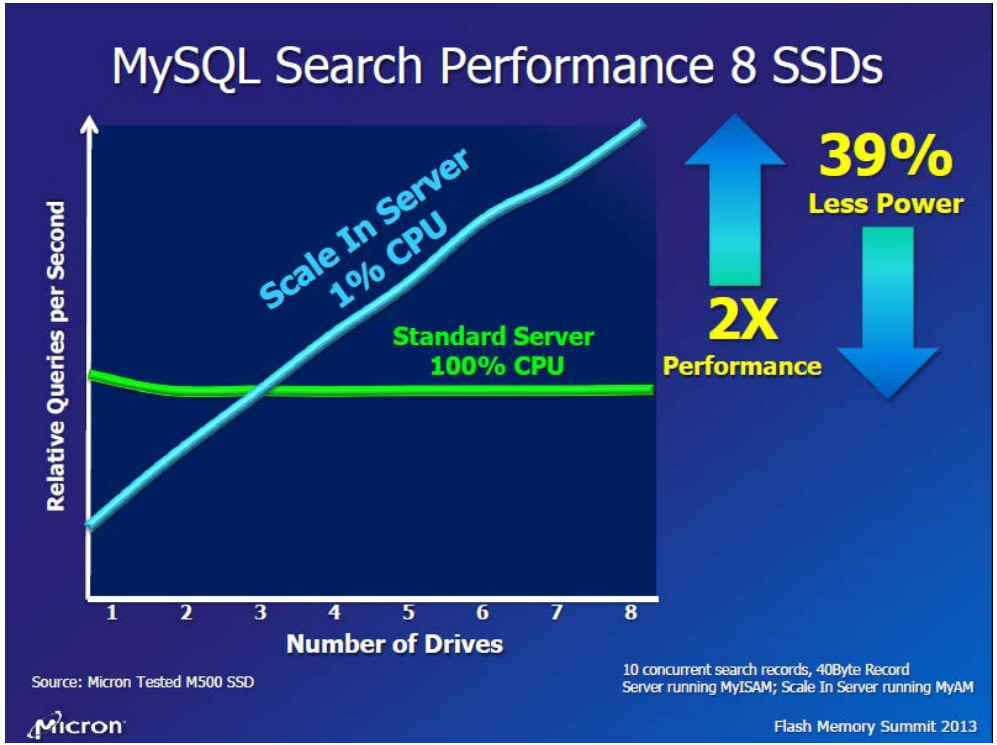

Compute In Memory

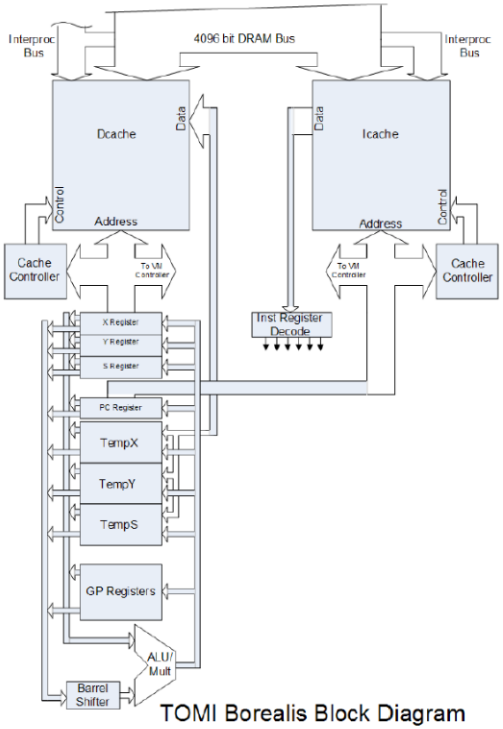

TOMI Borealis Block Diagram

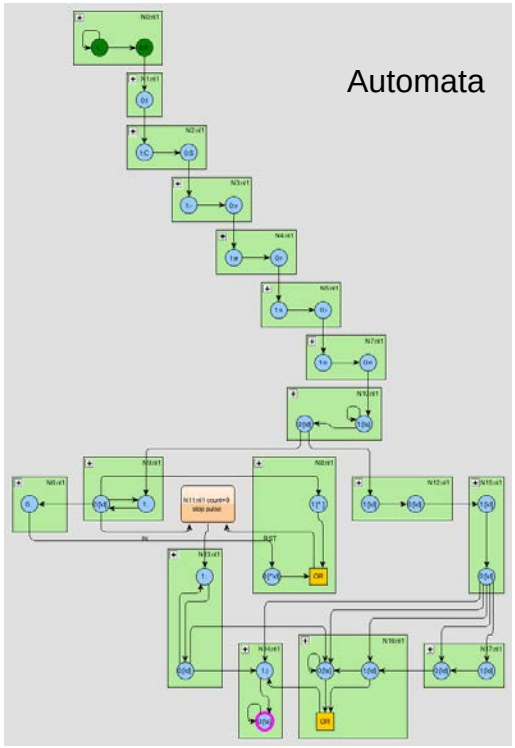

Compute In Memory

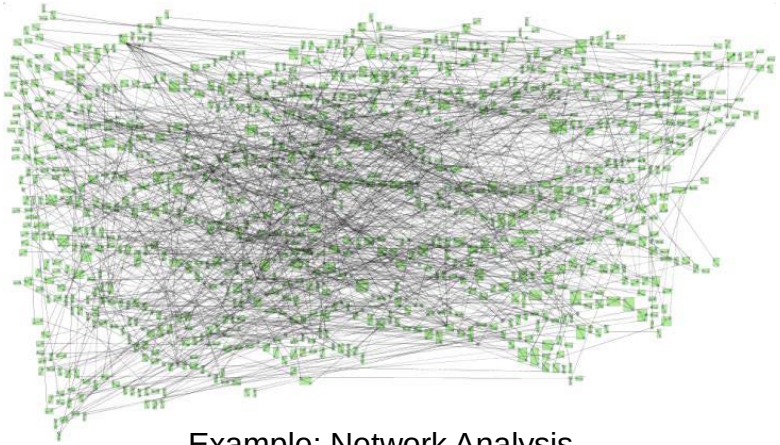

Example: Network Analysis

Compute Appliances

Schooner

MySQL

memcached

Virident

Violin Memory with Microsoft Windows Server

All Require Support Throughout The System

- Software support
- Firmware (BIOS) support
- CPU support
- Even new pins on the bus!

Example: Let's look at Persistent Memory (PM)

PM Needs New Software

- Move away from “Storage vs. Memory” approach
  - Store at the byte level, not blocks
  - Avoid the storage stack
  - Avoid things like flash translation

The SNIA Persistent Memory Programming Model

- https://www.snia.org/PM

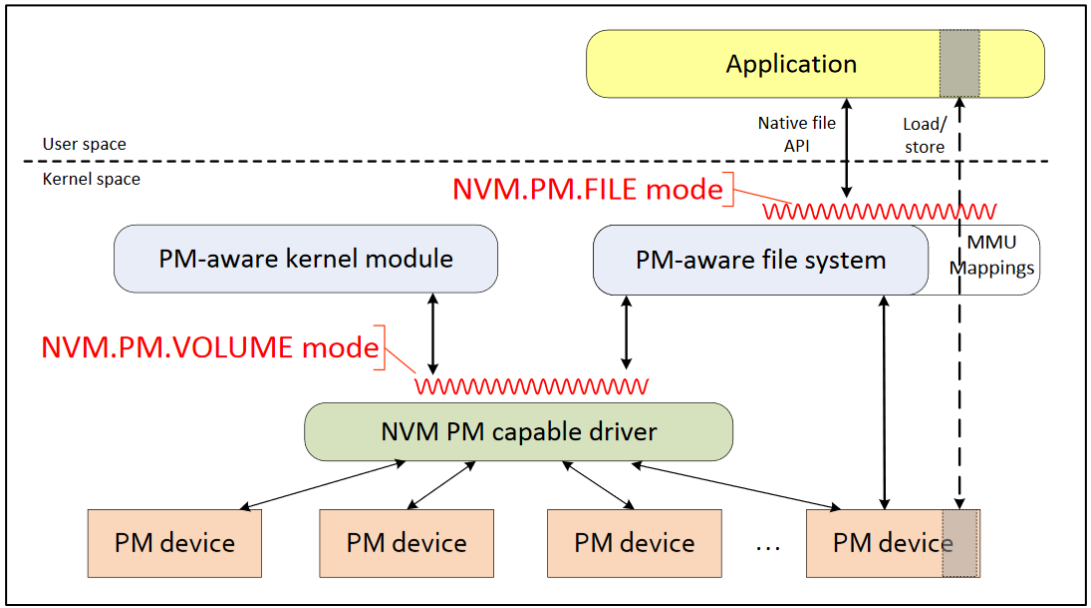

PM Needs BIOS Support

Current BIOS

- DRAM corrupt at boot
- Power fail loses everything:
  - Memory
  - Registers
  - Cache

PM BIOS

- DRAM may boot with valid data
- Push DRAM to PM at power fail:
  - Memory
  - Registers
  - Cache

PM Needs Processor Support

- Intel's new IA instructions
  - Cache & write buffer flush
  - NVDIMM-N "Commit": Copy DRAM to NAND
- Context switch issue?

Context Switches Become The Issue

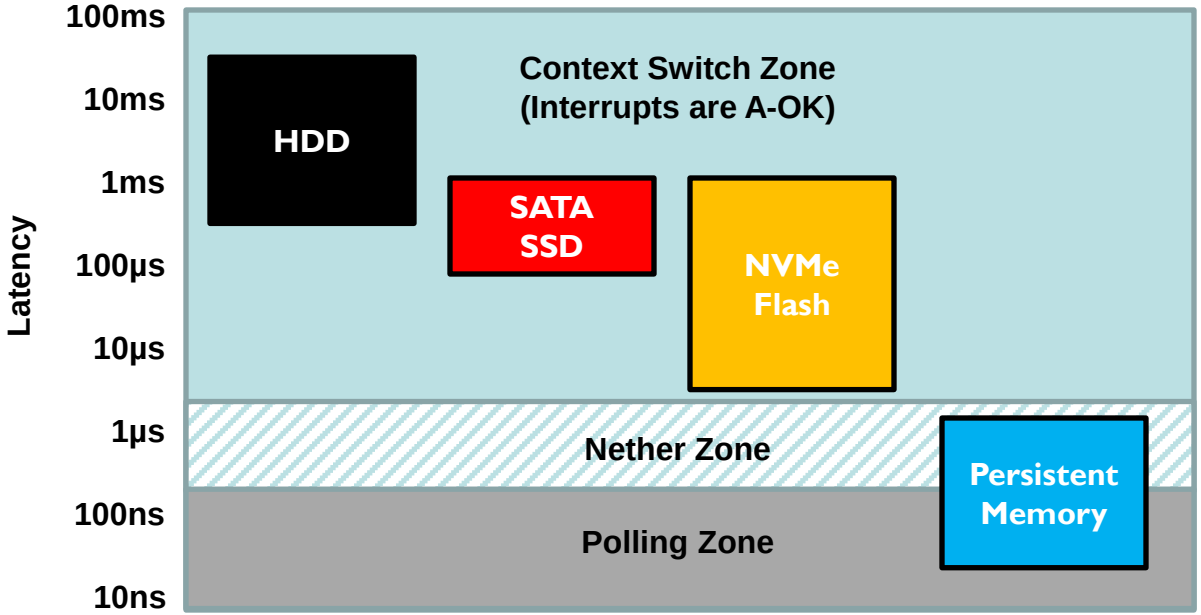

Advent of SSDs Was Similar

- New software required
  - Caching/tiering
    - Blossomed into Software-Defined Storage
  - Faster I/O stacks
- New hardware
  - SATA 2 & 3
    - SAS kept pace
  - NVMe for PCIe SSDs

Other Alternatives Will Need Similar Efforts

- In-situ processing:
  - Server vs. drive-based software
  - What to do about HA?
- How does this work with compute and memory virtualization?

Summary

- Change is in the air
  - Persistence everywhere
  - New ways to bring processing to storage
  - New approaches to computing
- A lot of support will be required
  - Software
  - Architecture
  - Hardware

SDC

STORAGE DEVELOPER CONFERENCE

SNIA SANTA CLARA, 2017

Questions?

Your Presenters

Thomas Coughlin
Coughlin Associates

Tom Coughlin, President, Coughlin Associates is a highly-respected storage analyst and consultant with over 30 years in the data storage industry in engineering and management at high profile companies.

Jim Handy
Objective Analysis

Jim Handy is a widely recognized semiconductor analyst, has over 30 years in the electronics industry. His background includes marketing and design positions at market-leading suppliers.

Source Material

- **Tom Conte, Georgia Tech,** The IEEE Rebooting Computing Initiative: Rethinking All Levels of How We Compute, Presentation at Sandia, July 2017
- **SNIA Persistent Memory Programming Model,** https://www.snia.org/PM
- **National Council on Aging**, https://www.ncoa.org