Session:

Enabling Nanosecond-Class Storage

Orthagonal Spin Transfer MRAM – A Better Approach

Next Generation Persistent Memory – Beyond NVDIMM N

Followed by a Panel Discussion

Enabling Nanosecond-Class Storage

Steffen Hellmold | Western Digital Corporation

FORWARD-LOOKING STATEMENTS

This presentation contains forward-looking statements that involve risks and uncertainties, including, but not limited to, statements regarding our addressable market, our product and technology positioning and compute platforms, the anticipated benefits of our new technologies, executing on our integrated strategic plans, realizing our strategic imperatives, including our memory card solutions and storage technologies. Forward-looking statements should not be read as a guarantee of future performance or results, and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all. Forward-looking statements are subject to risks and uncertainties that could cause actual performance or results to differ materially from those expressed in or suggested by the forward-looking statements.

Additional key risks and uncertainties include the impact of continued uncertainty and volatility in global economic conditions; actions by competitors; difficulties associated with go-to-market capabilities; business conditions; growth in our markets; and pricing trends and fluctuations in average selling prices. More information about the other risks and uncertainties that could affect our business are listed in our filings with the Securities and Exchange Commission (the “SEC”) and available on the SEC’s website at www.sec.gov, including our most recently filed periodic report, to which your attention is directed. We do not undertake any obligation to publicly update or revise any forward-looking statement, whether as a result of new information, future developments or otherwise, except as otherwise required by law.

Enabling Next Gen Compute Platforms

- Data wants to be close to the CPU
  - Rise of data centric compute architectures
- Big Data needs acceleration
  - GPU / FPGA accelerators augmenting CPU
- Disaggregation using high speed fabrics
  - Interconnects for compute, memory and storage

Data Centric Compute Architectures

Data Center

Rack Scale Architecture

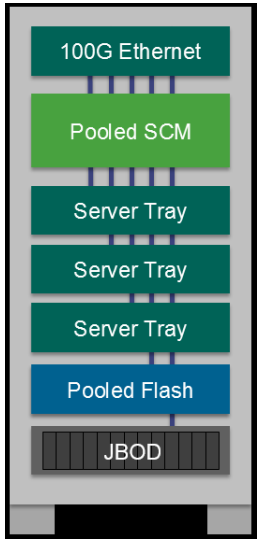

Data Center

Cheap CPUs Around PB of SCM

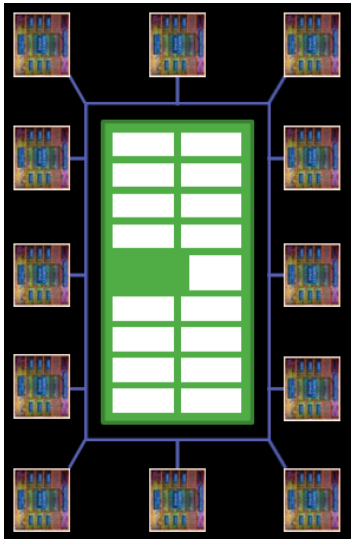

Client Compute

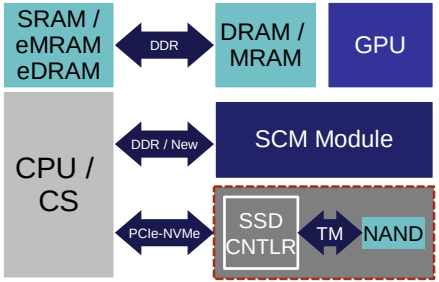

SCM complements DRAM for compute intensive clients

Mobile

Large memory requirements for Virtual Reality

TIME

Extending Storage to 250ns Latency

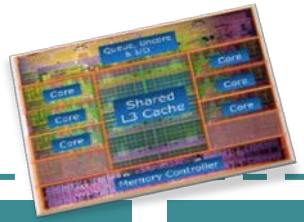

Core Register	Core L1 Cache	Core L2 Cache	Shared L3 Cache	DRAM	Storage Class Memory	Flash	HDD
Size	64KB	256KB	2-4MB	16-128GB	128GB-1TB	512GB-4TB	4-16TB
Speed	1ns	3-10ns	10-20ns	50-100ns	250-5,000ns	100,000ns-2,000,000ns	5-10,000,000ns
Cost				100x	20-25x	5x	1x

Source: Western Digital estimates

ns-NVM Catalyst Advancing Compute

- Fabric Attached Storage enables shared resource pool
  - New architectures optimize total cost of ownership
- On the memory bus enables near and far NVM
  - Near Memory – DRAM-like / Far memory close to DRAM speed
- Core register and caches enable all NVM compute
  - True instant-on / off enables most power efficient computing

Memory, Storage Fabric Standardization

- Transition from CPU to non-volatile memory centric architectures
- Industry standards efforts ongoing:
  - NVDIMM-P / NVDIMM-N
  - Gen-Z
  - CCIX
  - OpenCAPI
  - Rapid-IO

2018

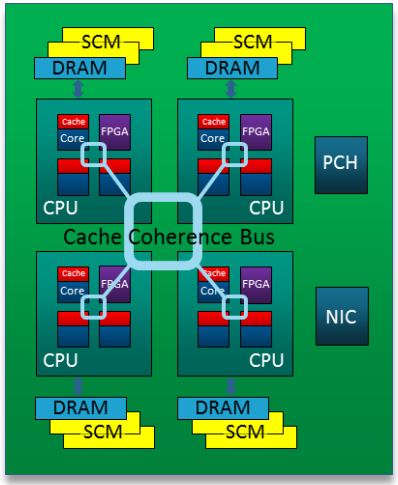

CPU-centric architecture

2020

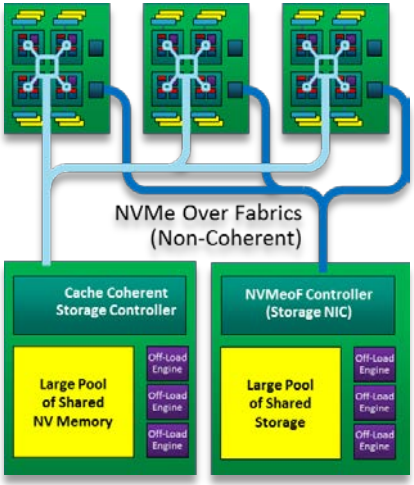

NVM centric architecture

Open industry standards key to broad-based nanosecond-class storage adoption and capital investments

Summary

- Nanosecond-class storage helps scale compute
- Introduces low and high-nanosecond storage tiers
- High-speed fabrics enable NVM and storage pools

Open industry standards for memory bus and fabric attached nanosecond-class storage will enable broad market adoption

Thank You

Spin Transfer Technologies
An Allied Minds Company

Orthogonal Spin Transfer (OST) MRAM
A Better Approach

Barry Hoberman | Spin Transfer Technologies

Leading ST-MRAM Development

- Exceptional Funding: $108M
- Exceptional Magnetics and CMOS team
- Unique technology to create fastest/highest-endurance STT-MRAM
- Complete, fast-cycle magnetics R&D fab

Momentum: In 15 months, STT has gone from 1st pMTJ development to pMTJ customer samples

Advanced Technology: Benchmark Performance at sub-50nm pMTJ Feature

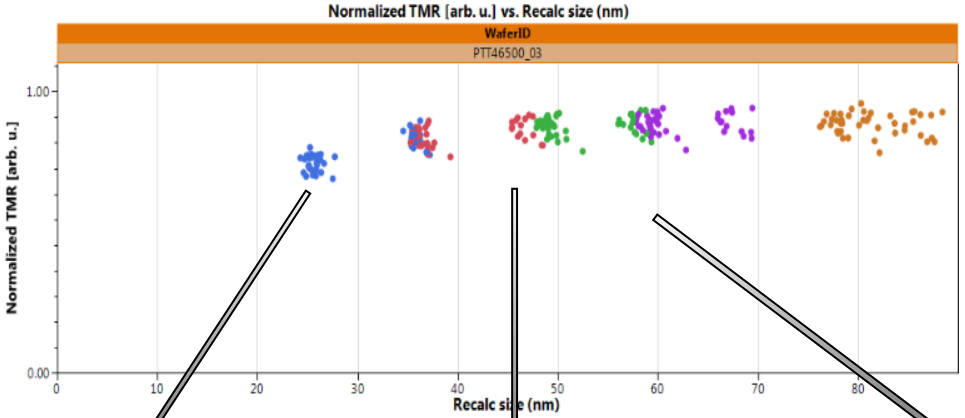

~20 nm pillar process demonstrated

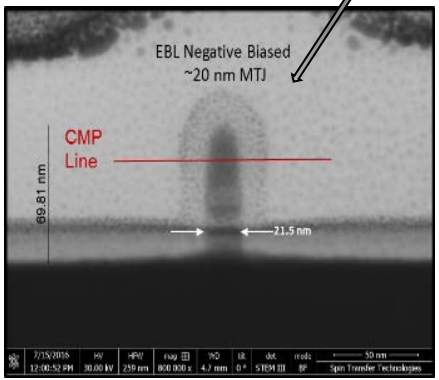

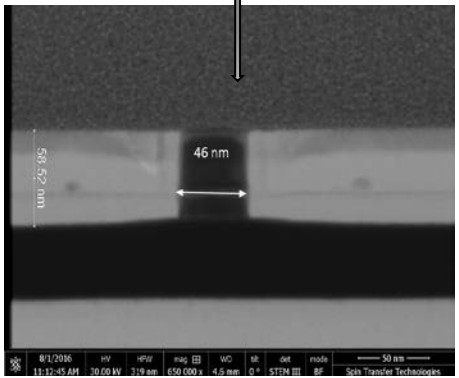

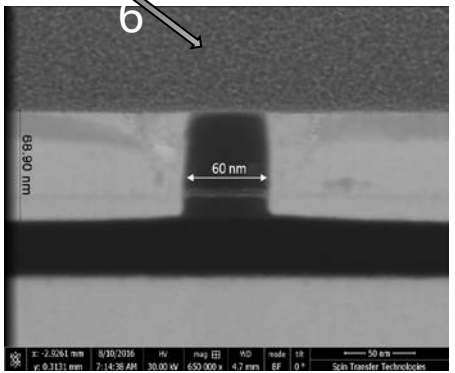

ST-MRAM Challenges

- Probabilistic effects
  - Write errors, bit flips, etc.
- Fabricating millions or billions of pMTJ's… with semiconductor yields
  - Decades of process experience in legacy CMOS structures
- Different application requirements
  - DRAM / SRAM replacement
  - eNOR replacement
- Emergence of high throughput 300/200mm wafer process equipment
- Many others…

Tug-of-War Among Specs

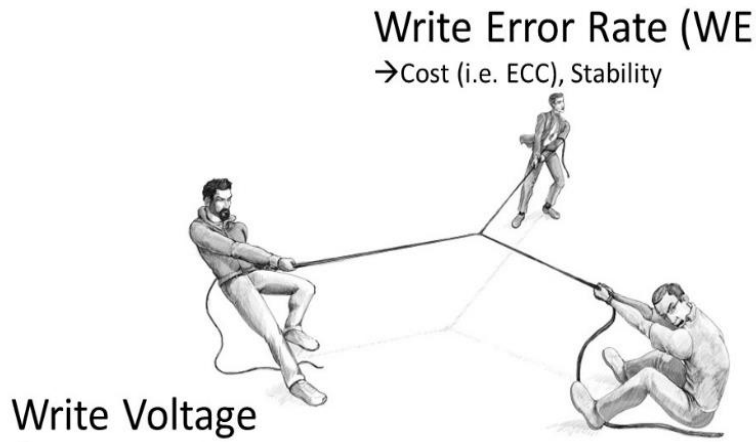

ST-MRAM Challenges: Resolved

Conventional ST-MRAM

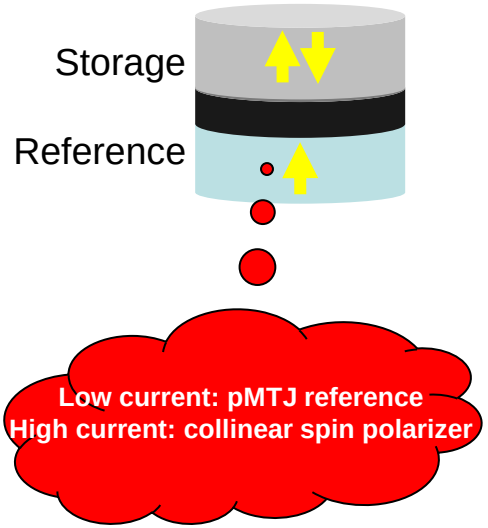

Spin Transfer Technologies
An Allied Minds Company

+

Proprietary architecture for writing, reading

High current: 'orthogonal' spin polarizer

Materials + Magnetics

Uniquely Combined With

Circuits + Systems

Produces the best <u>simultaneous</u> combination of:

- **Read/write speed**
- **Endurance**
- **Retention**
- **Low Power**
- **Low Cost**
- **Manufacturability**

Trends in ST-MRAM Market Uptake
-- Standalone

- 2018-2020
  - 1 Gigabit, DDRx…production
- Watch for:
  - Specialty memory in low energy, high endurance
  - Trade-off challenges in retention, speed/throughput, and endurance

Trends in ST-MRAM Market Uptake
-- Embedded

- embNVM
  - Foundry adoption in 2x nm and lagging nodes
  - Strong market pull in ‘low energy’…IoT, battery, etc
    - >>1000x write energy advantage relative to embFlash, plus performance advantage due to absence of ‘block erase’
- embRAM
  - Be wary of sacrifice of retention or endurance in high-speed (sub 20ns R/W cycle) implementations
    - Conventional ST-MRAM has no work-around

Summary

- The advent of pMTJ technology matches roadmap requirements below 65nm CMOS, for both standalone and embedded memory
- Timelines are settling now for mass production in the next 2-3 years

Next Generation Persistent Memory Evolution beyond the NVDIMM-N

Doug Finke, Xitore, Inc.

NVDIMM-P can make current DRAM/NAND Competitive with more Exotic New Media

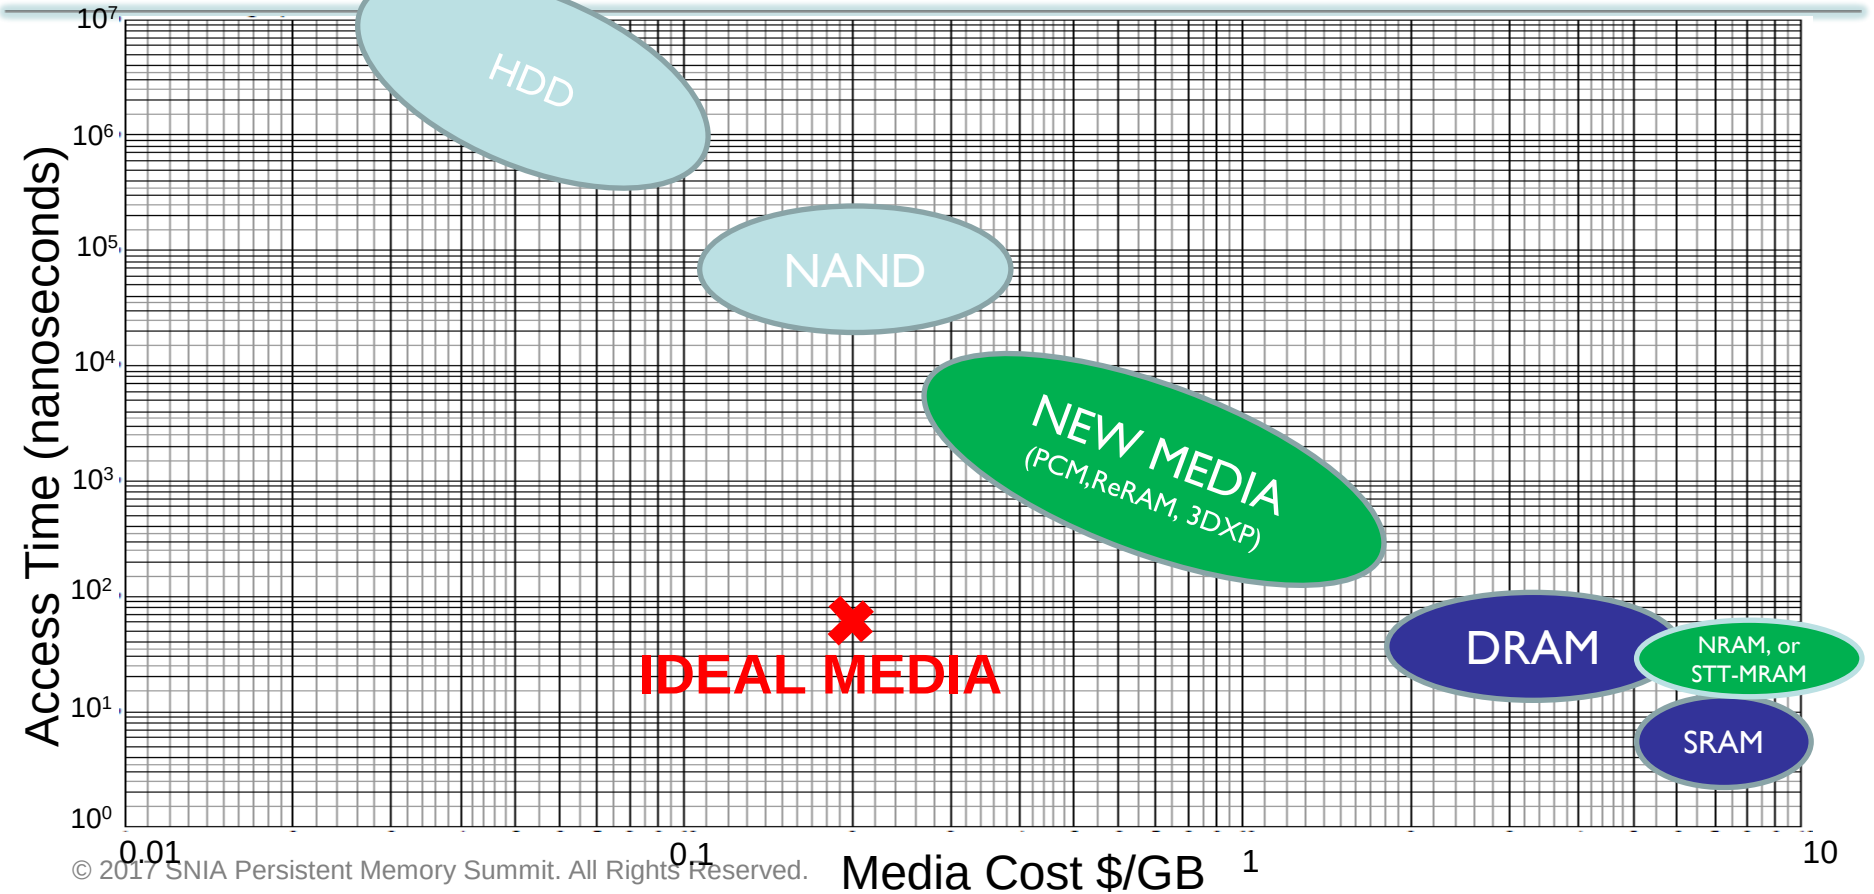

NVDIMM-P can make current DRAM/NAND Competitive with more Exotic New Media

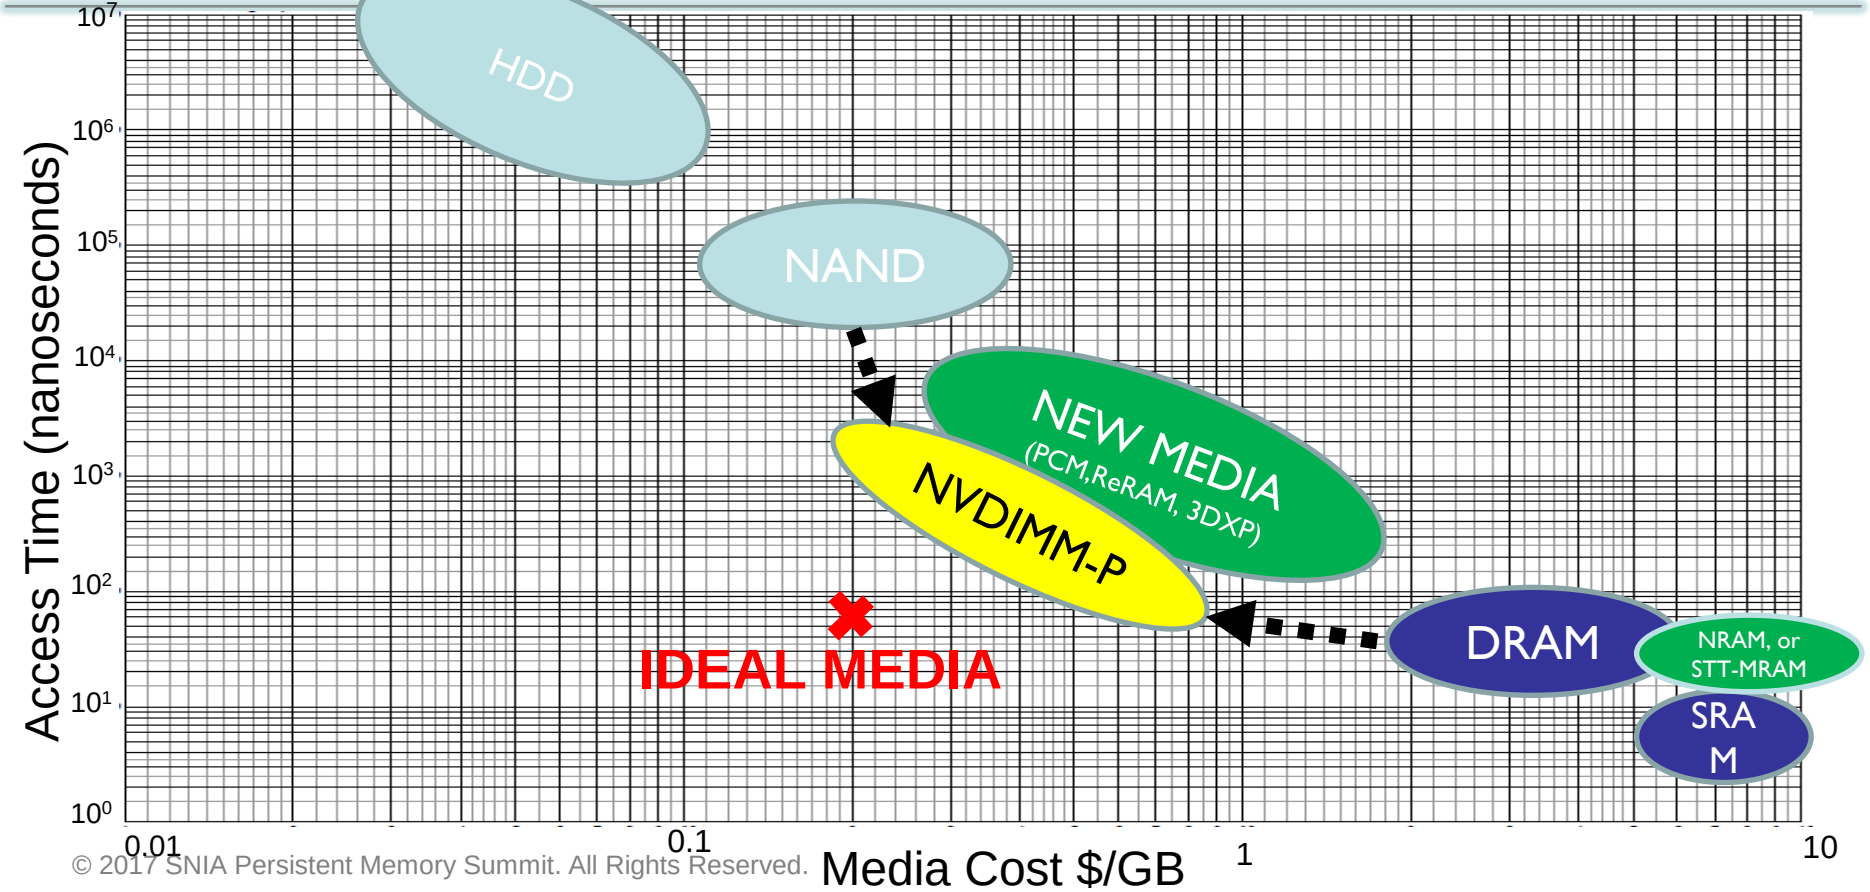

Key Elements of a Cache-Based NVDIMM Architecture

- NVDIMM-N is a shadow memory based architecture
  - DRAM and NAND size are about the same and cover the same address space
- NVDIMM-P is a protocol to allow non-deterministic accesses to enable a cache-based NVDIMM architecture
  - DRAM size is a fraction of the NAND size; **Significant cost advantages**
- Tiered, heterogeneous memory on the same DIMM
  - Both fast and slow memory connected in a cache architecture
- Private on-DIMM memory bus between the fast and slow memories
  - Doesn't tie up the main memory bus for line fetches and cast-outs
- Implement a command queuing capability like HDDs
  - Helps hide latencies when there is a cache miss

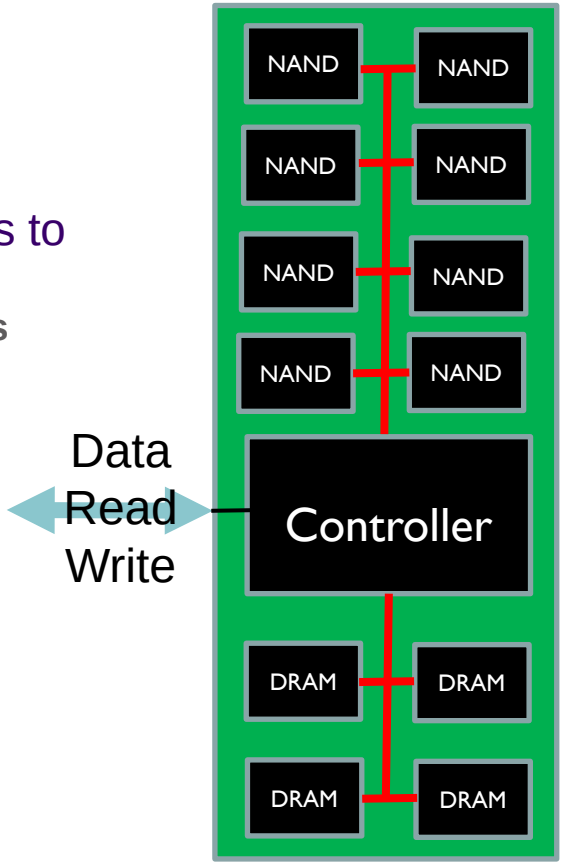

Challenge – Deterministic DRAM Bus

Potential Solutions

- Industry working on adding new control pins and associated protocol to a DDR4 extension and the standard DDR5 busses to handle variable access times

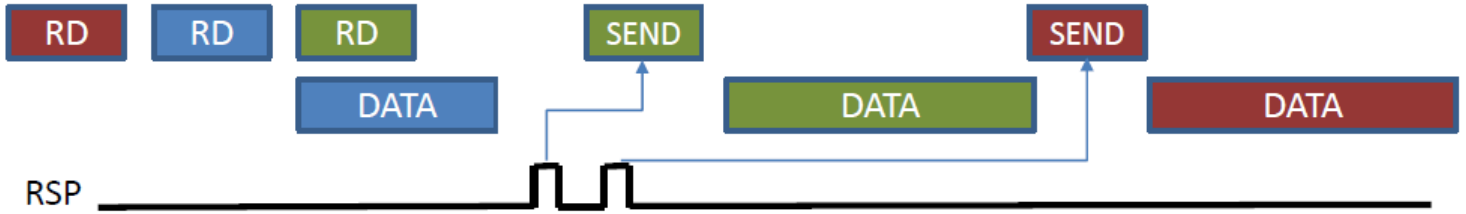

- Use a special driver and programming of the CPU page table
- Provide a backwards compatibility mode and driver to work on existing systems in a block mode

Implications for New Media

- Large scale adoption of new media technology will be slowed due to the onset of cache-based NVDIMMs
- Impact will raise the bar for the new technologies and they will require forcing further cost reductions
  - Published reports of costs/bit are roughly 50% of DRAM, but 500% that of NAND
  - Will need to achieve costs structure closer to NAND before going mainstream
- Once costs get closer to those of NAND the same cache-based NVDIMM concepts will be used by replacing the NAND/DRAM with the appropriate new technology

NVDIMM-P can make current DRAM/NAND Competitive with more Exotic New Media

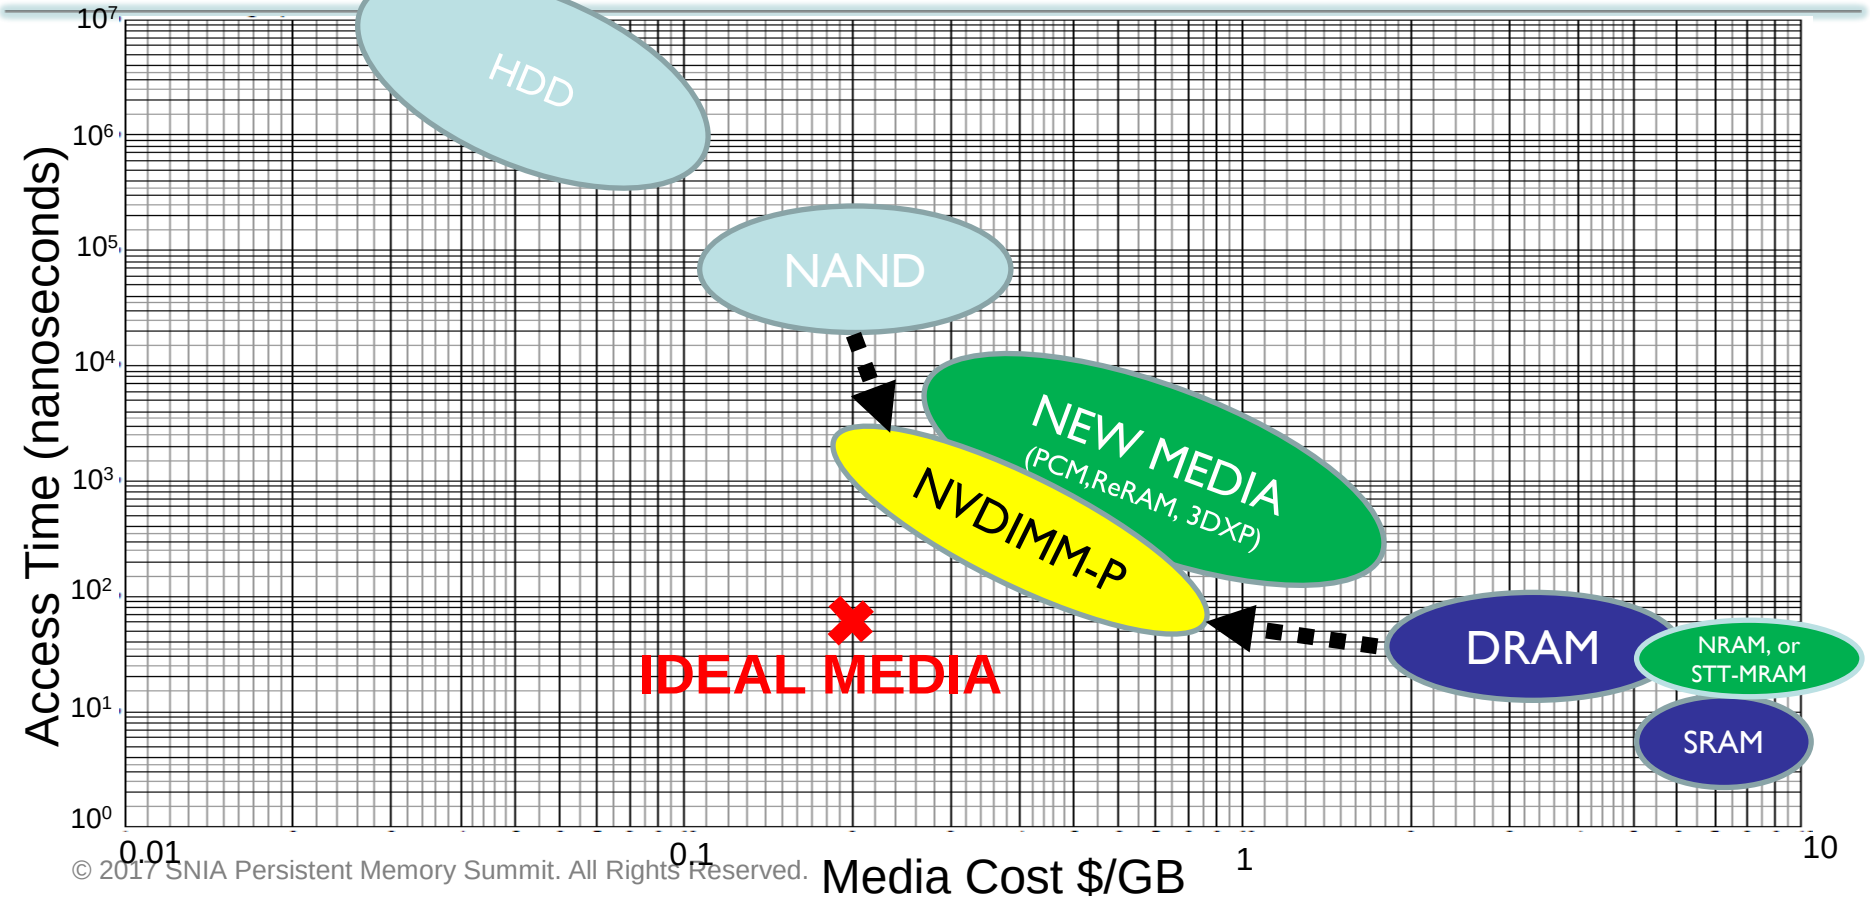

Use Same NVDIMM-P Architecture when the New Media Becomes More Cost Effective

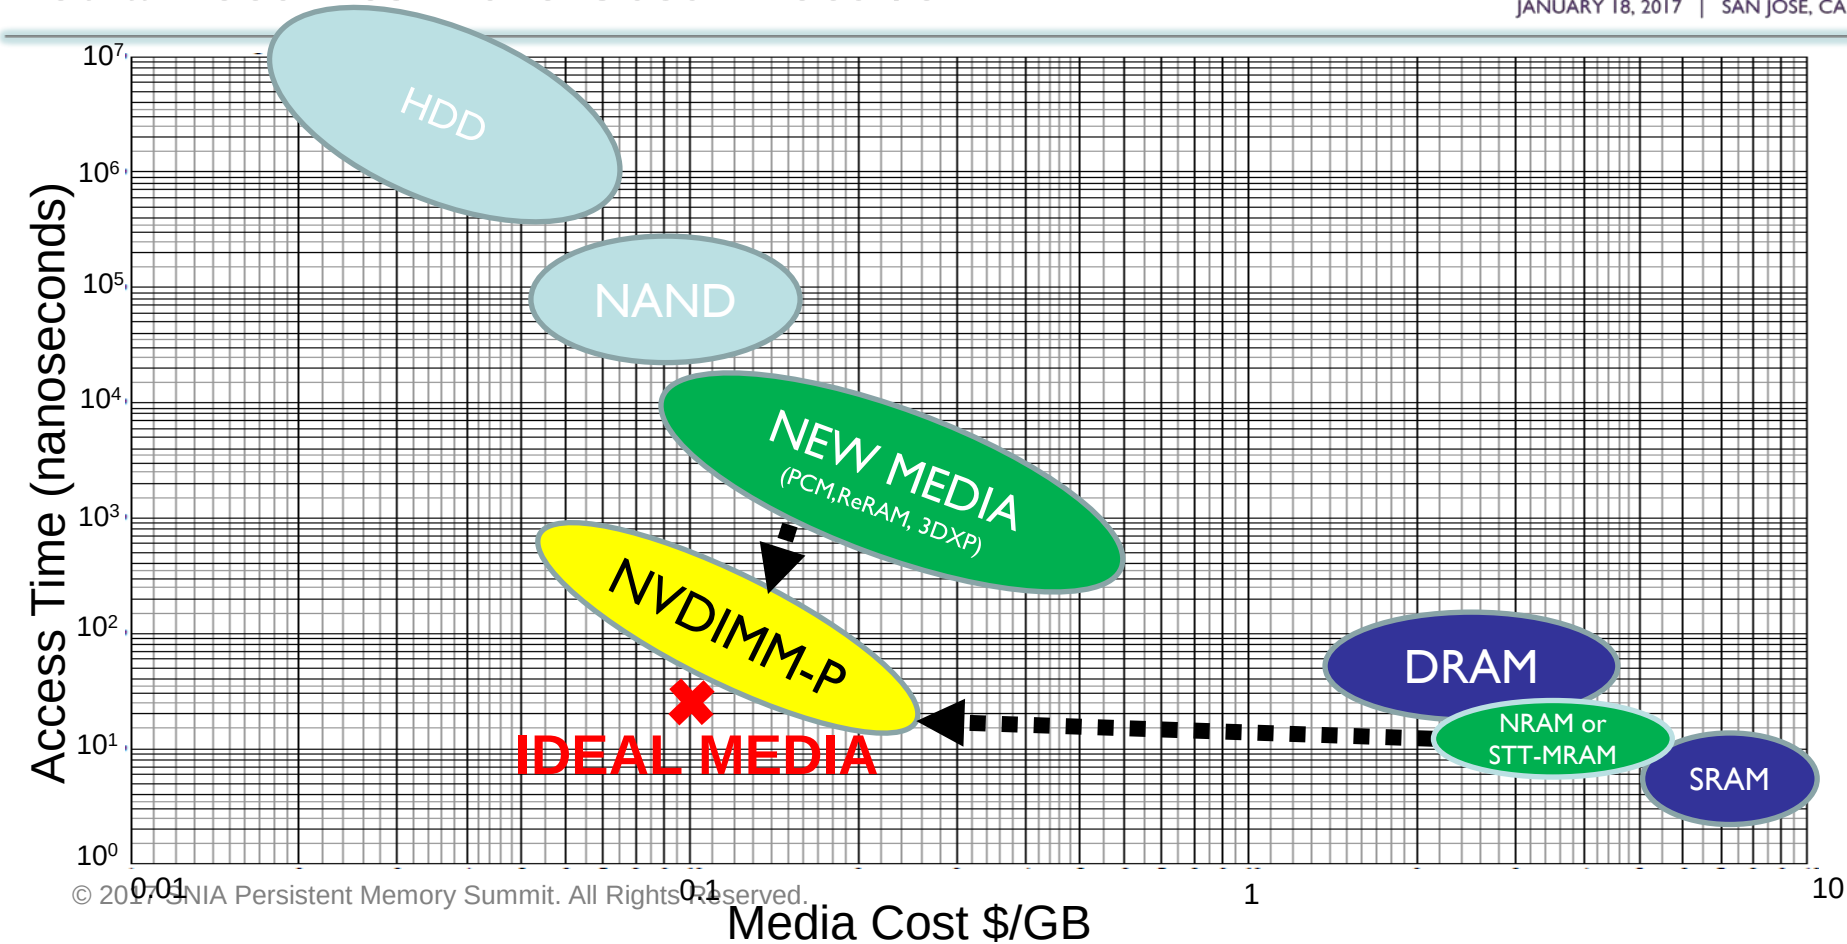

Conclusion – Persistent Memory in the 2020's

- Enterprise SSD interfaces will always migrate to the fastest bus available
  - NVMe → DDR4 → DDR5 → GenZ or OpenCapi
- Byte mode (or DAX) access will be heavily adopted and provide further performance improvements
- On-module cache-based architectures with fast and slow memories will become the standard
  - DRAM/NAND initially, with potential new media replacements later on
- Command queuing, not used with DRAMs today, will become a standard
  - Enables execution out-of-order when latencies are variable
  - Will help to hide latencies associated with cache misses
- **Cache-based NVDIMM architectures will be the predominant interface overtaking NVMe within the next 5-10 years in the race for performance**

NVDIMM **NVMe**

Panel Discussion

- Chair
  - Steffen Hellmold, Western Digital
- Panelists
  - Barry Hoberman, Spin Transfer Technologies
  - Doug Finke, Xitore