HEPiX Techwatch WG: Memory

Andrea Sciabà (CERN), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation), First Last (Affiliation)

"Money can buy bandwidth, but *latency is forever*", John R. Mashey (SGI), USENIX 99, June 1999.
"It's the Memory, Stupid!", Richard Sites, Microprocessor Report, Vol. 10, No. 10, Aug 1996.

Staging Area

Following are links to articles that cover topics that should be integrated into the document.

- SemiAnalysis overview of DRAM
- UltraRAM
- Emerging memory roadmap
- Multi-Rank Buffered DIMM (MRDIMM)
- MRDIMM
- Multiplexer Combined Ranks DIMM (MCR DIMM)
- SRAM Scaling
- More SRAM Scaling
- Even more SRAM scaling
- EUV Lithography and DRAM
- EUV at Micron
- DRAM Scaling
- Die Stacking yield implications
- GDDR Generation table
- RowPress
- RowHammer
- CRAM
- GDDR7
- 3DRAM
- HBM4
- DRAM Fabrication Challenges
- Per Core memory bandwidth scaling
- Flash Memory Summit

Topics in need of updating

- 2D, 2.5D, 3D interconnect and other packaging.techniques
- Discussion of power consumption, particularly with I/O to memory (motivation for HBM and GDDR memory)
- DDR5 and future DDR6
- GDDR7
    - PAM 3 signaling
- LPDDR5
- MR DIMM/MCR DIMM
- HBM3e/HBM4
- Impact of Unified Memory and connectivity (NVLink/Infinity Fabric/CXL)
- Emerging memory ? E.g. CRAM
- Persistent memory, is it dead ?

Suggested Revised Outline

Discussion of the future of DRAM and SRAM memory technology and changes to the packaging and placement, both physical and logical, of these memory components in the memory/storage hierarchy in compute systems. DIscussion of applicability of these new memory components to HTC and HPC computing.

1. Executive Summary
2. Burning Questions
3. SRAM
    a. On die SRAM scaling
    b. Alternatives
4. DRAM
    a. DRAM Technology
    b. Packaging Technology/Standards
        i. HBM
        ii. GDDR/LPDDR
        iii. DDR
        iv. Disaggregated (CXL.mem)
    c. General competitive environment
        i. Samsung
        ii. SK Hynix
        iii. Micron
    d. Security
        i. RowHammer
        ii. RowPress

Executive Summary

Semiconductor memory, specifically volatile memory, feeds the data and instructions needed to keep CPUs busy. Keeping the CPU fed has become increasingly difficult as core counts have increased, semiconductor die sizes have decreased, per “pin” I/O baud rates have stalled and restrictions in motherboard real estate and power consumption have become an issue. This report covers the current state of memory technology used in the design of computing systems and developments in memory on the horizon to keep up with CPU performance that may impact HEP/NP computing.

Burning Questions

1. What are the key changes to watch in the memory space ?
    a. A wide range of memory system configurations using DRAM and SRAM are being deployed in “commodity” computing systems, driven by increasing demands of inference and training of AI models for data, bandwidth and latency.
        i. CPU configurations with expanded L3 cache, made possible through SRAM chiplets stacked on processor core dies, are now available (e.g. AMD EPYC 3DV Cache) on a few CPU models.
        ii. Substantially higher bandwidth DRAM subsystems are being deployed by CPU and system vendors, but they come with some limitations. These memory subsystems may be configured by designers as cache, main memory, or a high performance partition of main memory. Examples are:
            1. On-package HBM memory (e.g., Intel Xeon Max processor, Fujitsu A64FX), although no CPUs in the latest generation have it
            2. Co-packaged LPDDR5X memory (e.g. Nvidia Grace processor) to augment the system DDR5 DRAM
    b. “Unified” memory systems, where multiple devices have coherent access to a common pool of memory. The typical application is for CPU-GPU systems.
        i. Sharing of memory can be accomplished with CXL, NVLink, and Infinity Fabric between the CPU and accelerators
    c. Higher bandwidth and higher capacity memory modules
        i. 3DS DDR4 DIMM - Allows higher memory capacity in a standard DIMM package through the use of “stacked” DRAM die.
        ii. Multiplexer Combined Ranks DIMM (MCR DIMM)
        iii. Multi-Ranked Buffered DIMM (MR DIMM)
2. Is traditional off package memory dead, i.e. DIMM based memory ?

a. Not anytime soon. Other types of packages address very specific use cases (e.g. very large bandwidth or capacities).

Introduction

(Updated 9/24) Memory is critical to the operation of any computing or storage system. The code being executed by the CPU, the data being manipulated, and the state of the system are all stored in memory. As computational logic and memory technology has evolved over the past two decades, the gap between CPU performance and memory performance has grown. This document covers the current state of memory devices and technologies used in building memory systems to mitigate the deficiencies of the available memory technologies.

Memory Latency

Over the years, CPU performance has increased at a substantially faster rate than memory performance, as schematically shown by the following graph.

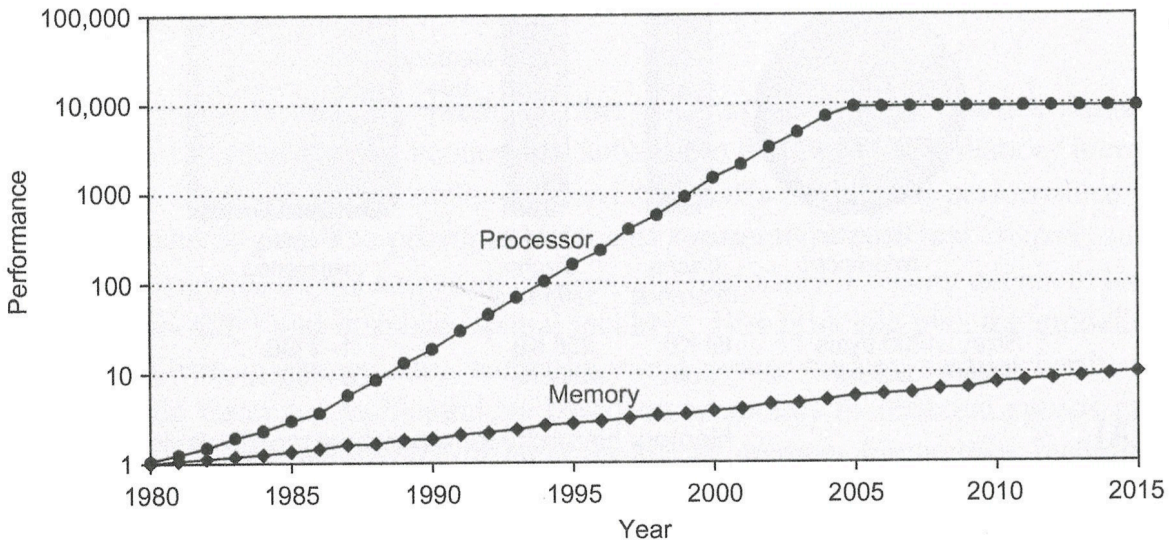

Qualitative graph of the gap between CPU and memory performance from Hennessy and Patterson, Computer Architecture: A Quantitative Approach.

The impact of this growing gap in performance between memory and processor has been referred to as "hitting the memory wall". In order to mitigate the effects of hitting the wall, a computing system (CPU + memory subsystem) needs to be designed to compensate for this performance gap, or the application (software) needs to be re-written with algorithms and coding

practices that access memory in a more efficient manner. The former approach is used by general purpose CPUs and their associated multi-level memory subsystems. The latter approach is used by software that is tailored for GPUs and AI/ML processors. For both paths, minimizing memory system latency and maximizing bandwidth are paramount for achieving maximum performance.

Memory Technologies and Hierarchies

Modern CPUs and computing systems utilize a hierarchy of memory subsystems with different characteristics to feed the CPU with data at the lowest possible latency and at the highest possible bandwidth. These memory subsystems are composed of memory devices that differ in the following characteristics:

1. Access latency and bandwidth
2. Storage capacity
3. Power consumption (Watt/GB)
4. Density (GB/nm^2, GB/package)
5. Cost ($/GB)
6. Addressability (block vs byte)
7. Persistence (volatile vs non-volatile)

The first two parameters are important when considering the performance of the memory subsystem being designed. The third, fourth and fifth parameters are important when considering the economic viability of the memory subsystem. Collectively, these parameters influence the performance (bandwidth and latency) and capacity (GB) of memory that can be deployed in a given physical or logical location in the memory hierarchy.

The following table shows the relative tradeoff between access bandwidth and latency for the different types of memory in a typical modern CPU. Note the increase in latency and decreased bandwidth when accessing memory that is physically more distant from the CPU.

Memory type	Bandwidth	Latency

In the next sections we will cover the memory devices, static (SRAM) and dynamic (DRAM) RAM, and memory packaging technologies, (e.g., HBM, DDR) that are used to build memory

subsystems in current computing systems. In addition, emerging memory technologies, both devices and packaging are also covered.

Static RAM

Static random access memory or (SRAM) technology has been the memory technology of choice for CPU registers and on-chip cache due to its speed. SRAM is substantially faster than DRAM and does not require refreshing like DRAM. SRAM utilizes a six transistor (or four transistor) bi-stable circuit to hold data (0 or 1) as long as power is applied. In contrast to DRAM, which uses capacitors to store data in the form of an electrical charge, SRAM is built only from transistors, making them more compatible with semiconductor “logic” processes used to build the rest of the CPU. Historically, this has led to the integration of SRAM cells on the same semiconductor die on which the CPU cores reside.

With the advancements in process technology, SRAM has become more problematic for use on chip. The area scaling behavior of SRAM in newer FinFET processes is weaker than with older planer transistor processes, resulting in smaller SRAM area reductions for a given reduction in process feature size (See the graph below comparing scaling of SRAM and logic as transistor feature size shrinks). Controlling SRAM power consumption has also become harder with newer processes. To exacerbate the power consumption and area problem, emerging applications like machine learning are driving the need to have more on-chip memory. On-chip SRAM, in the form of L1, L2, and L3 cache, consumes a substantial fraction of area on a modern CPU die and a significant fraction of the power.

The following graph shows SRAM memory density for 16nm and 28 nm processes as a function of memory frequency. A graph of SRAM memory density for recent processes from Intel, Samsung, and TSMC can be found at WikiChips.

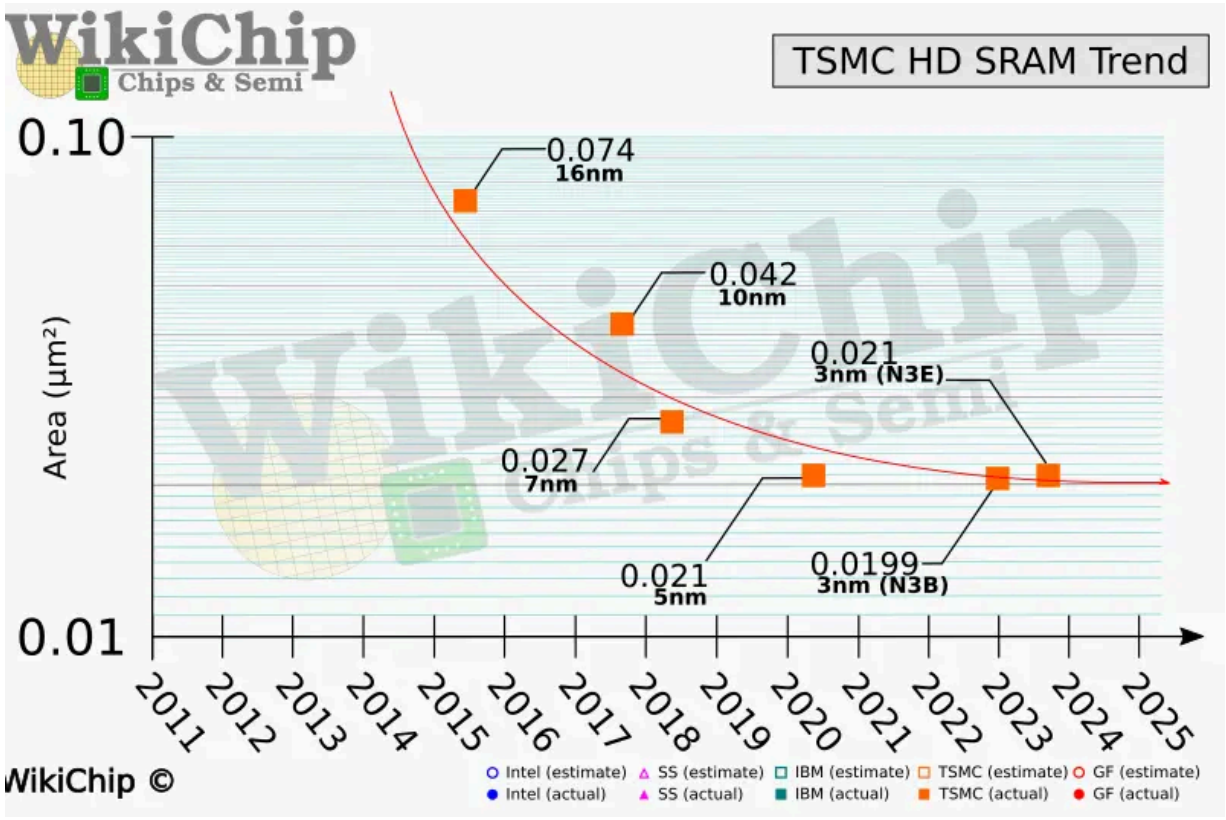

"IEDM 2022: Did We Just Witness The Death Of SRAM?", WikiChip.com

Mitigation Strategies

As SRAMs are mostly embedded in CPU die, advances in SRAM are in the hands of the semiconductor foundries, namely Intel, TSMC, and Samsung. The switch from FinFETs to GAAFETs, aka ribbon FET or nanosheet FET, in newer semiconductor processes is one technology that will help keep on die SRAM scaling as logic processes advance. An alternative is to move SRAM off the processor (core) die.

SRAM Chiplets

Recently, AMD is using SRAM chiplets stacked on top of their core chiplets with their 3D V-Cache technology. With a dedicated SRAM chiplet, AMD can leverage semiconductor processes that are better suited for memory, instead of the more expensive, logic optimized processes used for the core chiplets. This also allows for more SRAM and CPU cores in a physical package compared to a monolithic CPU/SRAM die. There is a marginal increase in latency with this approach, but the manufacturing process is more complex. However, 3D

V-Cache is available on only a handful of AMD SKUs, like the Genoa-X family, and at this time it is not foreseen for the upcoming AMD Zen 5 server CPUs.

Alternatives to SRAM

eDRAM (Embedded DRAM)

eDRAM or embedded DRAM is DRAM memory that is fabricated on the same silicon die as the CPU. eDRAM has been championed by IBM and is used in IBM's Power 7, Power 8 and Power 9 processors to implement the L3 cache. According to IBM, the benefits of eDRAM are 3x higher bit density (bits/nm^2) and substantially lower power consumption than SRAM, but with similar performance characteristics. According to IBM, adding eDRAM requires minimal changes to the manufacturing process for the CPU. With eDRAM, a substantially larger on-die L3 cache can be implemented on a CPU at the same die size as one with a smaller SRAM L3 cache, or an equivalent L3 cache can be implemented on a CPU with a smaller overall die size. The former has the potential to increase CPU performance, the latter has the potential for lower cost (due to more die/wafer and higher yield).

Use of eDRAM has not proliferated in CPU designs, despite the advantages professed by IBM and no current CPUs adopt it. Likely reasons are be high cost or a process specific implementation. In addition to IBM, Intel tried eDRAM with their Haswell processor, albeit with a separate eDRAM die connected to the CPU in package

There is considerable research being conducted to find alternatives to SRAM, although all of them are still way too expensive to be viable. These alternatives are covered in the Emerging Memory Technologies section of this document.

[1] "The Trouble with SRAM", EETimes, Dec 17, 2018
[2] "The Race is On", EETimes, Jan 6, 2019
[3] "SRAM Scaling Issues, And What Comes Next", Semiconductor Engineering, Feb 15, 2024

Dynamic RAM

Overview

Dynamic random-access memory or DRAM has been the memory technology of choice for the bulk of the memory in computing systems. Although slower than SRAM, DRAM is significantly denser (higher bits/area) and cheaper ($/GB). Through the years, DRAM capacity (bit density) and bandwidth have increased, while cost has decreased, even if it is subject to very large variations. However, memory latency has effectively remained constant.

The memory bandwidth per core was approximately constant in the last decade, as it can be seen by the following plot.

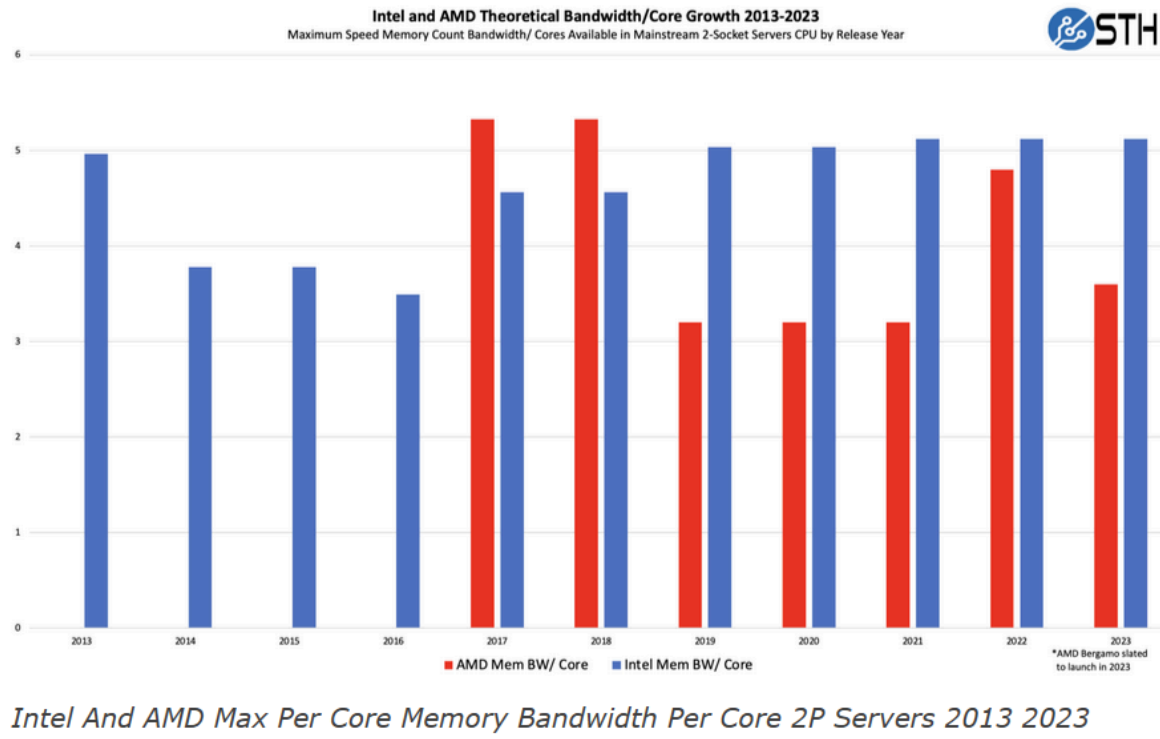

Memory bandwidth per core for Intel and AMD vs. time (ServeTheHome)

The following table from Micron Technology shows memory latency, the “true latency” column, for the different DRAM technologies that have been available over the years from oldest, SDR (Single Data Rate), to the newest DDR4 (Dual Data Rate Gen 4)

SPEED VS. LATENCY AS MEMORY TECHNOLOGY HAS MATURED
(INDUSTRY STANDARDS)

TECHNOLOGY	MODULE SPEED (MT/s)	CLOCK CYCLE TIME (ns)	CAS LATENCY (CL)	TRUE LATENCY (ns)
SDR	100	8.00	3	24.00
SDR	133	7.50	3	22.50
DDR	335	6.00	2.5	15.00
DDR	400	5.00	3	15.00
DDR2	667	3.00	5	15.00
DDR2	800	2.50	6	15.00
DDR3	1333	1.50	9	13.50
DDR3	1600	1.25	11	13.75
DDR4	1866	1.07	13	13.93
DDR4	2133	0.94	15	14.06
DDR4	2400	0.83	17	14.17
DDR4	2666	0.75	18	13.50

DRAM memory latency for several generations of DDRx memory. Micron Technology

Moving forward, manufacturers of DRAM chips continue to reduce cost/GB and increase chip density and speed.

Discussion on the current state of the art for DRAM implementations, challenges moving forward including:

1. Reduced stored charge per bit (reduction in cell capacitance), impacting bit readout noise margins and charge retention times
2. ~~Difficulties in manufacturing~~
3. ~~Cell to cell crosstalk~~
4. https://blocksandfiles.com/2020/04/13/dram-is-stuck-in-a-10nm-process-trap/

Basic DRAM Technology

DRAM is arranged as a rectangular array of charge storage cells consisting of a capacitor and a transistor per bit, with the cells needing to be refreshed periodically. While the original DRAM was of the asynchronous type, it was eventually replaced from 1997 on by synchronous DRAM (SDRAM) due to its support for multiple internal banks that can be read in parallel.

The DRAM cell area is given as nF^2, where “F” is a distance measurement roughly corresponding to the minimum feature size in the manufacturing process and “*n*” is a number derived from the cell design.

In order to increase bit density, DRAM manufacturers adopt new process technologies to shrink cells and, when possible, adopt new cell designs. While all current DRAM uses $6F^2$ cells and FinFET transistors, there are ways forward to considerably increase the density in the next years, including adopting a $4F^2$ design, sub-10nm processes, three-dimensional transistors like

vertical channel transistors (VCTs), EUV lithography and eventually stacked DRAM (Tom's Hardware).

The $4F^2$ cells utilize new devices (Vertical Channel Access Transistor or VCAT) and different cell geometries to reduce the cell size. A $4F^2$ cell yields 33% more bits for a given area compared to a $6F^2$ cell. This allows for increased bit density with existing lithography equipment, delaying the need to transition to EUV lithography.

The following diagram from Tech Insights DRAM Technology Roadmap report shows the expected technologies in DRAM chips in current and future DRAM chips.

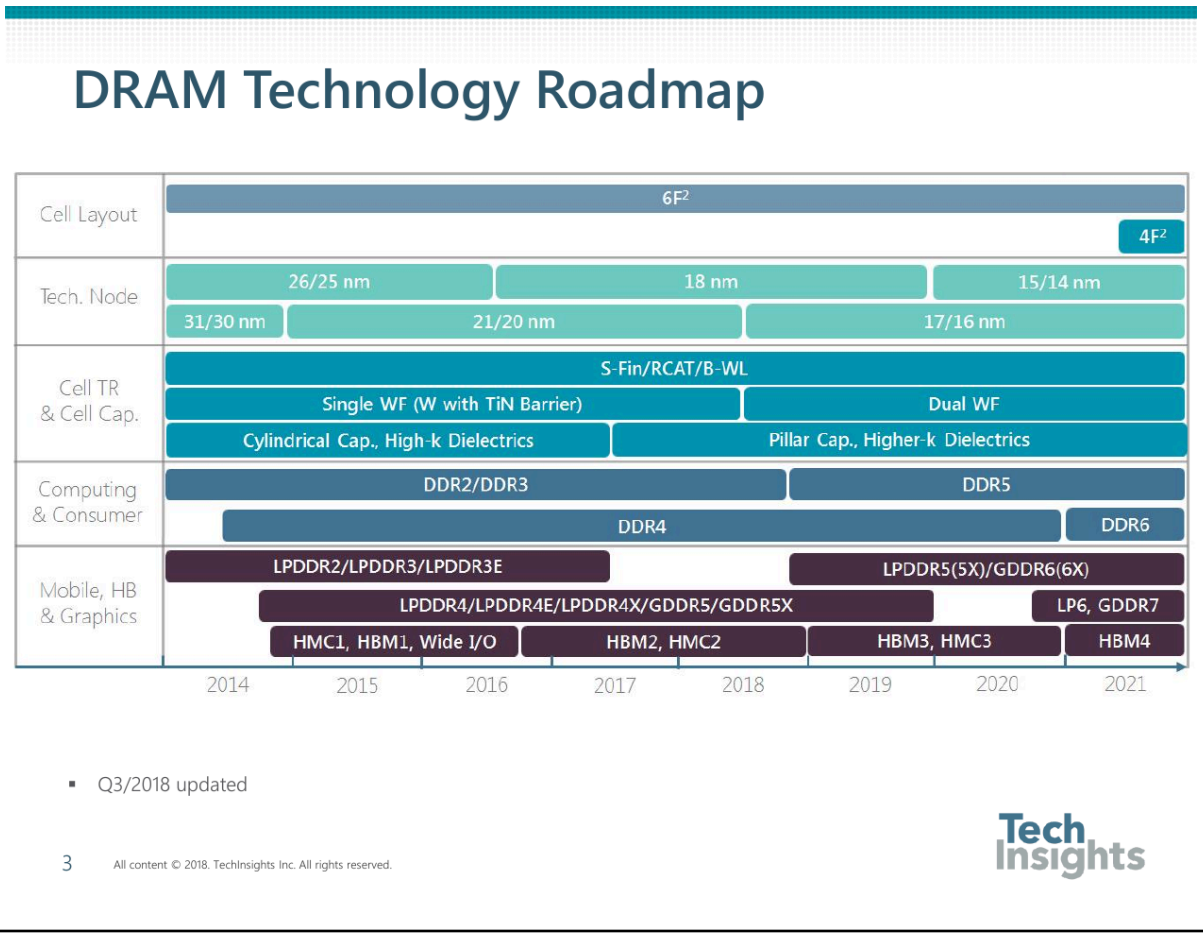

DRAM technology roadmap from TechInsights

The development of DRAM fabrication faces several challenges in further shrinking the structure sizes, which explains the lag in DRAM nodes with respect to logic nodes. Below a certain threshold in feature sizes, electrical resistivity increases nonlinearly, requiring new materials and new capacitor shapes to be evaluated. Similarly, crosstalk between cells becomes an issue, which can be mitigated by using new types of dielectrics. Lithography patterning and exposure

eventually will require either EUVL or self-aligned quadruple patterning (SAQP), with reducing tolerance for defects (Entegris).

Memory Packaging

The range of DRAM-based memory systems designs represent different tradeoffs in performance, cost, and capacity to satisfy specific needs. However, fundamentally they are all a response to the fact that communication circuitry has not matched logic scaling with regards to speed and power, specifically:

1. I/O energy consumption per bit has been roughly constant over time
2. I/O energy per bit increases with distance
3. I/O circuitry does not benefit from process feature size shrink
4. Size of I/O circuitry per lane remains roughly constant

The following table shows the change in energy (pico-joules) per logic operation vs energy per bit transmitted over various distances.

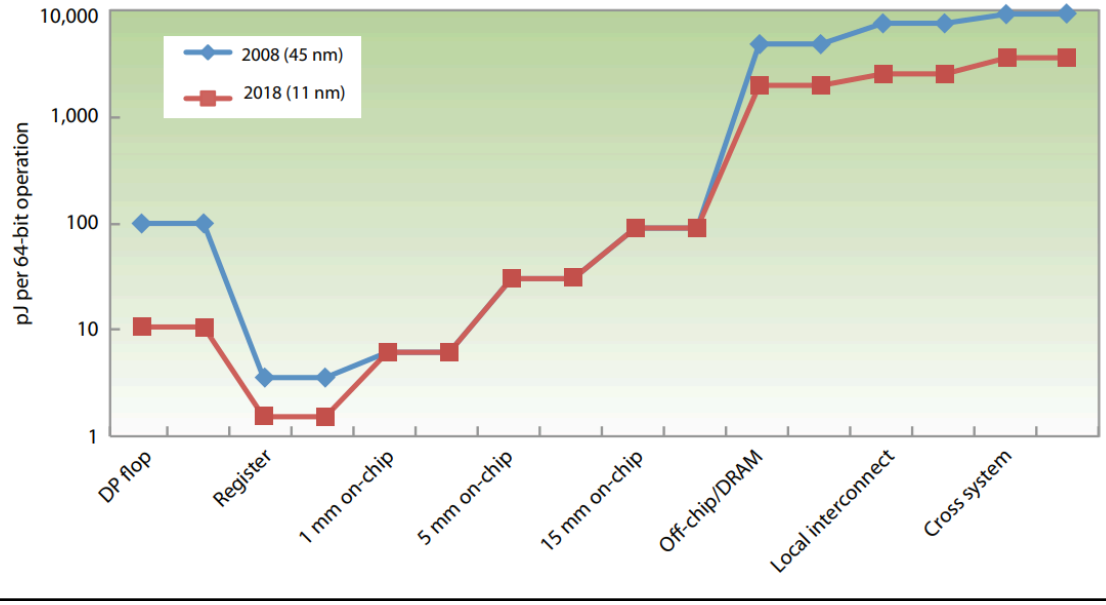

Exascale Computing Trends: Adjusting to the "New Normal" for Computer Architecture, IEEE Computing Society, Computing in Science & Engineering, Nov/Dec 2013.

Stacked Die memory (L3)

On Package memory

- HBM/GDDR/LPDDR memory

Off Package Memory

- DDR DIMM

- CXL/NVLink/Infinity Fabric memory

Disaggregated Memory

- CXL Fabric attached memory

HBM2e/HBM3 Technology

High Bandwidth Memory (HBM), as the name suggests, is a high bandwidth memory technology. The original HBM memory, initially standardized by JEDEC in 2013, utilizes up to eight vertically stacked 2Gb DRAM dies to create a compact memory module that can be placed in the same multi-chip module (MCM) as the CPU or GPU or even stacked on top. The following diagram from an AMD presentation on HBM schematically shows the packaging of HBM memory.

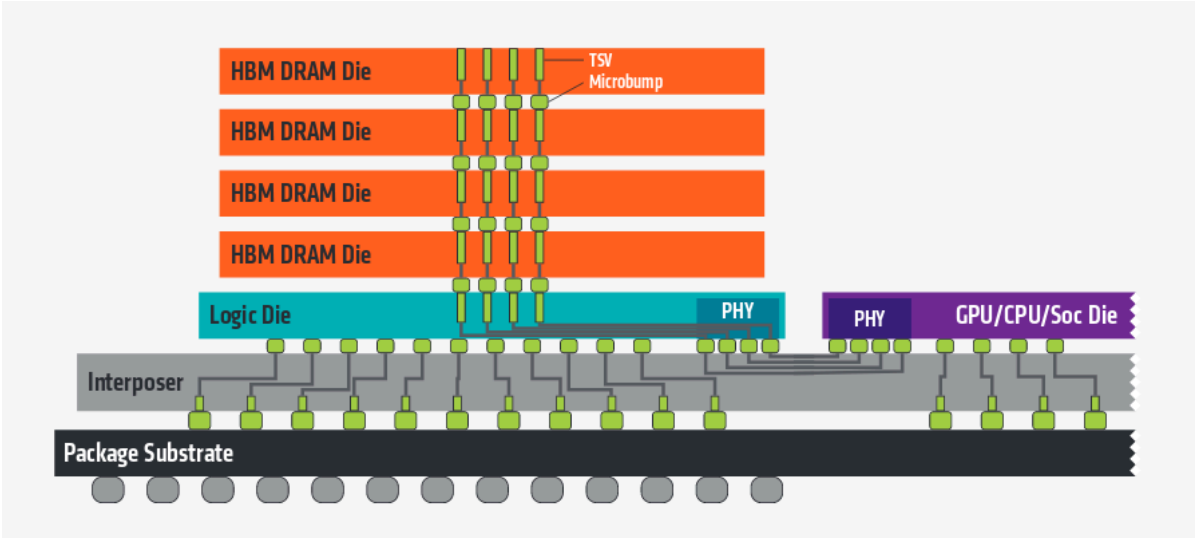

Schematic view of an HBM die stack from AMD

Although the “per pin” bit rate is low compared to GDDR5 (1 Gbps for HBM vs 10 Gbps for GDDR5), HBM achieves higher bandwidth by utilizing a 1024-bit wide interface per stack vs 32 bits per GDDR5 chip.

HBM2, the follow on to HBM, was published by JEDEC in 2016. HBM2 maintains the 1024 bit interface of HBM but upgrades the other parameters in the original HBM standard including:

1. Increasing per pin bit rate from 1Gbps in HBM to 2Gbps for HBM2
2. Increases DRAM die capacity from 2Gb to 8Gb
3. Increases the die stack from 4 die to 8 die.

Maximum bandwidth for HBM2 is 256GB/sec per stack compared to 128 GB/s for HBM. In December of 2018, JEDEC released an updated HBM2 specification that increased per pin bit rates to 2.4Gbps and die stack from 8 die to 12 die.

A comparison of bandwidth between DDR4, GDDR6, and HBM2 is shown in the following diagram from Rambus taken from an article in Semiconductor Engineering. The diagram shows the hardware required to achieve a memory bandwidth of 256GB/sec for each of the memory technologies.

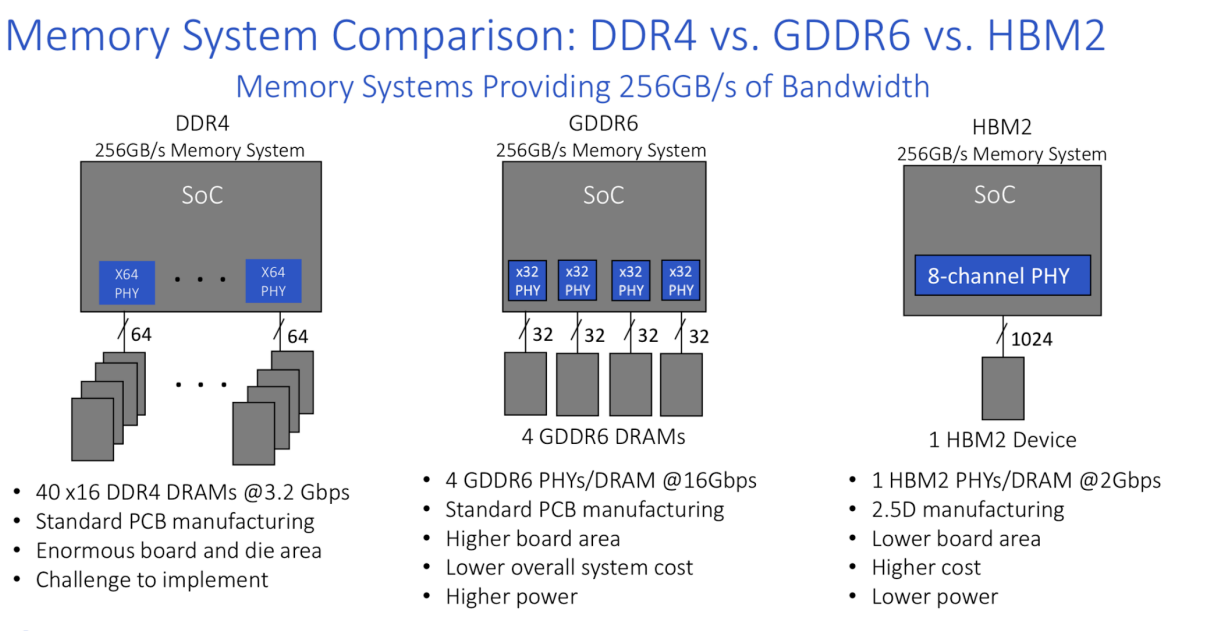

Comparison of different DRAM technology types. Semiconductor Engineering.

GDDR5/GDDR5X/GDDR6 Technology

Dual Data Rate Graphic DRAM version 5 or GDDR5 is DRAM technology that is tailored to high bandwidth environments, traditionally associated with graphics applications, and more recently with GPU compute environments. The GDDR5 standard specifies significantly higher per pin data rates for memory chips, up to 10 Gbps per pin. Maximum memory capacity per chip is 8Gb. Note that GDDR5 does not include a module level packaging standard like the DDR4 DIMM memory module, this is because GDDR5 chips are meant to be soldered to a PCB board. GDDR5X is a modification to GDDR5 that increases per pin data rates to 14 Gbps, and adds a 2nd quad data rate 16n prefetch mode, in addition to the standard 8n prefetch provided by

GDDR5. Note that GDDR5 chips are packaged in 170 pin BGA (Ball Grid Arrays) while GDDR5X is packaged in 190 pin BGAs.

The following diagram from a presentation by Jin Kim of Samsung at the 2016 Hot Chips symposium shows higher bandwidth of GDDR DRAM relative to standard DDR DRAM.

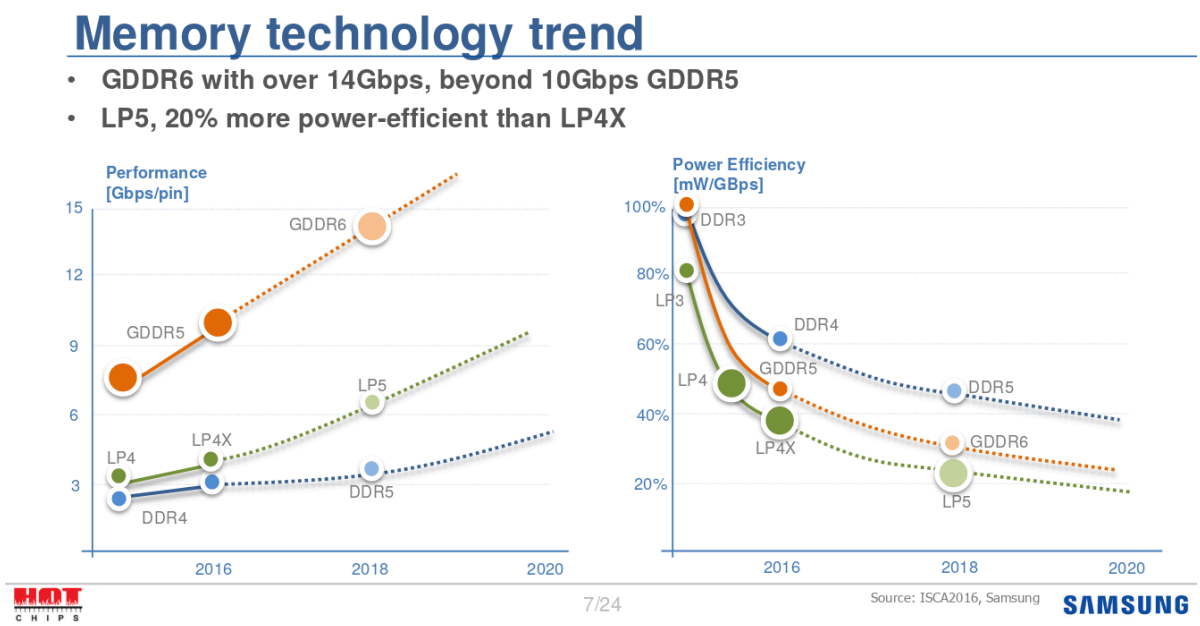

DRAM memory trends Jin Kim, Samsung, 2016 Hot Chips symposium.

Since the introduction of the GDDR5 standard in 2008, JEDEC has been working on its successor, GDDR6. This new JEDEC standard was finalized in 2017. Memory capacity chip remains unchanged from GDDR5/5X at 8Gb. GDDR6 chips have been announced by Samsung and Micron in 2018 with data rates of up to 16 Gbps per pin.

DDR4/3DS DDR4/DDR5 Technology

The main memory interface technologies used in servers are the various generations of the Dual Data Rate (DDR) defined by JEDEC.. Current CPUs use the DDR4 standard, while newer generation CPUs use either the DDR5. 3DS DDR4 (“3 Dimensional Stacked”) memory is a modification to the DDR4 standard that allows for higher capacity DIMMs, up to 512GB per DIMM for 3DS DDR4 compared to 64GB for DDR4.

Through various design changes, DDR5 allows for higher bandwidth and large capacity DIMMs than DDR5. DDR4 supports 1600, 1866, 2133, 2400, 2666, and 3200 MT/sec (mega-transfers/sec) at 8 bytes per transfer on DDR4 DIMMs. In contrast, DDR3 maxes out at 2133 MT/sec. In terms of capacity, DDR4 support 16Gb DRAM chips vs the 8Gb limit for DDR3.

JEDEC is currently working on the DDR5 standard. According to a JEDEC press release, "*JEDEC DDR5 memory will offer improved performance with greater power efficiency as compared to previous generation DRAM technologies. As planned, DDR5 will provide double the bandwidth and density over DDR4, along with delivering improved channel efficiency.*" DDR5 is expected to support data transfers from 3200 to 6400 MT/sec and 32Gb DRAM chips, effectively double the capabilities of DDR4. At this point in time DDR5 prototypes have been demonstrated by both Micron and SK Hynix. Both companies, along with Samsung have all announced DDR5 product availability in late CY2019.

The expected market life for the various versions of DDRx memory is shown in the following graph from a Flash Memory Summit 2018 presentation by Bill Gervasi from Nantero

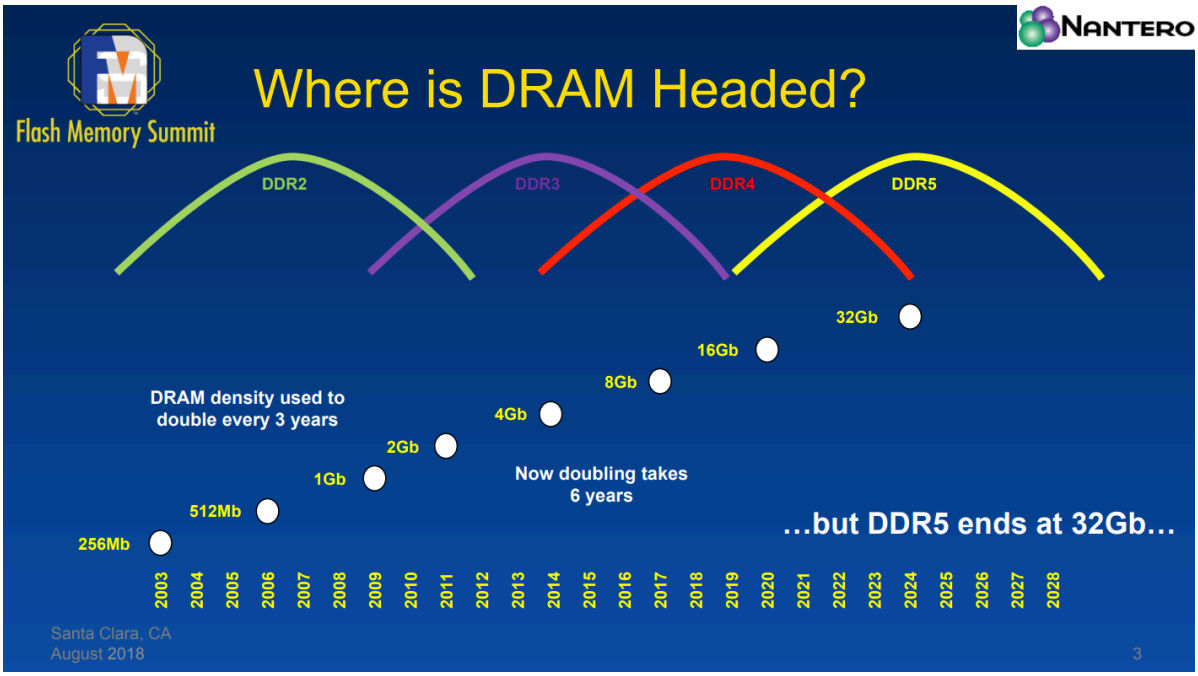

Market lifecycle for DDR generations. Bill Gervasi, Nantero, Flash Memory Summit 2018.

Security

Row Hammer

RowPress

Market status and outlook

Most recent DRAM market share information (CY2019) from DRAMeXchange can be found in the DRAMeXchange Weekly Report published on Aug 26, 2020

Table 1: Global Memory Module Revenue Ranking, 2021

Company	DRAM Ranking		DRAM Revenue (US$1 Million)			DRAM Market Share (%)	
	2021	2020	2021	2020	YoY %	2021	2020
Kingston Technology	1	1	14,261	13,200	8.0%	78.7%	78.0%
ADATA Technology	2	2	640	540	18.5%	3.5%	3.2%
Corsair	3	--	543	560	-3.1%	3.0%	3.3%
SMART Modular Technologies	4	6	459	401	14.6%	2.5%	2.4%
POWEV	5	5	442	406	8.9%	2.4%	2.4%
Kimtigo	6	4	438	493	-11.2%	2.4%	2.9%
Ramaxel	7	3	343	510	-32.7%	1.9%	3.0%
Team Group	8	7	168	174	-3.4%	0.9%	1.0%
Innodisk	9	10	164	100	64.7%	0.9%	0.6%
Apacer Technology	10	9	145	106	36.2%	0.8%	0.6%
Others			510	430	18.6%	2.8%	2.5%
Total DRAM Revenue			18,113	16,920	7.0%	100.0%	100.0%

Note 1: Due to the diversified nature of module manufacturers' business activities, above data based solely on individual DRAM revenue
Note 2: Above ranking is based on data provided by ranked brands
Source: TrendForce, Aug. 2022

(Updated 1/23) 2021 Memory Module Revenue by vendor (DRAMeXchange)

Q3 2022 revenues for memory manufacturers are as follows, also from DRAMeXchange

Table: Global Ranking of Branded DRAM Manufacturers by Revenue, 3Q22 (Unit: US$ Million)

Ranking	Company	Revenue			Market Share	
		3Q22	2Q22	QoQ	3Q22	2Q22
1	Samsung	7,400	11,130	-33.5%	40.7%	43.5%
2	SK hynix	5,242	7,011	-25.2%	28.8%	27.4%
3	Micron	4,809	6,271	-23.3%	26.4%	24.5%
4	Nanya	362	613	-40.8%	2.0%	2.4%
5	Winbond	150	239	-37.4%	0.8%	0.9%
6	PSMC	38	63	-40.0%	0.2%	0.2%
	Others	186	268	-30.5%	1.0%	1.0%
Total		18,187	25,594	-28.9%	100.0%	100.0%

Note 1: 2Q22 - - US$1:KRW$1,261; US$1:NT$29.4
Note 2: 3Q22 - - US$1:KRW$1,341; US$1:NT$30.4
Source: TrendForce, Nov., 2022

(Updated 1/23) 2022 Q3 DRAM Manufacturer Revenue by vendor (DRAMeXchange)

Table: Price Projections for Different Categories of DRAM Products, 4Q22~1Q23

	4Q22E	1Q23F
PC DRAM	DDR4: down 13~18% DDR5: down 18~23% Blended ASP: down 15~20%	DDR4: down 15~20% DDR5: down 18~23% Blended ASP: down 15~20%
Server DRAM	DDR4: down 23~28% DDR5: down 30~35% Blended ASP: down 23~28%	DDR4: down 15~20% DDR5: down 18~23% Blended ASP: down 15~20%
Mobile DRAM	down 18~23%	down 10~15%
Graphics DRAM	down 20~25%	down 18~23%
Consumer DRAM	down 18~23%	down 18~23%
Total DRAM	down 20~25% (DDR5 Blended)	down 13~18%

Source: TrendForce, Jan. 2023

(Updated 1/23) Projections for DRAM costs for 2022Q4 and 2023Q1 from TrendForce.

According to DRAMeXchange, DRAM prices are dropping dramatically, with contract prices dropping by over 25% and contract lengths shortening from quarterly to monthly. DRAMeXchange is blaming lower demand, resulting in significantly higher inventory. On the PC side, Intel's processor shortage is one cause for a drop in demand. The processor shortage is expected to last through the third quarter of 2019. DRAMeXchange expects that there will be additional pressure on DRAM prices as new DRAM production capacity is expected to be brought on line over the next two years by several DRAM vendors.

Historical DRAM market share information can also be found at Statistica on the Global DRAM chip vendors' market share 2011-2018 web page.

Spot price for DRAM can be found at DRAMeXchange, while price history is available by subscription only. PCPartPicker has a price history for selective DRAM types on their trends web page.

Market penetration for the different incarnations of DRAM, actual as well as estimated, is shown in the following graph from Cadence Design Systems:

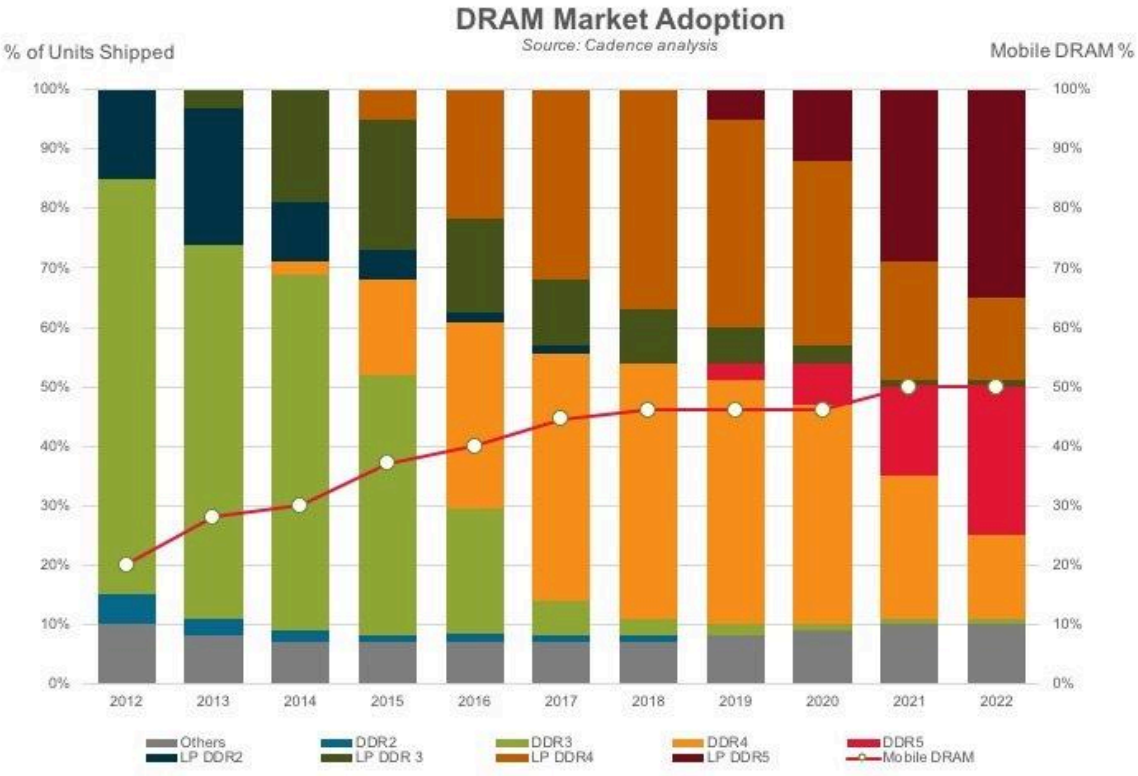

DRAM adoption by type. Cadence Design

Samsung

Samsung investor relations and link to presentations

SK Hynix

SK Hynix investor relations and research web sites

Micron

Micron investor relations

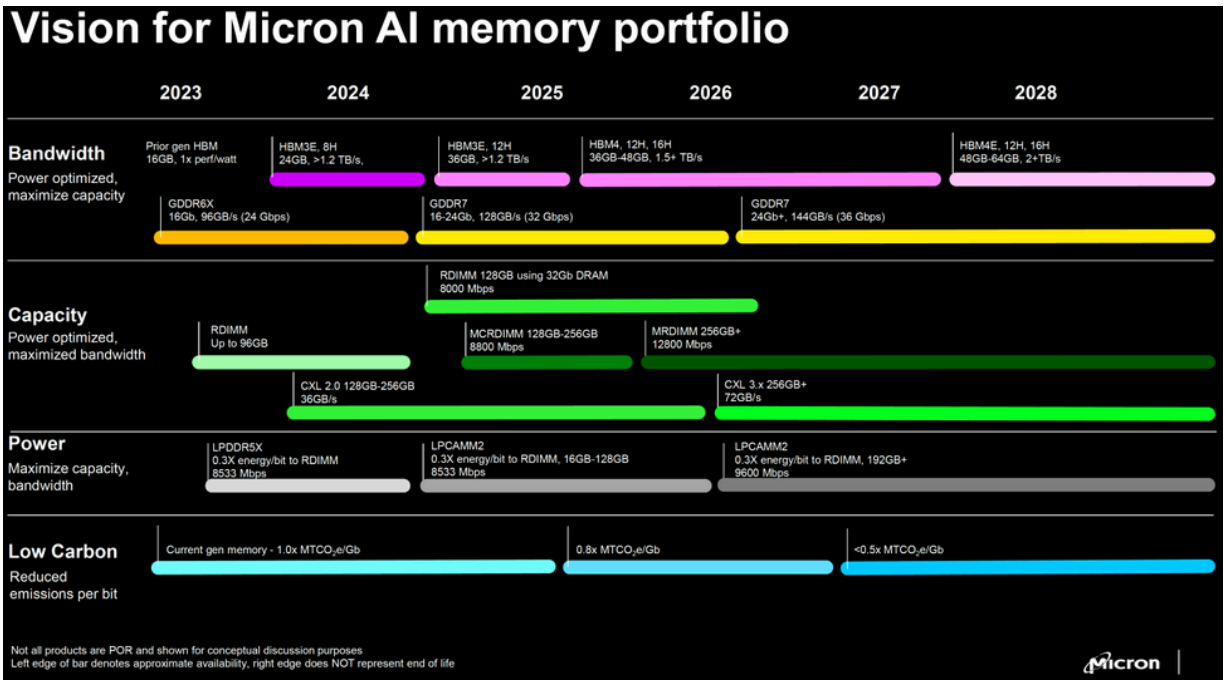

Micron DRAM roadmap from Tom's Hardware

Documentation

Specification for DRAM memory interfaces are managed by the Joint Electron Device Engineering Council (JEDEC). They are responsible for the various versions of the Double Data Rate (DDR) DRAM standards, with DDR4 being the most recent and DDR5 in development. They are also responsible for various versions of the High Bandwidth Memory (HBM) specification, with HBM2 being the most recent, as well as the Graphics Double Data Rate (GDDR) standards. The relevant JEDEC committee is JC-42

1. JEDEC DDR DRAM standards
2. JEDEC GDDR DRAM standards
3. JEDEC HBM standards

DRAMeXchange, part of TrendForce, is a market research company that focuses on the DRAM market. In addition to tracking DRAM spot pricing, DRAMeXchange generates reports on the DRAM market and also has pointers to other DRAM reports through their Weekly Research web page.

Memory Infrastructure

CXL

<ADD Content>

NVDIMM

<ADD Content>

Zoned Namespaces

<ADD Content>

Emerging Memory Technologies

Overview

Emerging memory technologies covers memory technologies that aren't SRAM, DRAM or NAND flash. These technologies may fit anywhere within then computer memory hierarchy, be it "close" to the CPU (low latency applications) to "far away" from the CPU, e.g., computer system main memory and is influenced by how the technology is packaged. This document is focused on the fundamental technology and not the packaging of the technology. For various reasons, most of the discussion of emerging memories is in the HEPiX Storage Working Group document. The primary reason for this is that, with the exception of Magnetoresistive RAM (MRAM) and Carbon Nanotube RAM (NRAM), the properties of these emerging memory technologies are more NAND flash-like, then DRAM-lie. Also, at this point in time, almost all applications for these emerging technologies are extensions of persistent, block storage.

From the discussion of SRAM and DRAM, it is clear that there is a problem with both technologies. SRAM suffers from power and memory density problems, going from to smaller and smaller geometries. This can be seen in the following density and leakage current graphs for SRAM from a Flash Memory Summit presentation by Les Crudele and Andrew J. Walker of Spin Memory.

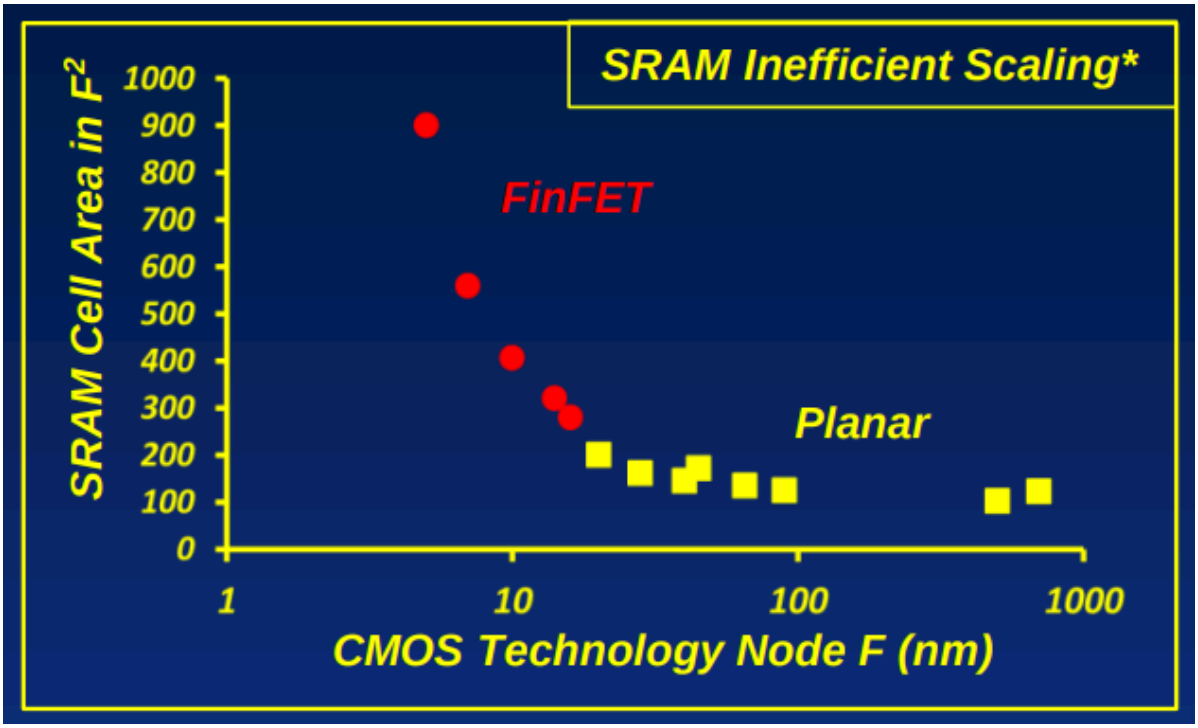

SRAM scaling. [Crudele and Walker, Spin Memory, Flash Memory Summit 2918].

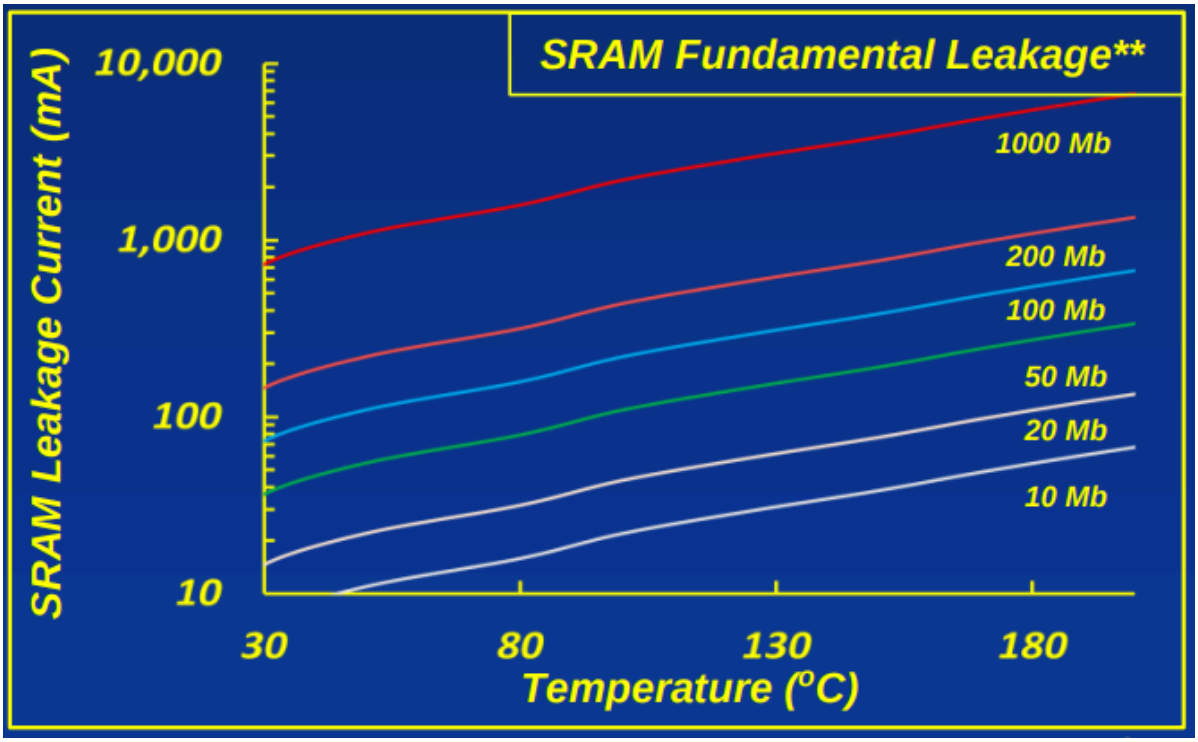

SRAM Leakage current as a function of temperature. Crudele and Walker

With DRAM, as mentioned previously, increases in performance have not kept up with CPU performance, particularly with latency. Access to DRAM, which resides off chip, also requires substantially more energy than accessing SRAM, which resides on chip, as shown in the following diagram from the same paper.

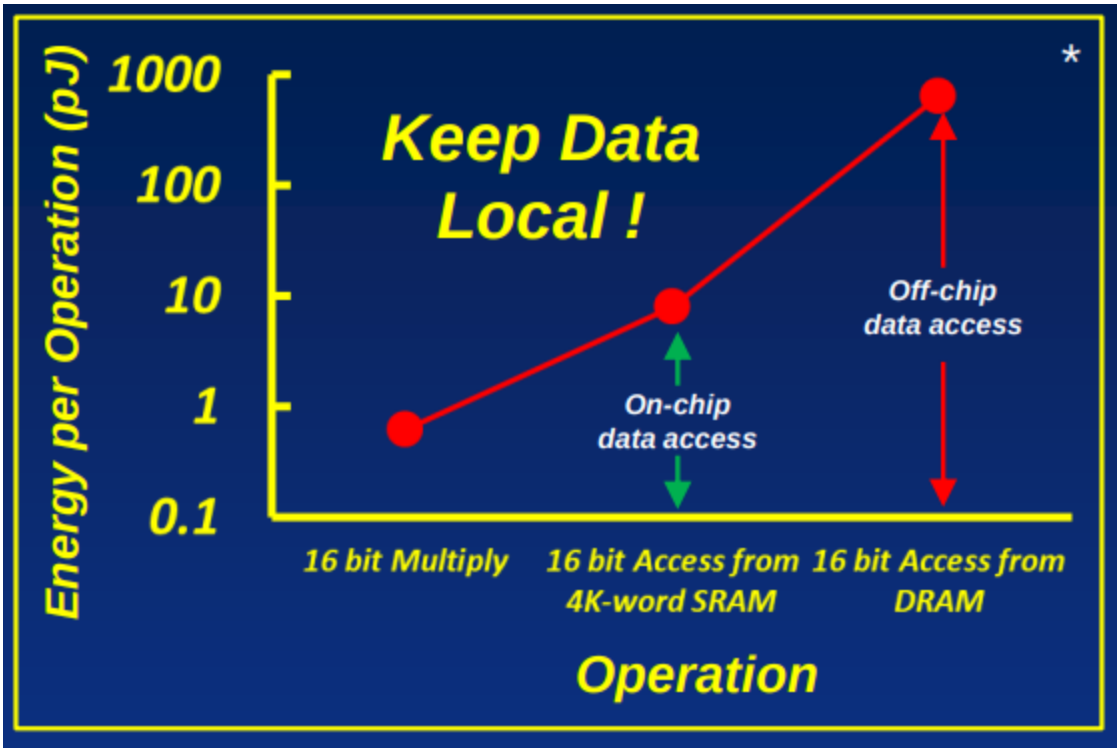

Energy consumed for on-chip SRAM and off-chip DRAM access relative to a logic operation. Crudele and Walker

Magnetoresistive RAM (MRAM)

Magnetoresistive RAM or MRAM, particularly Spin Transfer Torque MRAM (STT-RAM), is considered to be a viable replacement for SRAM and eDRAM for the last level cache (LLC) on a CPU. MRAM can also be configured to work as Persistent Memory, replacing DRAM or existing between DRAM and FLASH. The following slide from a MRAM presentation at the 2017 SNIA Persistent Memory Summit, by Andrew Walker of Spin Memory (formerly Spin Transfer Technologies), shows these two uses of MRAM in the memory hierarchy.

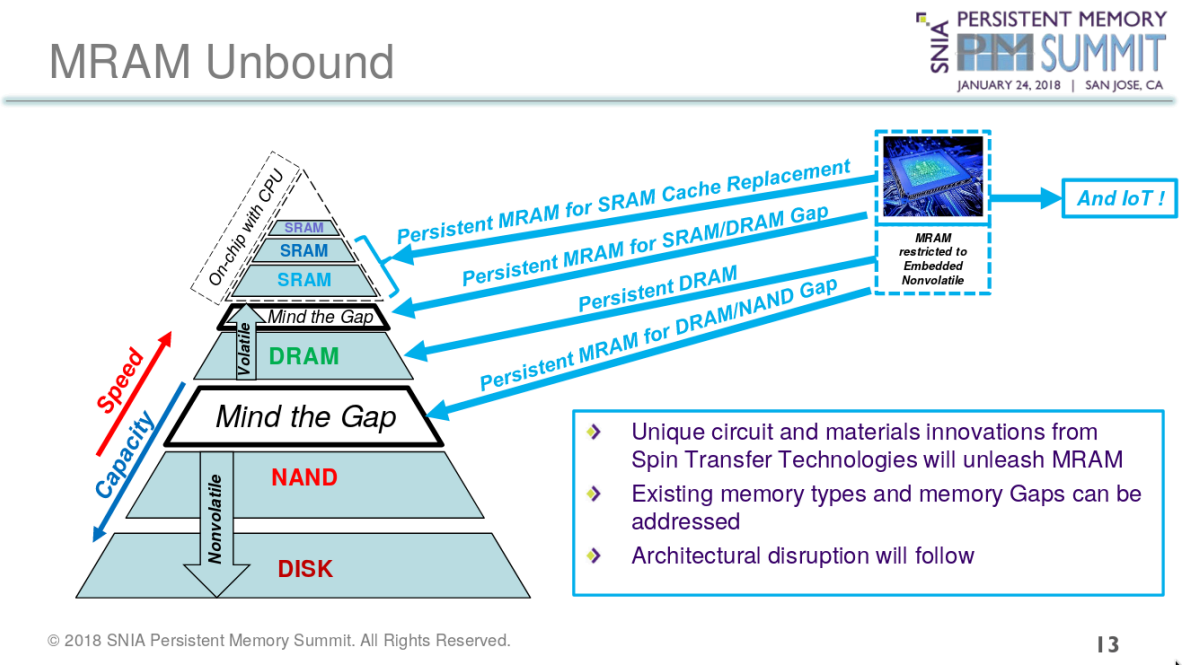

Positioning of MRAM memory in the memory hierarchy. Andrew Walker, Spin Memory, 2017 SNIA Persistent Memory Summit.

The following table from a presentation by Luk Thomas of TDK-Headway Technologies at the 33rd International Conference Massive Storage Systems and Technology 2017 shows the size of embedded MRAM (eMRAM) bit cells compared to eFlash, eDRAM, and SRAM. The table also shows the added processing steps (mask layers) required for each of the different memory technologies.

	eMRAM	eFlash	eDRAM	SRAM
Cell size (F^2)	30-50	30-50	30-90	100-300
Added mask layers	2-3	10-12	4-6	0

Cell size for various memory technologies. Luk Thomas TDK Headway Technologies, 33rd International Conference on Massive Storage Systems and TEchnologies 2017.

The presentation by Crudele and Walker shows a similar difference in cell size between eMRAM and SRAM, as shown below.

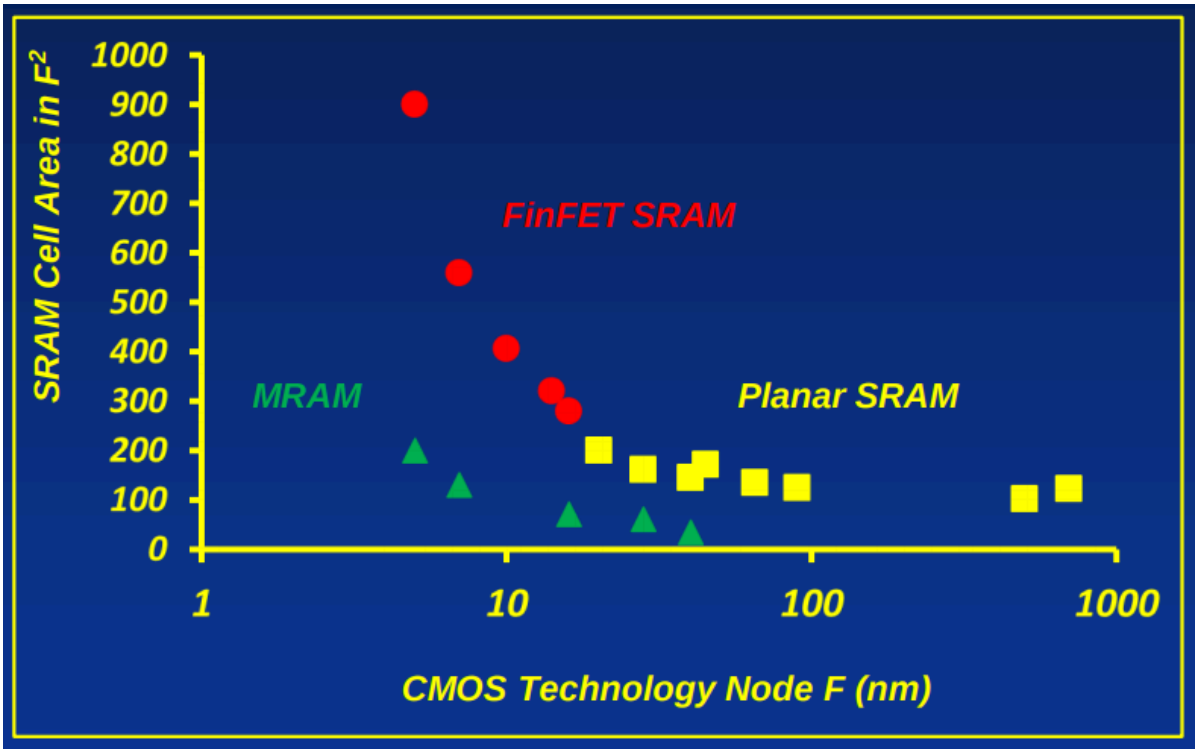

Scaling of MRAM cells relative to conventional SRAM. Crudele and Walker Flash Memory Summit

Carbon Nanotube RAM (NRAM)

Information about NRAM can be found in the presentation by Bill Gervasi of Nantero SNIA 2018 Persistent Memory Summit.

Market status and outlook

Target market for STT-MRAM and variants from Everspin Technologies is available in the Everspin 2019 and 2018 investor slides.

- % market, units/PB shipped, annual growth rate
- pricing evolution, (discount rates)

- vendors and their status and market share, sales, net revenues
- impact on HEP

Documentation

- The Memory Guy is a web site run by an industry analyst that covers memory technology. For those with money, a quick way to get an in depth view of the memory industry can be obtained by purchasing the analyst's memory technology reports.
- The Memory Guy's market research company is Objective Analysis. Although you need to pay for their in depth research reports, the outlines for the reports are freely available and provide an high level view of the important areas memory technology
- https://blocksandfiles.com/2020/01/23/universal-memory-candidate-technology/
- https://blocksandfiles.com/2020/08/12/spin-memory-universal-selector/

Persistent Memory Technology

NVDIMM

Persistent Memory technology covers non-volatile memory and storage technology that sits between traditional DRAM and NVMe SSD in the storage hierarchy. Persistent memory has been referred to as Storage Class Memory. The primary characteristics of Persistent Memory are as follows:

1. Non-volatility (i.e. persistent)
2. Between DRAM and NVMe SSD with regards to the following characteristics
    a. Bandwidth
    b. Latency
    c. Cost/GB
    d. Programming/Address models (byte and/or block addressable)

The following graph from the Stanford EE380 EECS Colloquium presentation by Rick Coulson of Intel shows the latency reduction possible from moving non volatile memory from NVMe connectivity to DDR connectivity.

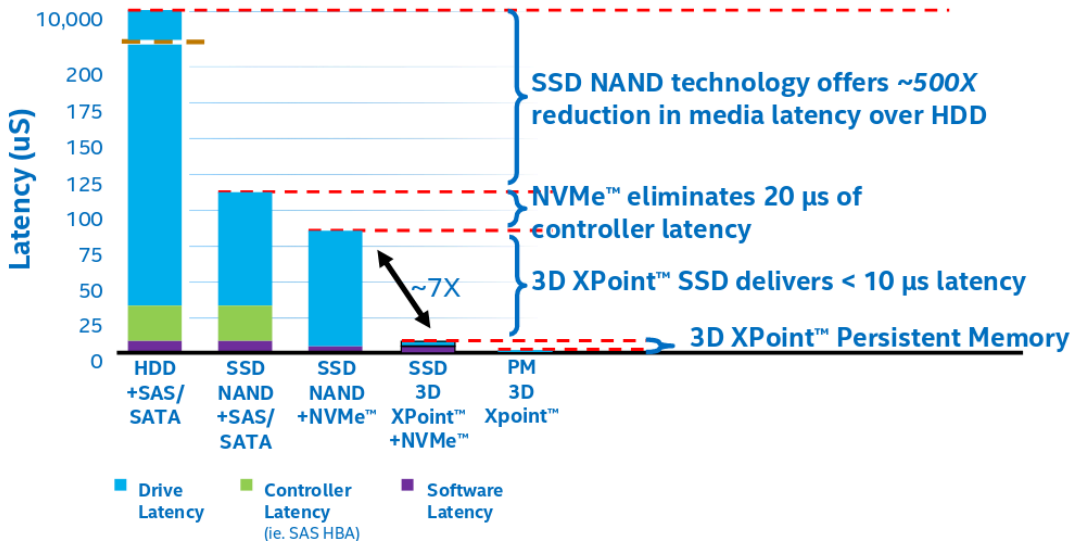

Storage Latency reduction with NVMe [Stanford EE380 EECS Colloquium by Rick Coulson]

The among the use cases for persistent memory are as follows:

1. Fast SSD - PM bypasses much of the I/O stack, both hardware and software, that slows down NVMe and SSD devices.
2. Non-volatile “DRAM” - PM can be used for application working data that needs to be persistent, but accessible like normal DRAM memory (byte addressable).
3. Expanded volatile memory - PM can be used to expand main memory (DRAM) capacity. Think of the NV memory in Persistent Memory as equivalent to traditional swap space, but significantly faster. This potentially allows for significantly large main memory, at a lower cost than a pure DRAM implementation.

The primary Persistent Memory Technology at this point in time are NVDIMM, that are non-volatile storage/memory that sits on the memory bus of the CPU (DDR4/DDR5). JEDEC categorizations of NVDIMM technologies are as follows :

1. NVDIMM-N - DRAM backed by non-volatile memory (transparent to software). Byte addressable.
2. NVDIMM-F - non-volatile storage connected to the CPU memory bus, like NVMe SSD, with DDR4 replacing NVMe/PCI-e. Block addressable.
3. NVDIMM-P - Enhanced memory bus protocol with explicit, standardized, support for persistent, block based memory. Expected as part of the JEDEC DDR5 memory standard.

NVDIMM-N is available from multiple vendors. Availability of NVDIMM-F is known, and NVDIMM-P awaits release of the JEDEC DDR5 standard.

Use of persistent memory technology is heavily dependent on BIOS/UEFI support and development of programming libraries. The former can be found in the ACPI standard and the latter can be found in the Persistent Memory Development Kit (PMDK) spearheaded by Intel. The PMDK allows applications to access persistent memory in ways defined by the SNIA NVM programming model.

Persistent Memory over Fabrics

The following diagram, taken from a presentation by Stephen Bates of Microsemi at the 2017 SNIA Persistent Memory Summit shows the use of non-volatile memory network fabrics, both legacy and new, to connect non-volatile memory to create new compute/storage architectures.

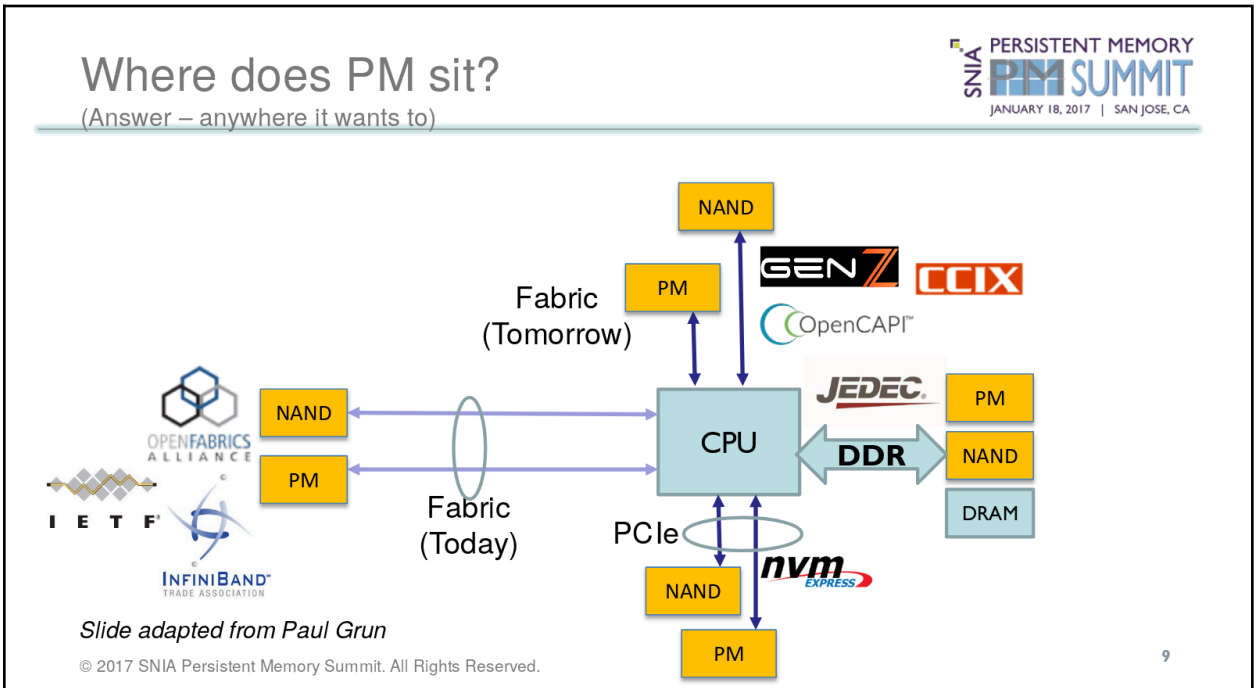

Integration of Persistent memory over fabrics in computing systems. Stephen Bates MicroSemi at 2017 SNIA Persistent Memory Summit.

Some additional information on Gen-Z, CCIX, and OpenCAPI is in the HEPiX CPU Working Group document. Some information about NVMe over fabric can be found below, in the Storage Infrastructure section of this document.

Possible impacts on HEP (and others)

Especially M-Type SCM is interesting for ultra fast persistency in the DAQ part of the experiments. Large scale storage systems have a strong demand for that type of systems to store and query file system metadata. Vendors have already shown system architectures capable of storing more than 100TB with ~100GB/sec squeezed in a 1U form factor - largely limited by the network subsystem to connect to the outside.

Summary

<Insert Summary Here>

Sources

- NVMexpress.org is the organization developing the NVMe standards
- The Joint Electron Device Engineering Council (JEDEC) is responsible for the three NVDIMM terminology, “-F”, “-N”, and “-P”. JEDEC is also responsible for DDR4 and DDR5 memory standards that include support for the various NVDIMM types.
- The ACPI specification part of the UEFI specification. UEFI support for NVDIMM support in ACPI 6.2 interface, including NVDIMM Firmware Interface Table (NFIT).
- SNIA NVM Programming Working Group has defined a programming model for non-volatile memory, the SNIA NVM Programming Model.
- The Persistent Memory Development Kit (PMDK) is an open source library for Linux and Windows that can be used by applications to access persistent memory in the ways outlined by the SNIA NVM Programming model. Development of the PDMK is managed by Intel. The web site for PDMK development is pmem.io
- An overview of the state of persistent memory technology can be found a the SNIA 7th Annual Persistent Storage Summit
- The annual Flash Memory Summit has a plethora of information on persistent memory. An archive of their proceedings since 2006 is at Flash Memory Summit Proceedings.
- SCMandMIEC_overview_12Feb2013
- http://smorgastor.drhetzler.com/wp-content/uploads/2014/08/SSD-Reliability.pdf

- The Memory Guy is a web site run by an industry analyst that covers memory technology. For those with money, a quick way to get an in depth view of the memory industry can be obtained by purchasing the analyst's memory technology reports.
- The Memory Guy's market research company is Objective Analysis. Although you need to pay for their in depth research reports, the outlines for the reports are freely available and provide an high level view of the important areas memory technology