A Brief History of Memory: From RAMAC to Flash and Beyond

IEEE-CNSV
Consultants' Network of Silicon Valley

Andy Walls, IBM Fellow, CTO and Chief Architect IBM Flash Storage

13 November 2018

Let’s start in the middle . . . . 1981 – 37 years ago

A young impressionable 22 year old new engineer begins his career straight out of UCSB in Rochester, MN

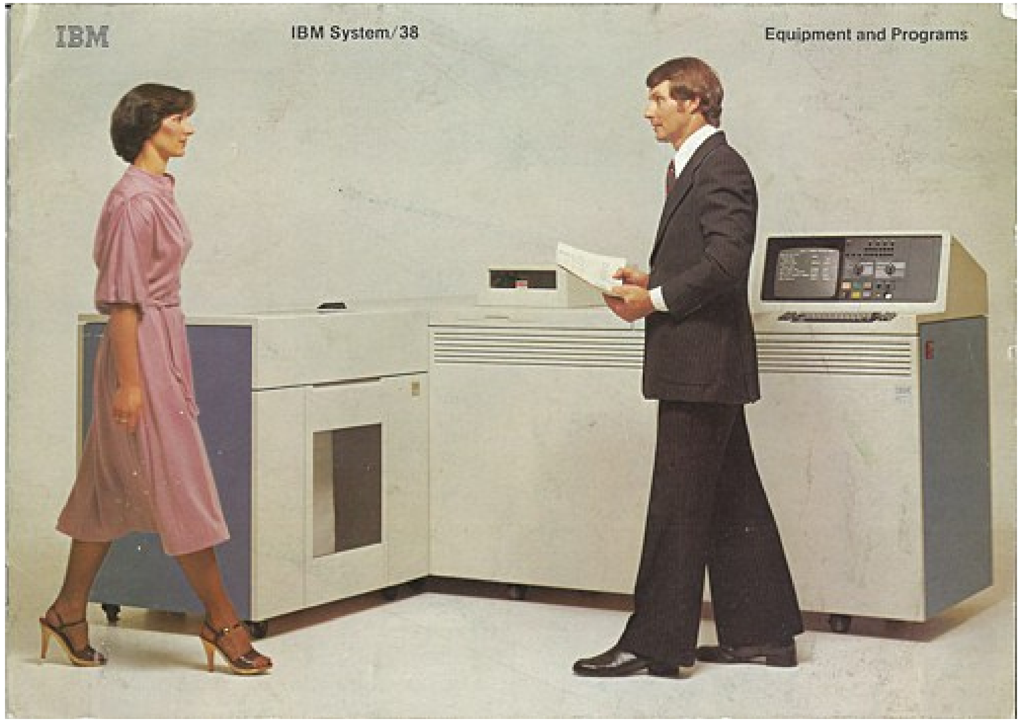

IBM System /38

IBM Rochester, MN

My First Day of Work

"Here is your System /38 – You will be owning and modify a backup and restore program for DASD to Tape.

IBM 3370 Direct Access Storage Device

571MB

IBM 3410/11 Tape Unit

IBM 62PC Direct Access Storage Device

65MB

The 3370 DASD

- 1979
- Introduced thin-film head technology to large disk files. Work on thin-film head structures was started in the late-1960s. (Watson Research)
- 571.3MB
- Up to 4 of these could be attached to a System 38 for a whopping 2.285.5 GB of DASD!!!
- Average seek time was 20 milliseconds and the nominal data rate was 1.859 megabytes per second.
- At announcement, a Model A11 could be leased for $900 a month, rented for $1,058 a month or purchased for $35,100. Corresponding charges for each B11 unit were $600, $705 and $23,400.
- ***<u>SO: $61,438 a GB!</u>***

The 3431 Tape Drive

- 1971 – withdrawn in 1987
- The 3410 subsystem had a new, desk-high design with tape reels mounted horizontally instead of vertically, as in most tape drives
- Control circuitry for as many as six tape drives was built into one of the tape units, eliminating the need for a separate control unit and saving floor space.
- 20K to 80K Read/Write Speed per second!!
- 800 bits per inch.
- Purchase prices ranged from $10,200 to $16,400.

The 62PC Piccolo – part of 3310 DASD

- 1979 – Hursley, UK
- The IBM 62PC (Piccolo) 65MB capacity drive introduced in 1979 used six 8-inch disks with one surface dedicated to servo position data.
- 360,000 sold between 1979 and 1990.
- Average seek time was 27 milliseconds and the nominal data rate was 1.031 megabytes per second.

Let's go back to 1956 . . . *Maybe the most important year in the history of technology*

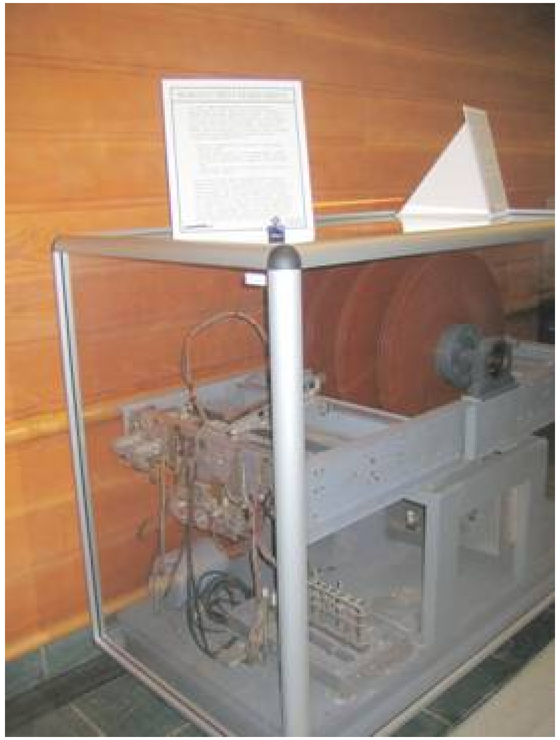

A Prototype of the first Disk Drive – IBM 350 part of RAMAC 305

The San Jose Mercury News ran a story under the headline, "A machine with super memory!

"The information on the discs can be added to, altered or erased at will. Card-sorting, one of the most time consuming office-machine processes, is eliminated or greatly reduced."

RAMAC was the first computer to use a random-access disk drive—the 350 Disk Storage Unit—the progenitor of every hard disk drive made in the 62 years since.

Reynold Johnson

How it all happened

- In 1952, with computer excitement in the air, IBM sent Reynold Johnson to San Jose to start a new research lab.
- Lab would focus on random access storage
- When the Air Force wanted a random access inventory system, Johnson set his 50-person lab in motion.
- They Tried everything—strips, rods, tapes, flat plates, you name it. In the early 1950s, no one had any idea how to make a fast, reliable random access memory machine.
- Lab settled on rotating magnetic disks.
- A single aluminum disk warped at high speeds
  - So they figured out gluing two together did the trick
- The arm was even tougher to figure out
- William Goddard and John Lynott invented an arm that shot out compressed air to keep it above the disk.

The Result

- In 1952, with computer excitement in the air, IBM sent Reynold Johnson to San Jose to start a new research lab.
- Lab would focus on random access storage
- When the Air Force wanted a random access inventory system, Johnson set his 50-person lab in motion.
- They Tried everything—strips, rods, tapes, flat plates, you name it. In the early 1950s, no one had any idea how to make a fast, reliable random access memory machine.
- Lab settled on rotating magnetic disks.
- A single aluminum disk warped at high speeds
    - So they figured out gluing two together did the trick
- The arm was even tougher to figure out
- William Goddard and John Lynott invented an arm that shot out compressed air to keep it above the disk.
- A single arm and head going to one of 50 platters.

The result

- A total of 5M Characters – each character was 6bits coded into 7.
- Every character could be read or stored independently
- 1000 sold
- 50 inches long, 68 inches high and 29 inches deep.
  - ***1 TON!***
- 1200 RPM
- 600mS seek time
- 100 BPI
- 8.8K per second transfer rate
- It is claimed that the Board of Directors canceled it due to threat to the punch card business.
  - San Jose kept working on it until the President approved it.

$10,000/MB

Patent 3,503,060 is the original patent for the HDD

The Impact

- It introduced the concept of instantly accessible information.
- Before RAMAC, information had to be entered by running a stack of cards through a punched card machine
- Answers would arrive in hours or days.
- RAMAC could find data in seconds, alter it, and move on to find a completely different piece of data. It let enterprises think about data in new ways, mixing and matching it on the fly.
- Random access made the relational database possible.
- **It introduced decades of improvements, cost reductions and innovation arguably unmatched in any industry.**
    - And yes Mr. Processor – I include you too!

One more incredible story. . . . IBM 1301

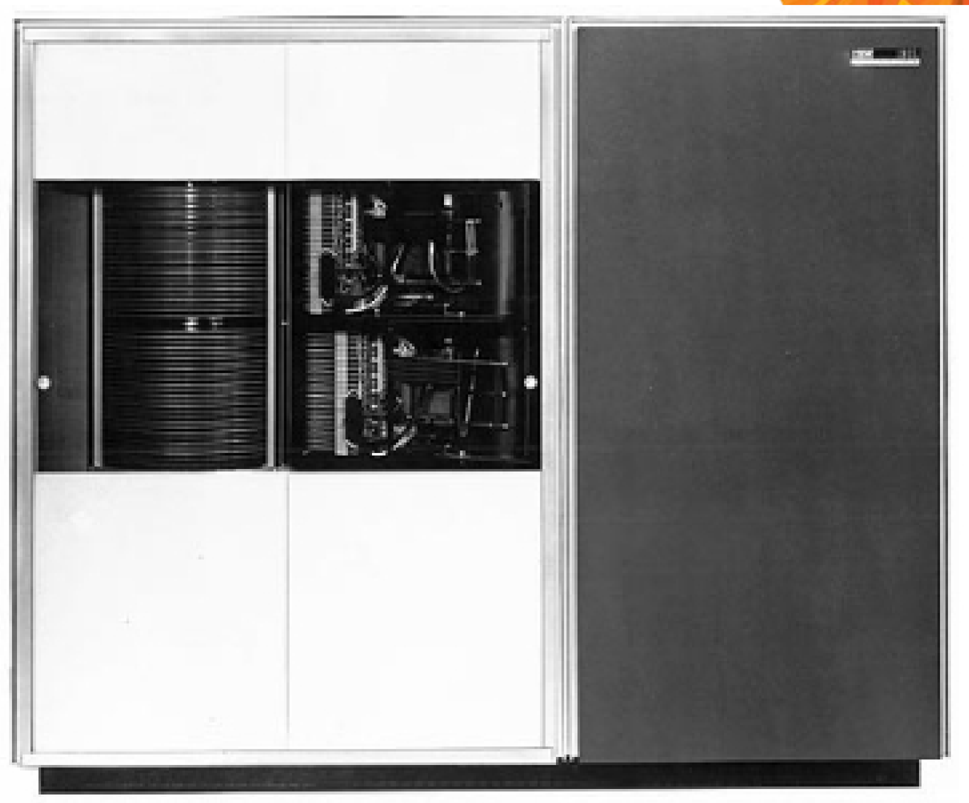

One more incredible story. . . . IBM 1301

- 1800 RPM – 90KB per second
- 50 TPI
- Head moved to 250 microinches from surface
- Pioneered two key things
    1. Self actting, air bearing slider technology
    2. Separate head per surface
- 28MB to the module
- 3Q1962 – Used in SABRE reservation system
- $115,000 - $4,142 per MB

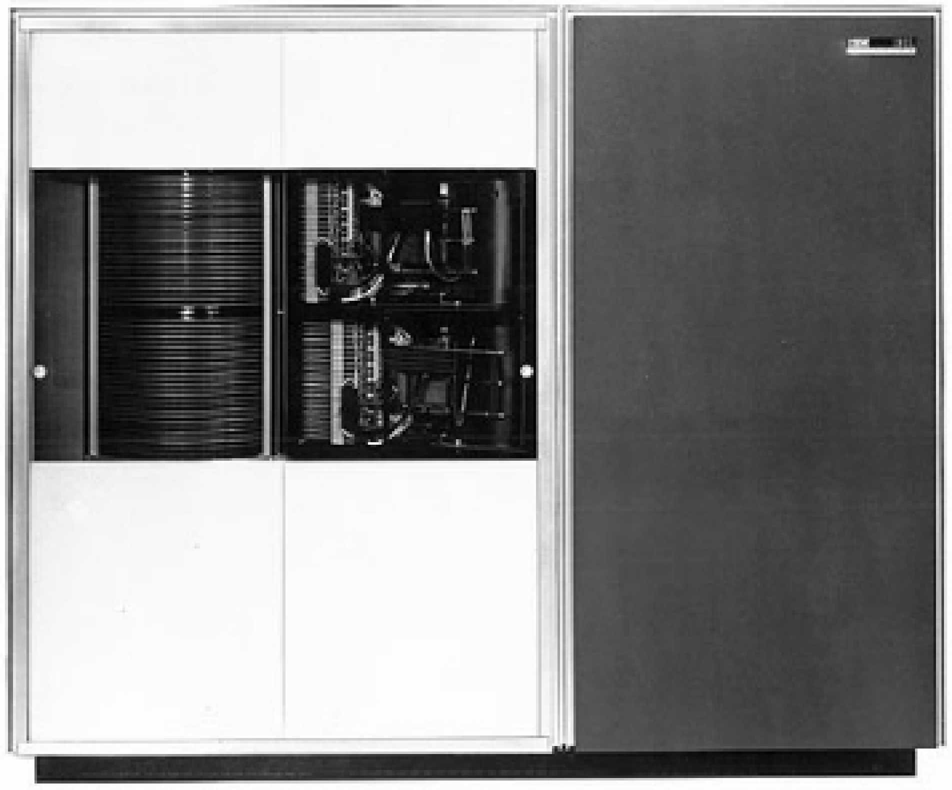

And then things get smaller, and denser, and faster. . . . .

What followed

- Some technologies came and went – Fixed Head
- 1973 – the Winchester drive
  - Heads landed on special area reducing cost of head actuators
- The PC Revolution
- SCSI happened – ATA happened
- Lots of competitors like HP, Seagate, Burroughs, Connor peripherals, Data General, Fujitsu, Hitachi, Honeywell Bull, Iomega, Memorex, NEC, Quantum, Sequel, Shugart, Siemens, Sony, Storage Tek, Syquest, Western Digital and many, many others
- Things start to commoditize – 5.25 then 3.5 then 2.5”

Radical Improvements Followed

Access Latency

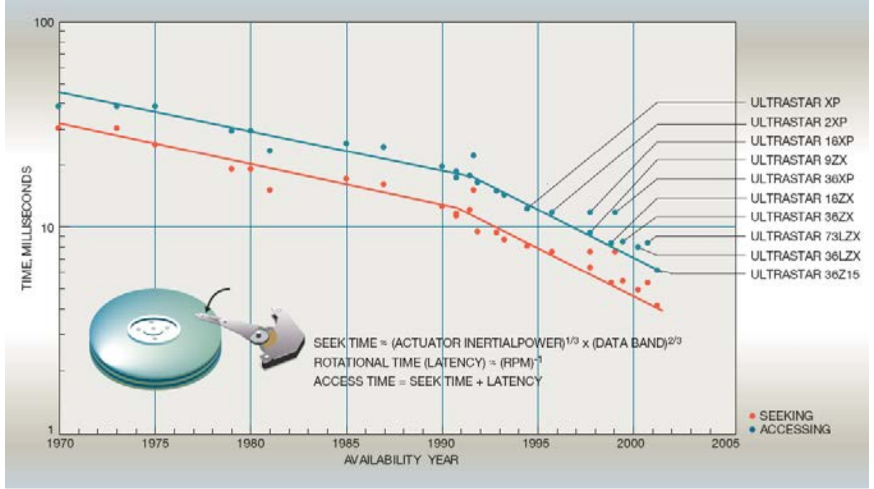

Data Rate

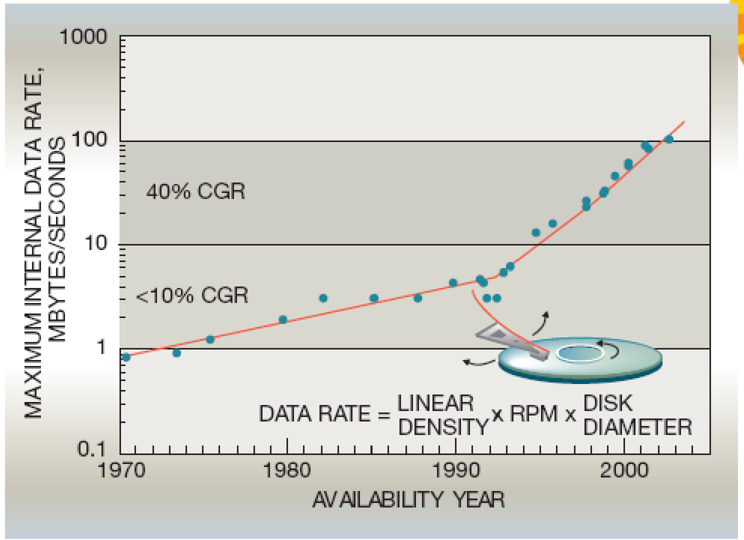

Areal Density

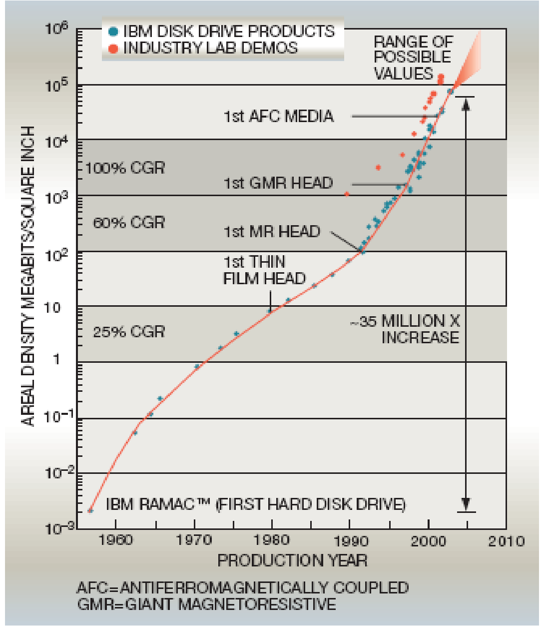

Areal Density perhaps best shows what has happened since

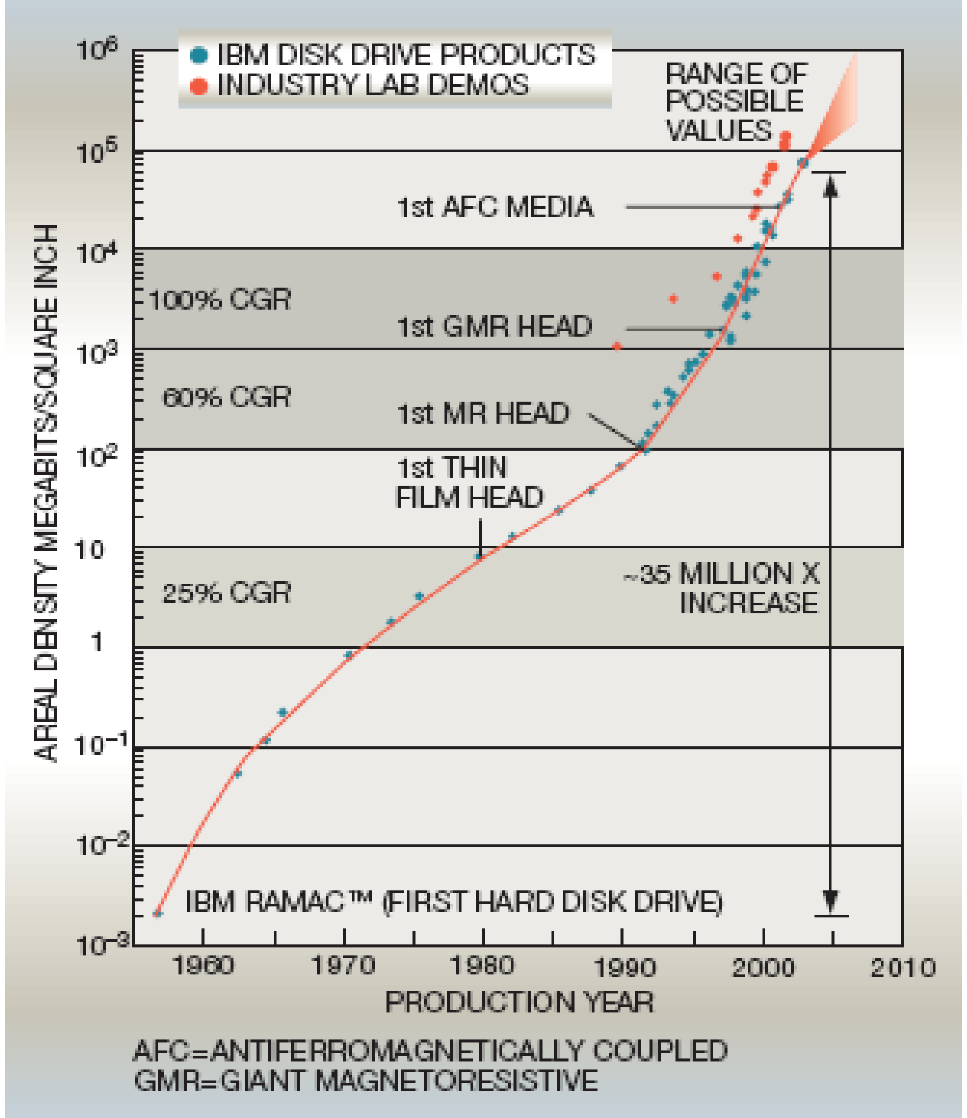

Storage Technologies Areal Density Trends

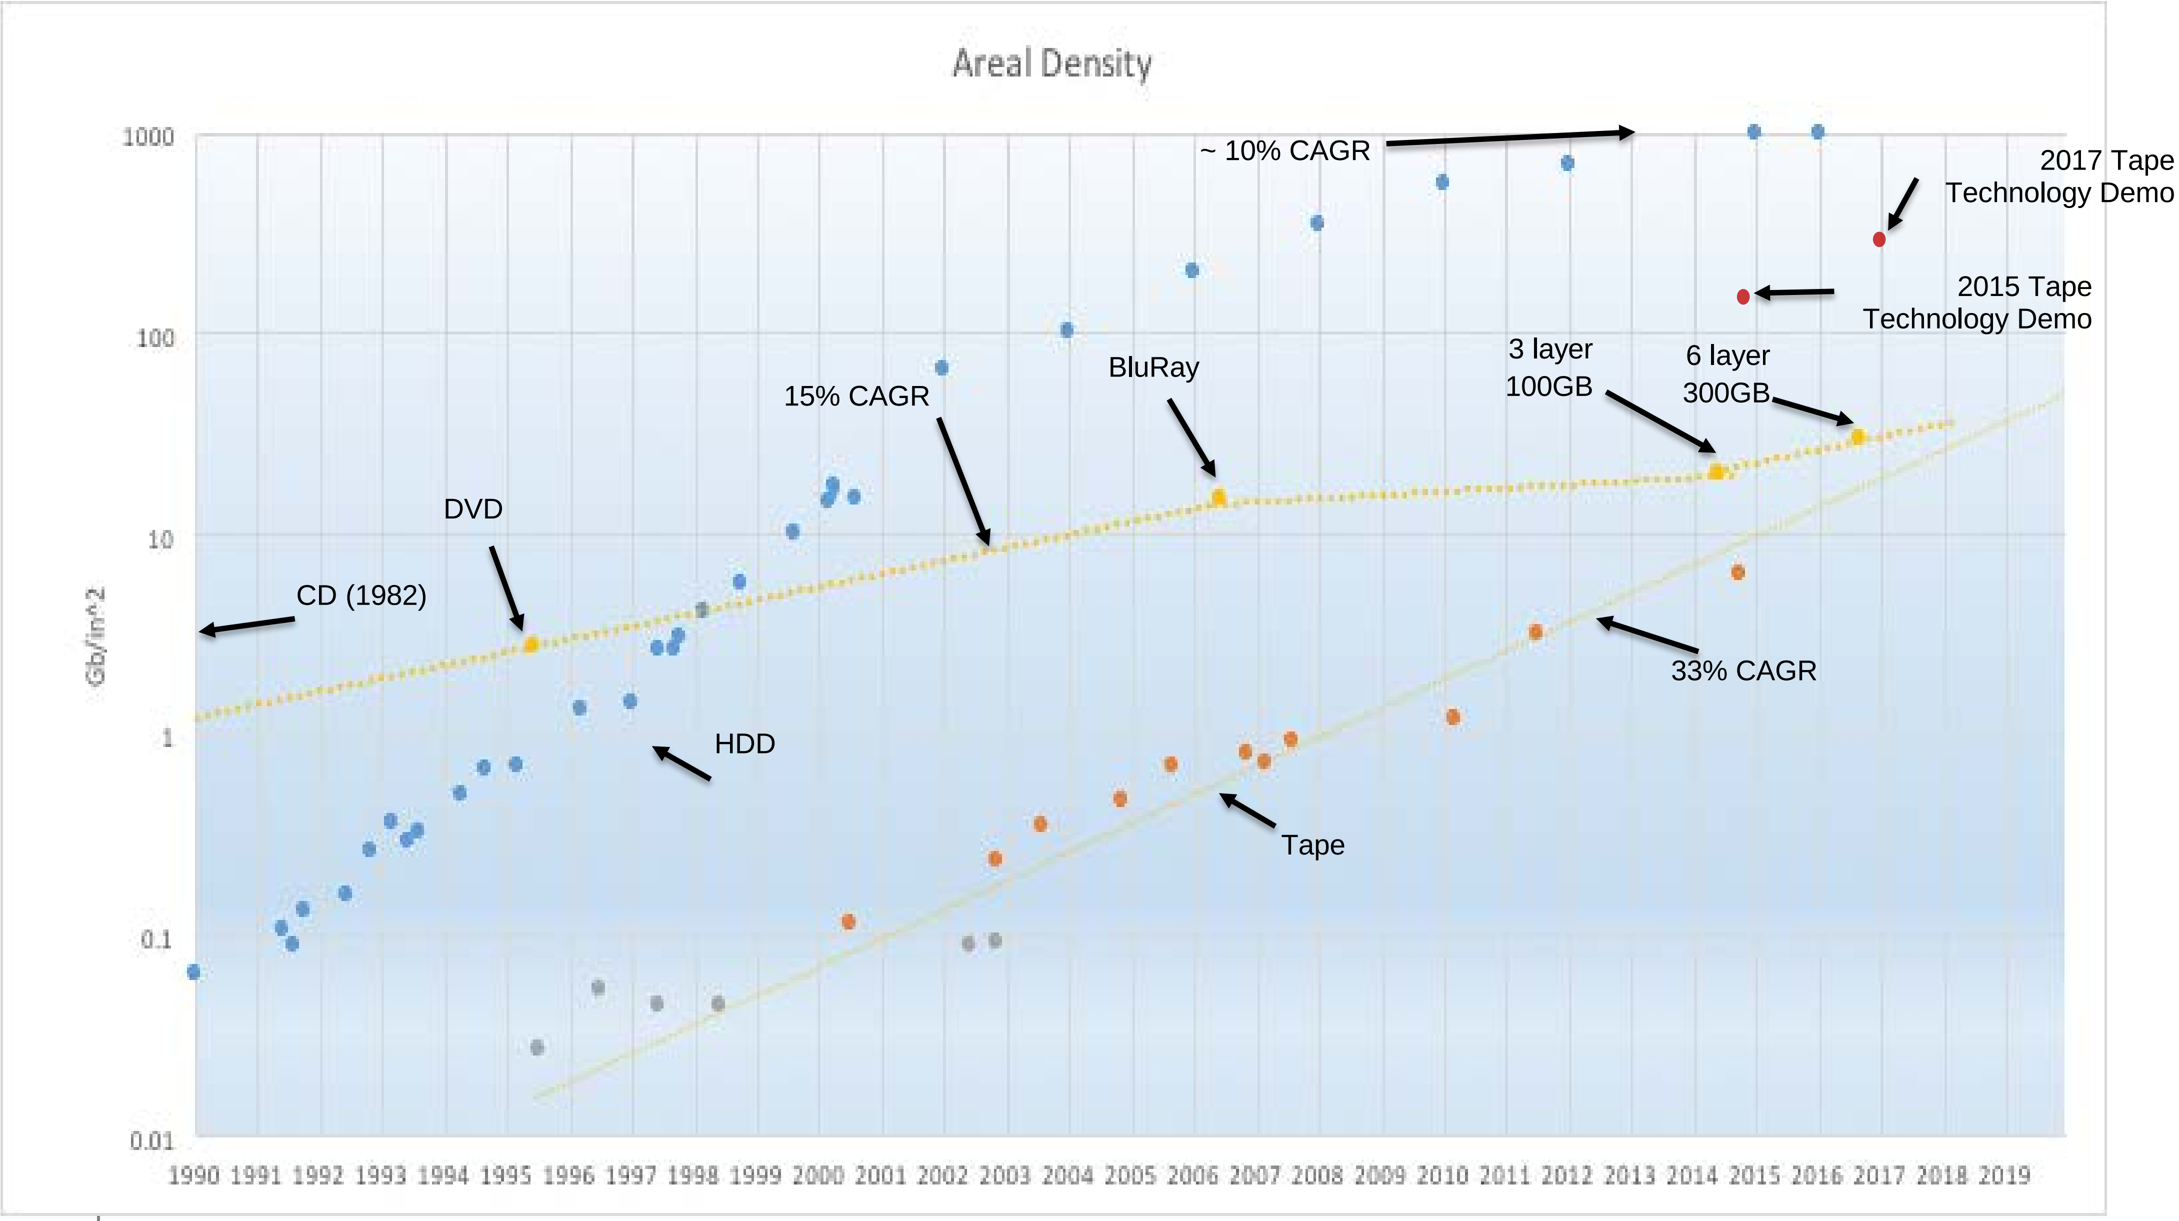

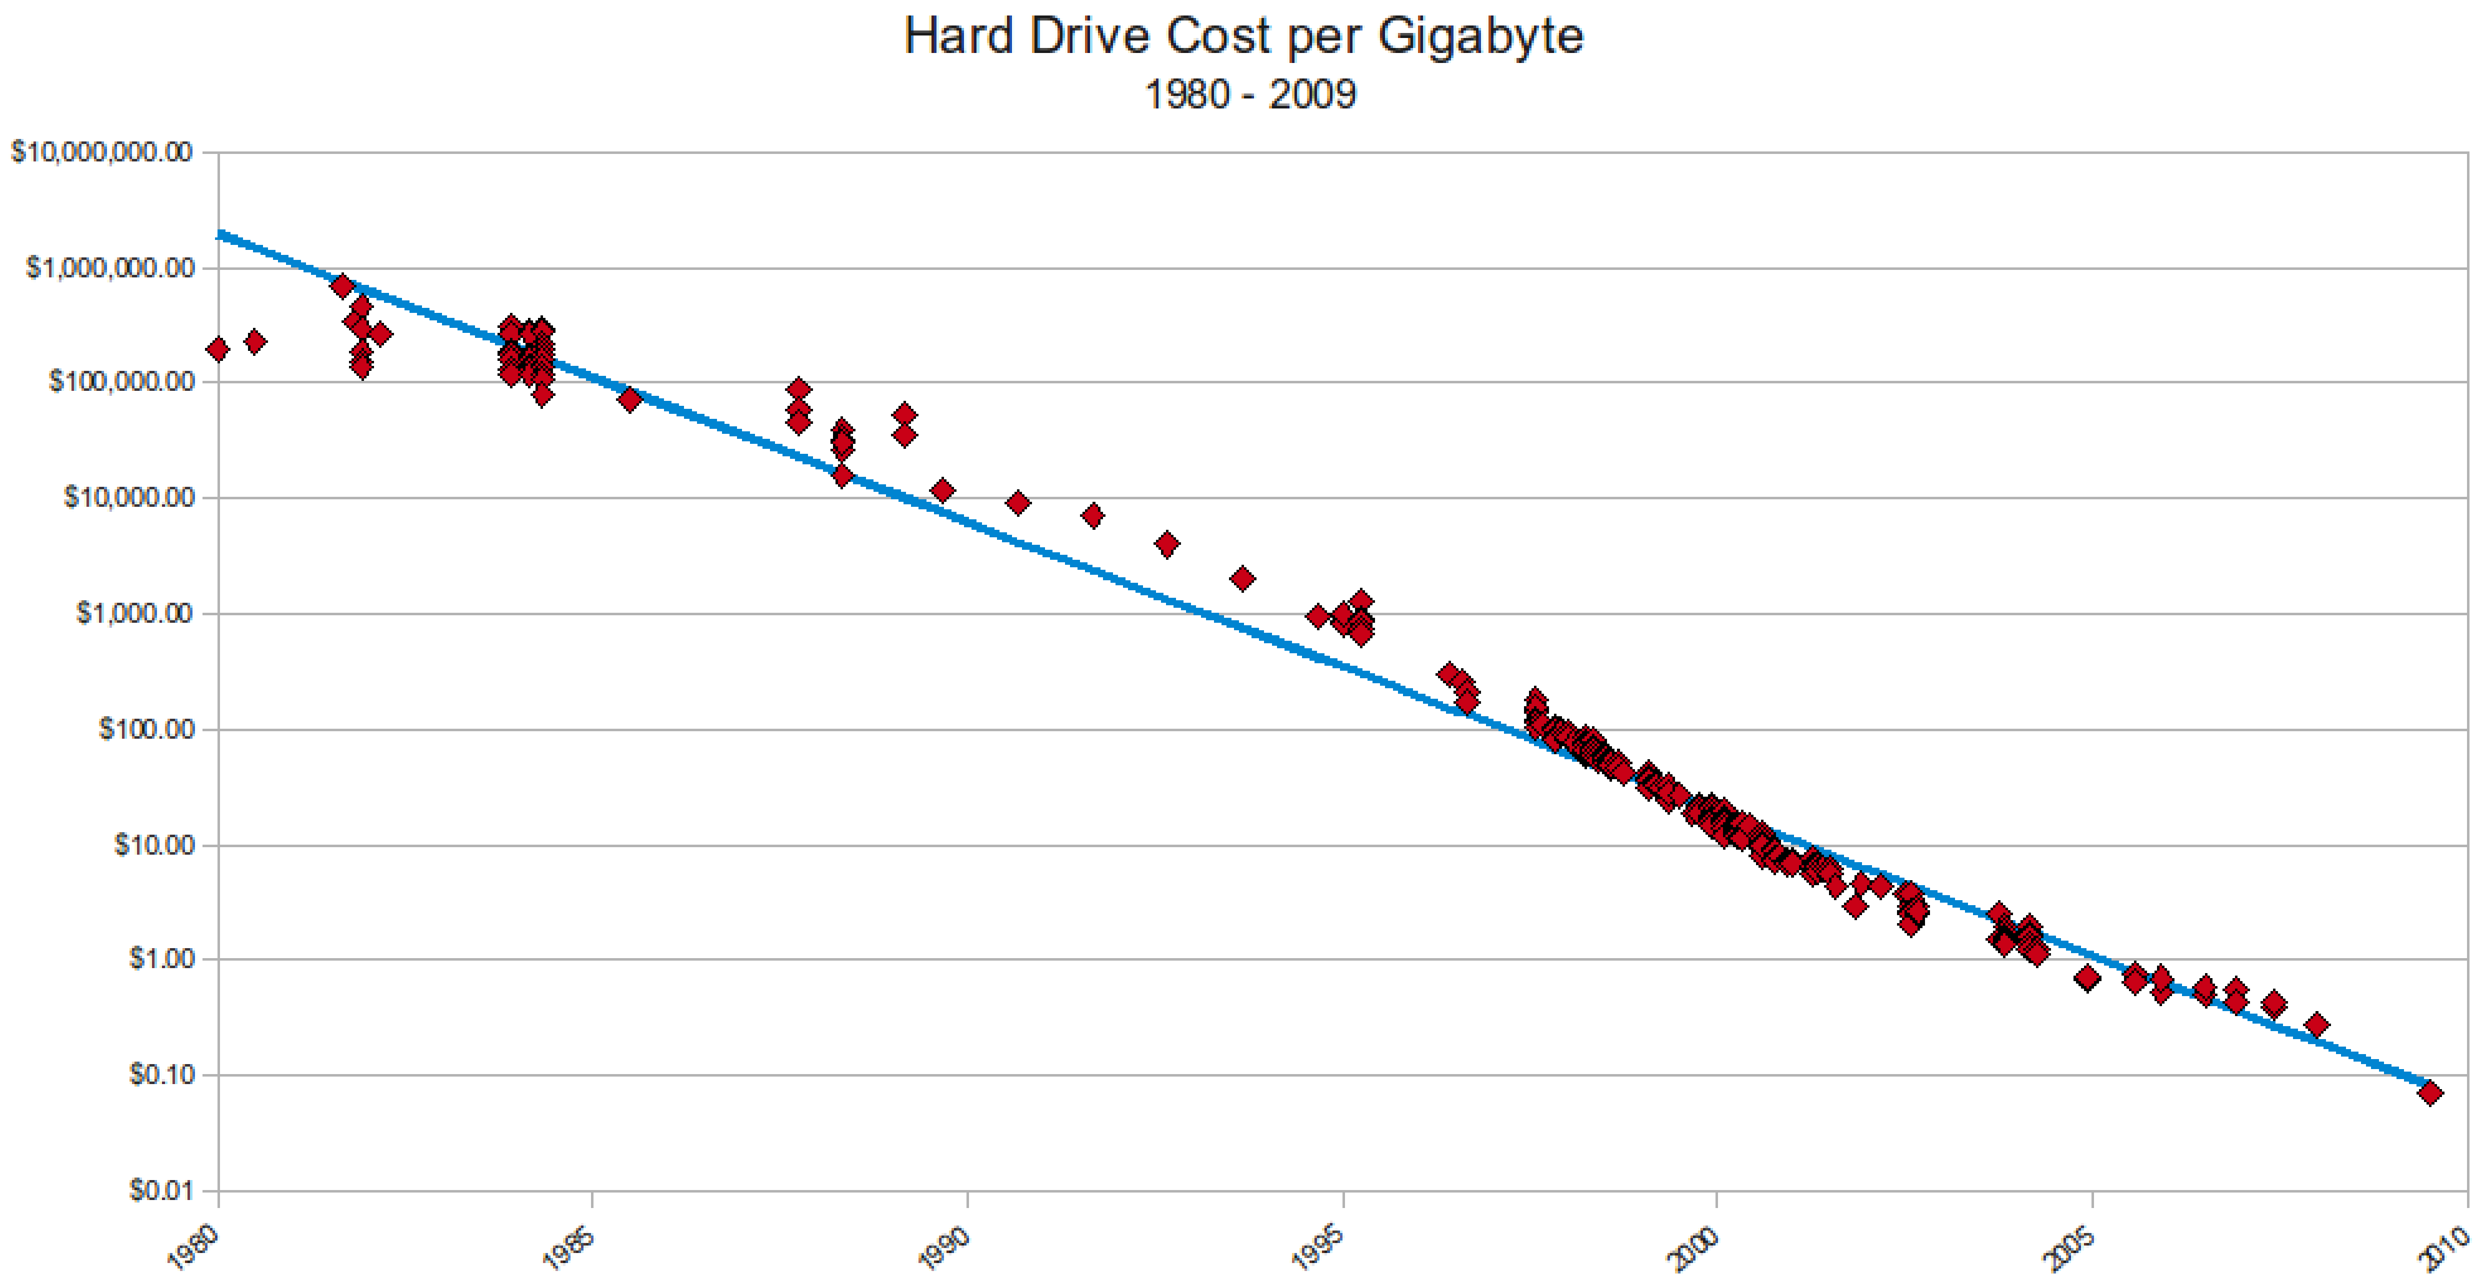

Hard Drive Cost per Gigabyte
1980 - 2009
$10,000,000.00
$1,000,000.00
$100,000.00
$10,000.00
$1,000.00
$100.00
$10.00
$1.00
$0.10
$0.01
1980
1985
1990
1995
2000
2005
2010

But the HDD itself is only part of the story. . . .

- Standard interfaces
- RAID gets invented
- Individual HDDs turn into arrays
- SANs – Fibre Channel and FICON
- Read and Write Cache
- Short Stroking,
- Snapshots and flashcopy
- Sync and Async remote copies
- Etc.

And then there was flash!

Historical Cost of Flash Storage

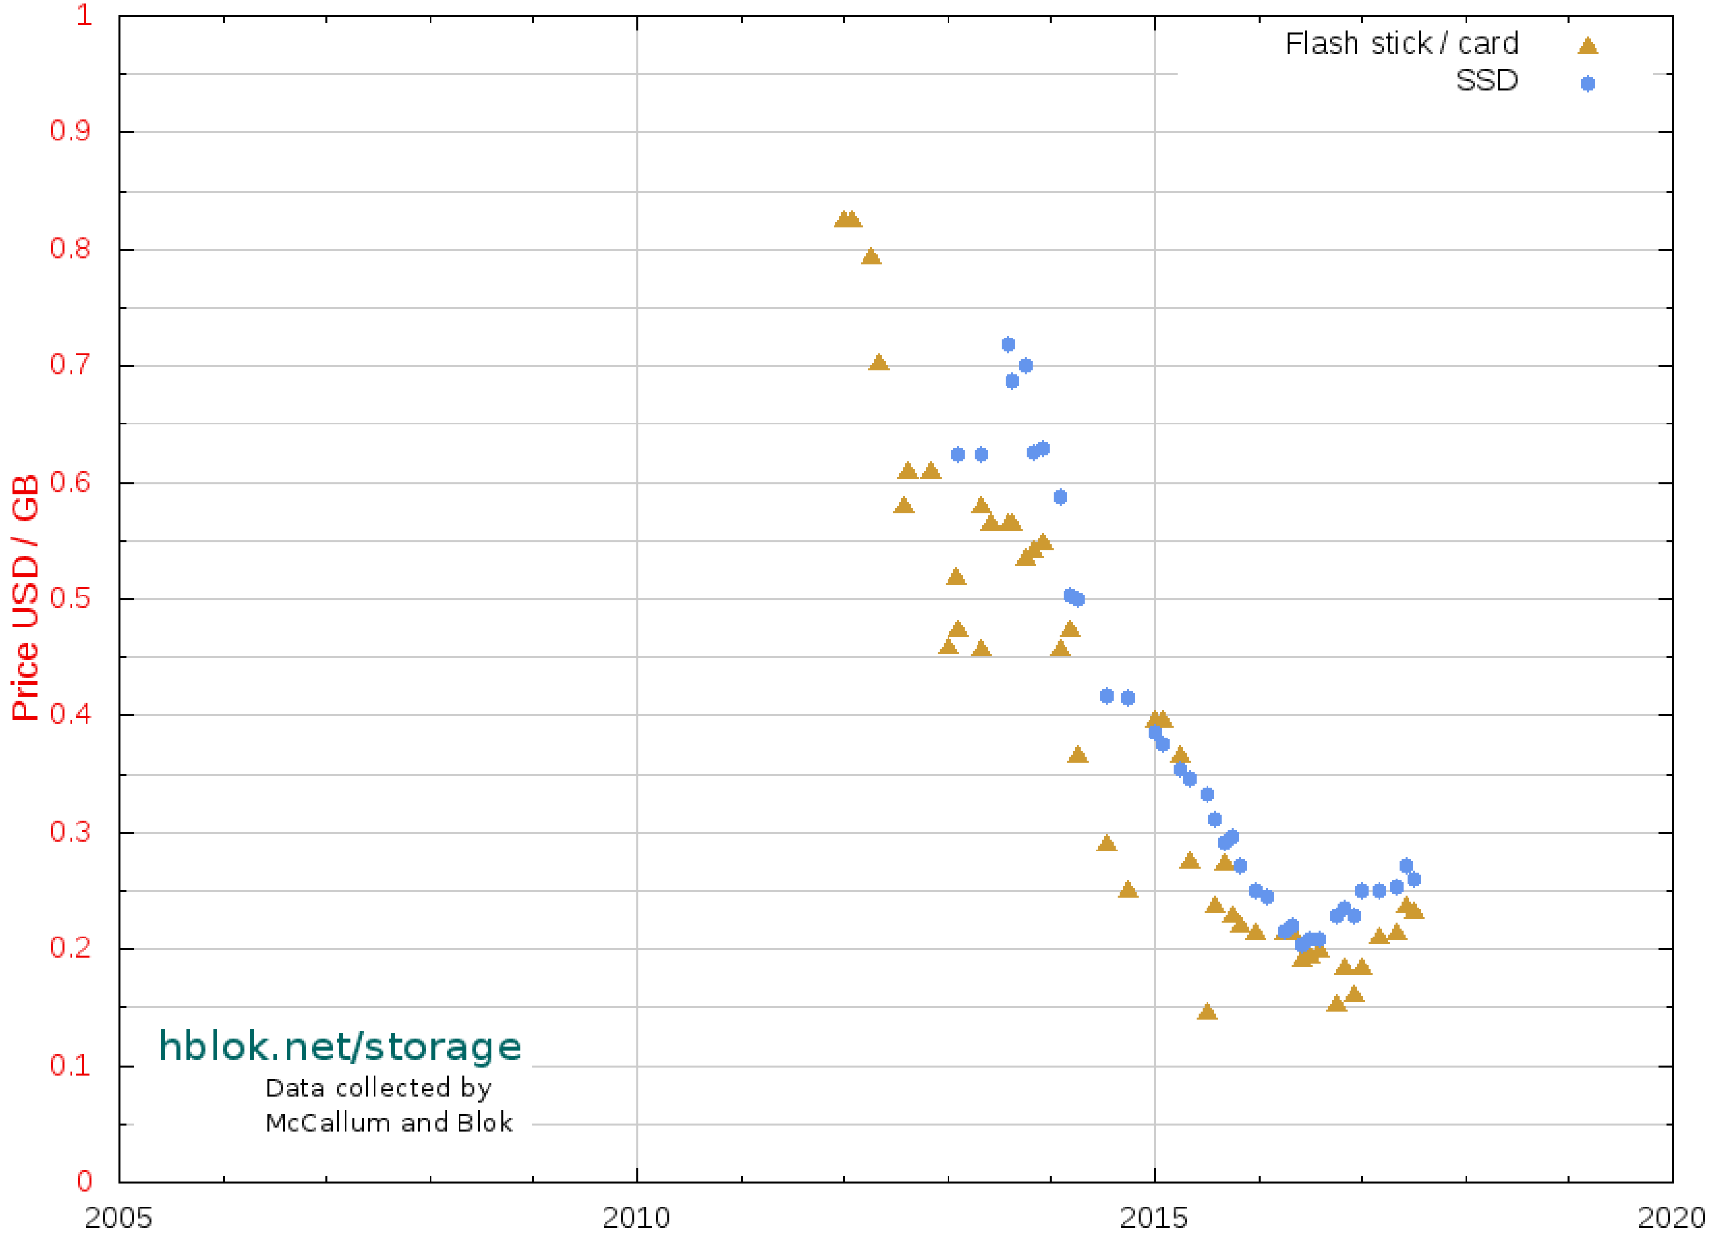

How Flash Gets Denser (and Cheaper)

1. Denser Lithography

Lower endurance

Limits to improvements – the end of Moore's Law

2. More Bits per cell

Signal processing to discriminate between 2, 4, 8 or more signal levels when a cell is read.

Improvement in ~2X density when a step is taken (about every 3-4 years)

Lower Endurance

Now on 3 bits per cell (TLC), QLC may deliver in 2020.

Not clear if we ever get 5 bits per cell for Enterprise

3. Skyscraper

Who builds higher?

Larger cells, higher endurance, MLC ok to use in 'healthy' 19~20 nm lithography

Current best is 64 high. 200 or more high may be possible.

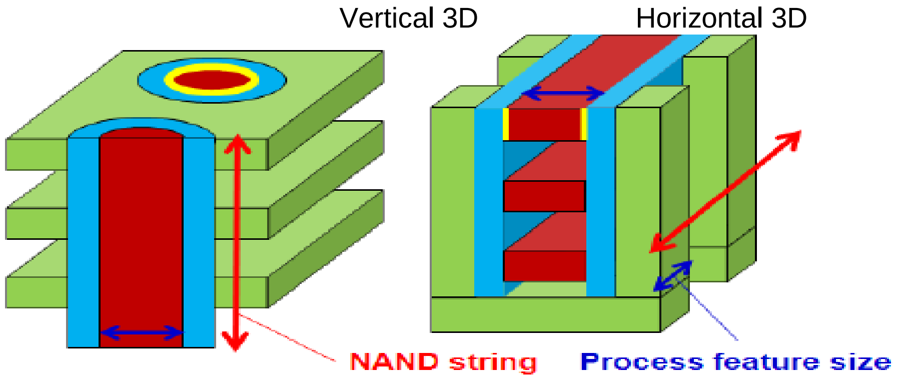

Source: Akira Goda, Micron: "Opportunities and Challenges of 3D NAND Scaling"

Source: Samsung

Source: Akira Goda, Micron: "Opportunities and Challenges of 3D NAND Scaling"

Directions in Storage Media

On-line data shifts to Flash, with $/IO/Sec advantage. HDD remains as 'active archive'
Efficiency technology is needed (compression, dedup, etc.)

- **Flash and other Solid State take all performance workloads**
- **3DXP – non-volatile memory extension**
- **Flash media cost improving at 25-30% CGR**
- **HDD improvement rate slows to 19% CGR or less**
- **Storage becomes Expensive**
  - Media cost improvements will not offset capacity growth
  - Technologies of heirarchies, compression, dedup, lifecycle management and defensible deletion are critical
  - A need for Cold Storage technologies with companion indexing, automated data placement and movement – Tape and Optical

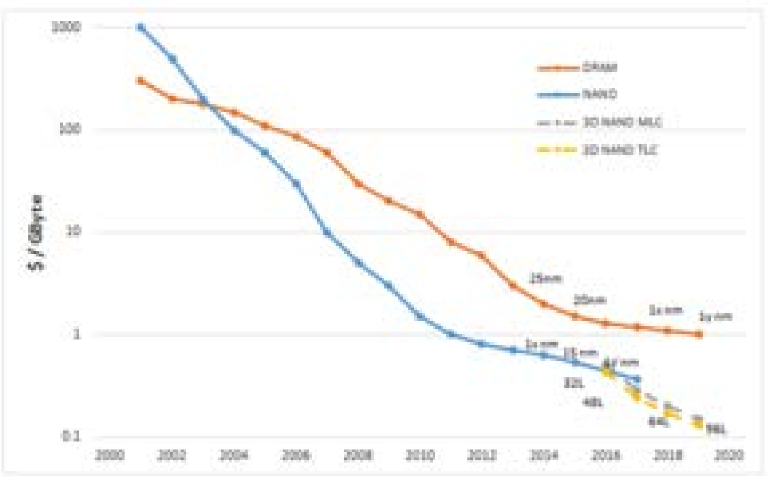

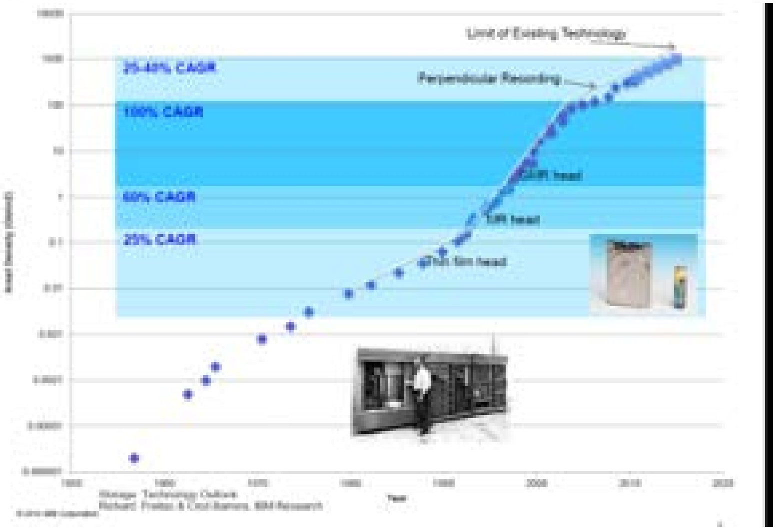

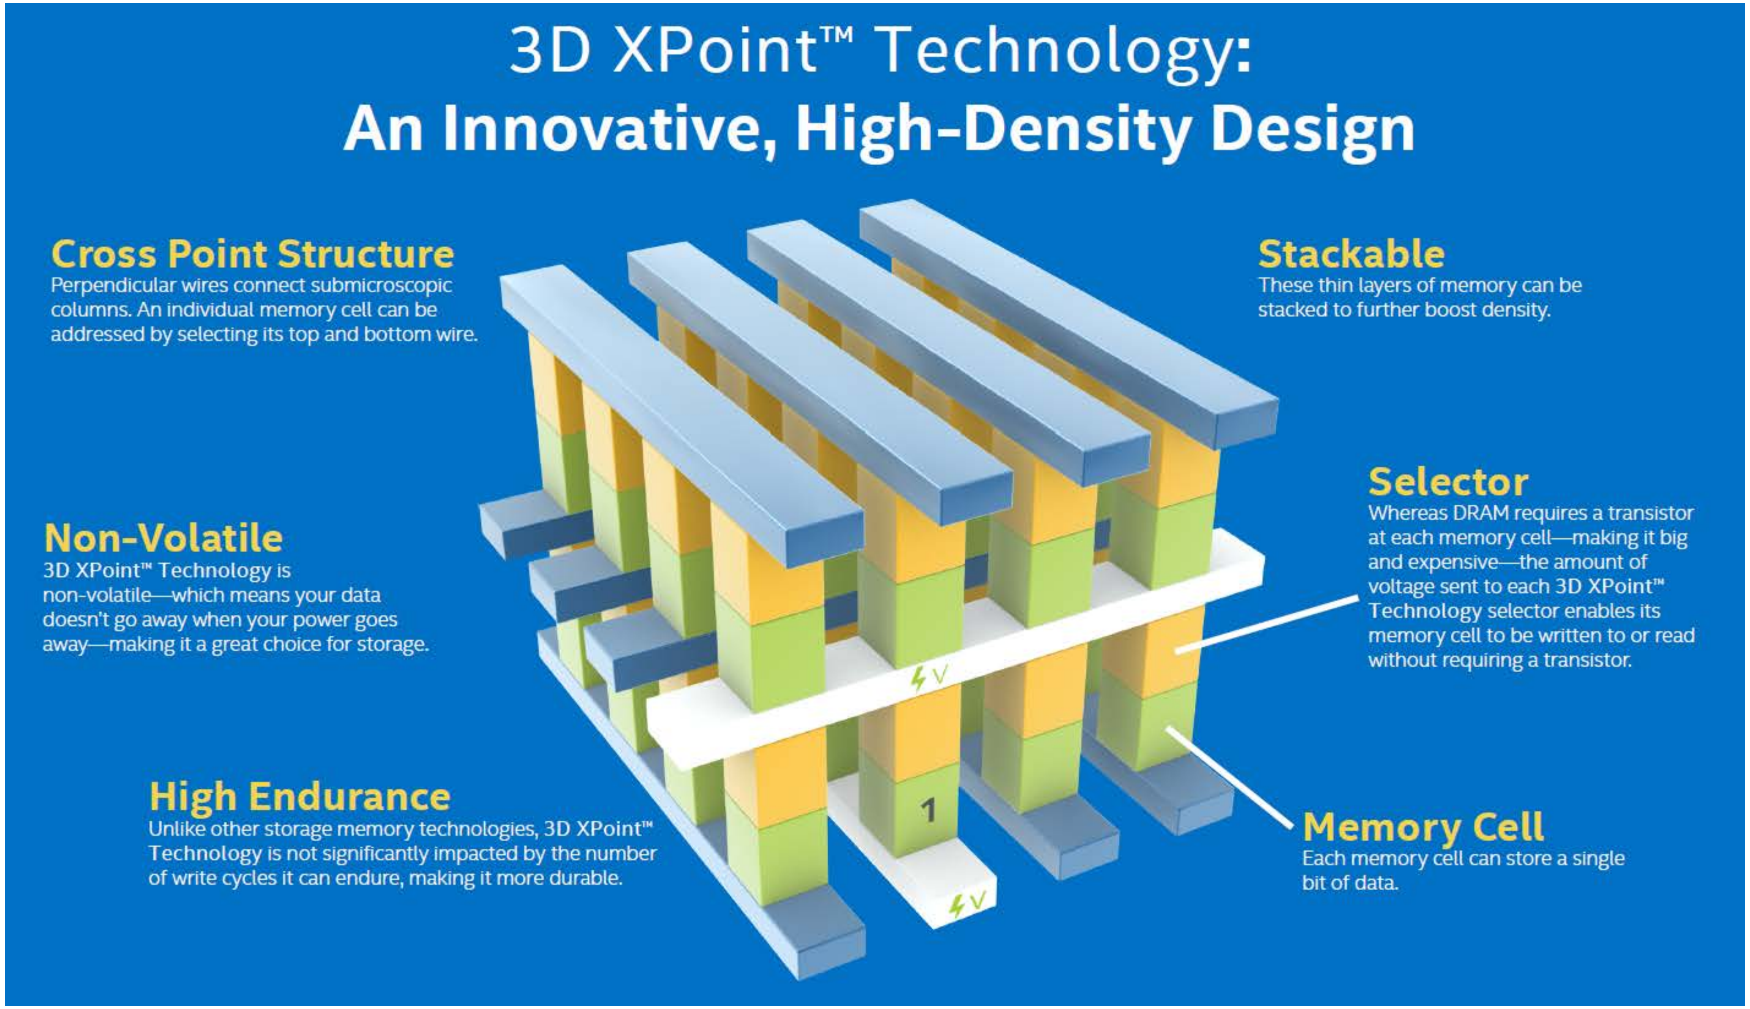
3D XPoint™ Technology:
An Innovative, High-Density Design
Cross Point Structure
Perpendicular wires connect submicroscopic columns. An individual memory cell can be addressed by selecting its top and bottom wire.
Stackable
These thin layers of memory can be stacked to further boost density.
Selector
Whereas DRAM requires a transistor at each memory cell—making it big and expensive—the amount of voltage sent to each 3D XPoint™ Technology selector enables its memory cell to be written to or read without requiring a transistor.
Non-Volatile
3D XPoint™ Technology is non-volatile—which means your data doesn't go away when your power goes away—making it a great choice for storage.
High Endurance
Unlike other storage memory technologies, 3D XPoint™ Technology is not significantly impacted by the number of write cycles it can endure, making it more durable.
Memory Cell
Each memory cell can store a single bit of data.
V
V
1

Storage Class Memories

- Bubble memory (Remember that)
- MRAM – actually very popular today
  - (magnetoresistive random access memory) is a method of storing data bits using magnetic states instead of the electrical charges used by dynamic random access memory (DRAM).
- ReRAM – Resistive RAM
-

Clients face explosive growth in Unstructured Data

- *Structured Data grows too, the need there is for dramatic improvements in IO latency and cost*
- *Unstructured data: Mobile applications, Internet of Things, Social Media, Analytics*
- *Problem: Legacy Storage Designed for Systems of Record lack the scale and cost we need*

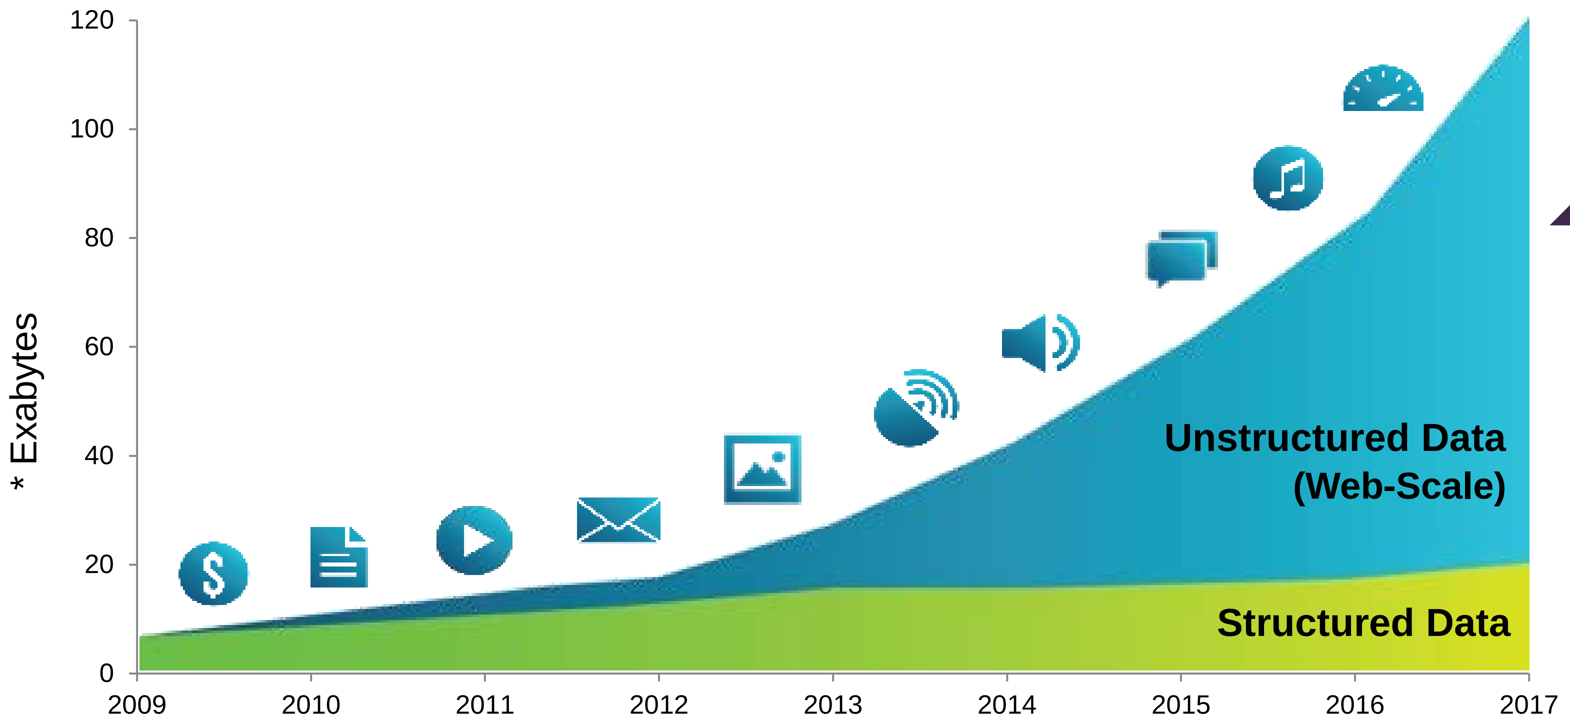

Source: IDC

Unstructured data growth of 60–80% per year creates Web-Scale storage needs

At 80% growth rate, we need a 44% cost decrease in Storage to maintain flat budgets

What is Phase Change Memory (PCM, 3DXP)?

- A new memory technology based on the resistance change of certain chalcogenide alloys upon application of heat (induced by application of current)
- Faster than Flash, but more expensive
- 3DXP now in use for server memory extension with NVMe over PCIe
- Phase-change materials are also used for rewritable data storage in optical disks such as DVDs and Blu-Ray disks - In optical disk technology, usage is made of the property of phase change materials to also change their optical reflectivity upon application of heat (induced by a laser beam)

From optical disks to semiconductors

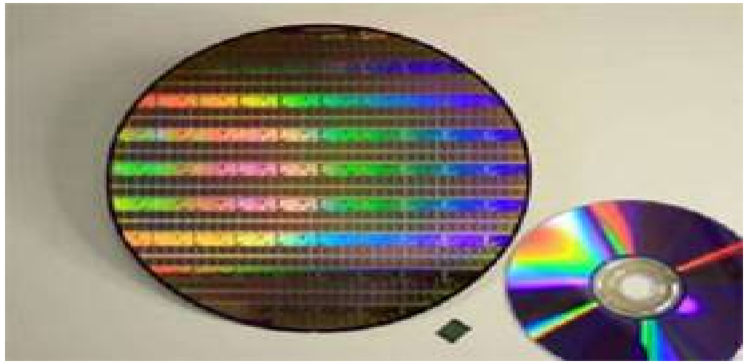

Phase change memory chip

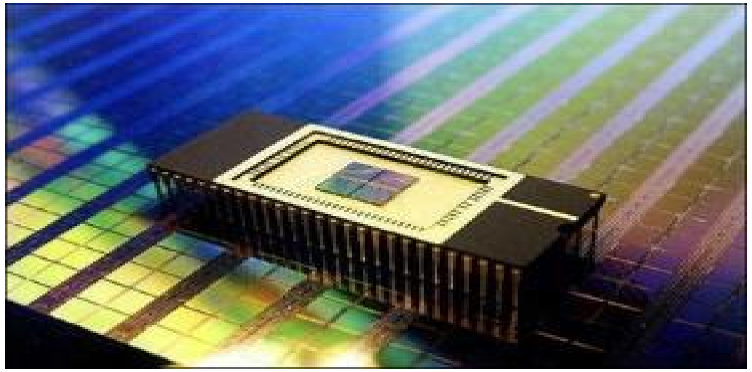

Phase change memory cell

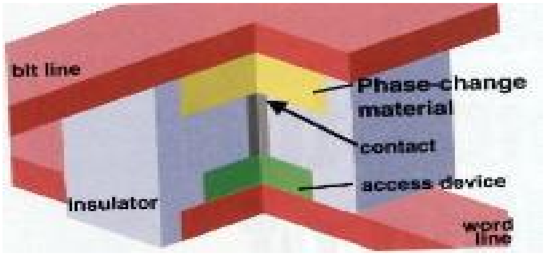

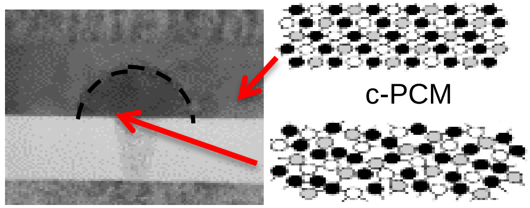

Atomic Scale Memory

World's Smallest Storage Device (IBM 2015)

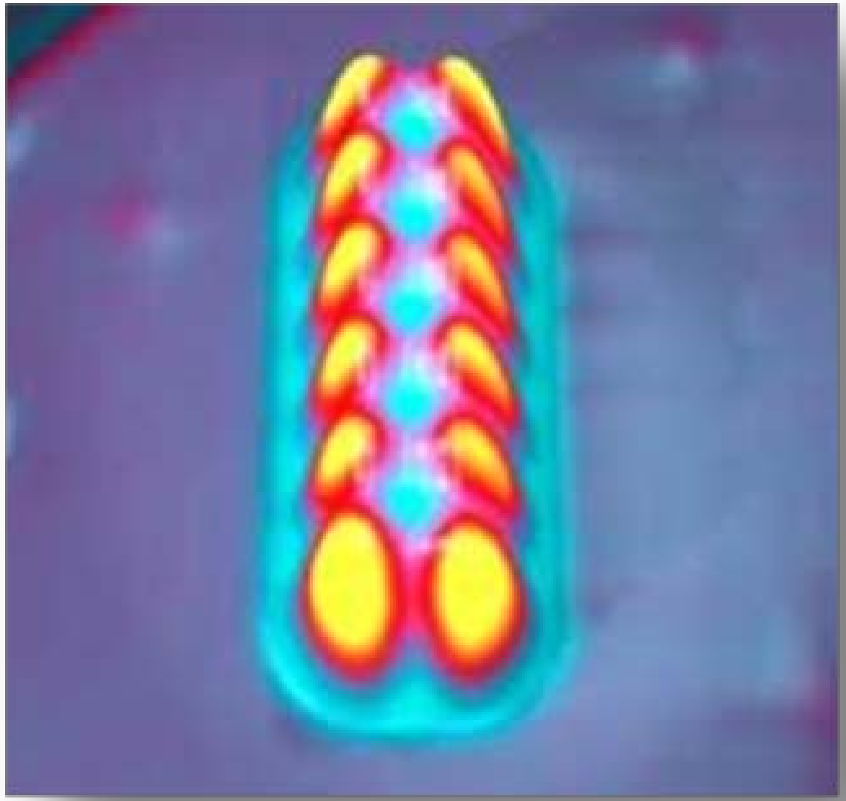

12 Atoms per Bit

Atomic-scale magnetic memory aerial density is potentially:

100x denser than today's hard disk drive technology

160x denser than NAND Flash

417x denser than DRAM

10,000x denser than SRAM

Single Atom Bit

March 2017

Magnetic bits lie at the heart of hard-disk drives, tape and next-generation magnetic memory," said Christopher Lutz, lead nanoscience researcher at IBM Research – Almaden in San Jose, California. "We conducted this research to understand what happens when you shrink technology down to the most fundamental extreme -- the atomic scale."

By starting at the smallest unit of common matter, the atom, scientists demonstrated the reading and writing of a bit of information to the atom by using electrical current. They showed that two magnetic atoms could be written and read independently even when they were separated by just one nanometer – a distance that is only a millionth the width of a pin head. This tight spacing could eventually yield magnetic storage that is 1,000 times denser than today's hard disk drives and solid state memory chips.

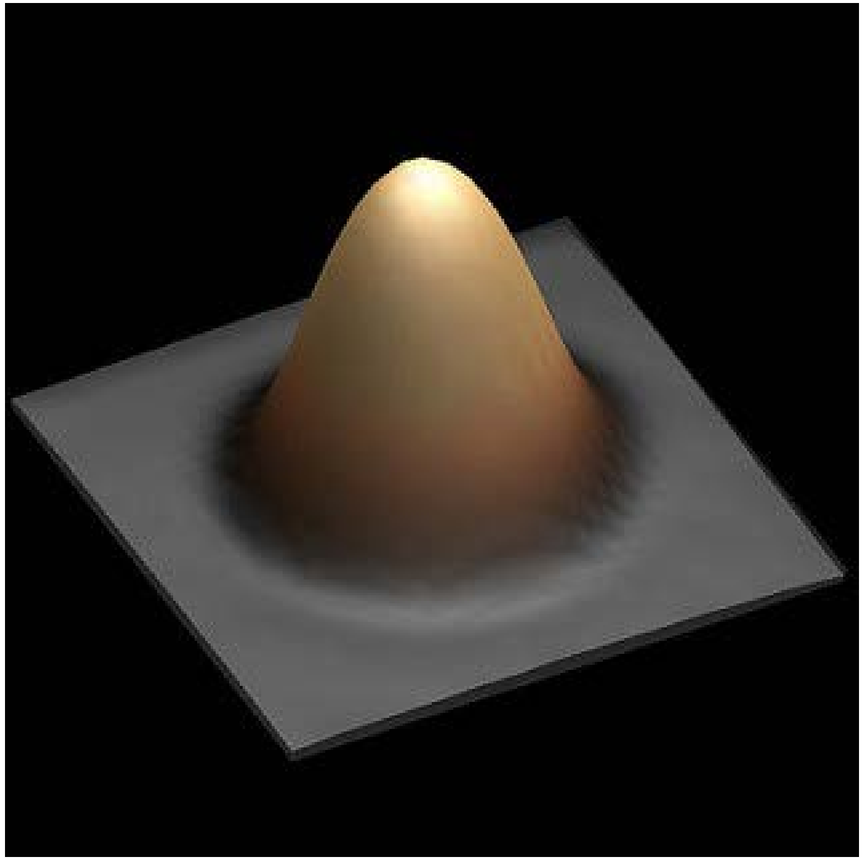

A view from IBM Research's Nobel prize-winning microscope of a single atom of Holmium, an element used as a magnet to store one bit of data. Photo credit: IBM Research - Almaden (San Jose, Calif.)

IBM