ECE566
Enterprise Storage Architecture

Spring 2025

Survey of Next-Generation Storage

Tyler Bletsch
Duke University

Software-Defined Storage

A “new” system architecture for storage

Software-Defined Storage

- "Software-Defined Storage" is part of the current "Software-Defined X" trend in IT
  - Remember we covered "software-defined networking"
- Key properties
  - Pool physical resources
  - Provision logical containers from shared pool
  - Service Level Agreements (SLAs) describe expected performance
- Do these sound familiar?
  - Same properties we covered in cloud!

- "Software-Defined Storage" and "Storage-as-a-Service in a Private Cloud" are basically the same thing

Advances in physical media

Lots of possible avenues...

- Wikipedia "list of emerging technologies" for storage:

Emerging technology	Status	Potentially marginalized technologies	Related articles
Emerging memory technologies	In development	Current memory technologies	T-RAM, memristor, Z-RAM, TTRAM, CBRAM, SONOS, RRAM, Racetrack memory, NRAM, Phase-change memory, FJG RAM, Millipede memory, Skyrmion, Programmable metallization cell, 3D XPoint, Ferroelectric RAM, Magnetoresistive random-access memory, nvSRAM
Emerging magnetic data storage technologies	In development (HAMR, BPM); diffusion (SMR)		SMR, HAMR, BPM, MAMR, TDMR, CPP/GMR, PMR, Hard disk drive

- That's a lot of things! Most won't pan out
- Temper your excitement, remember the hype cycle...

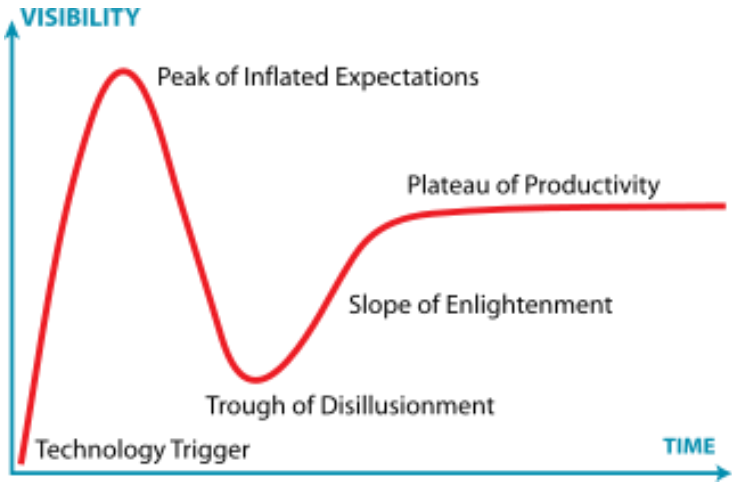

Areas of focus

- Improving HDDs
  - Perpendicular magnetic recording (PMR)
  - Shingled magnetic recording (SMR)
  - Heat-assisted magnetic recording (HAMR) (and friends)
  - Bit-patterned media (BPM)
  - Two-dimensional magnetic recording (TDMR)
  - Just pump a bunch of helium in there
- Improving SSDs
  - 3D NAND structures
- New solid-state memories
  - Phase-change memory (PCM)
  - Ferroelectric RAM (FRAM)
  - Magnetoresistive RAM (MRAM)
  - Resistive RAM (RRAM)
  - Conductive Bridging RAM (CBRAM)
- Memristors: are they a thing?
- Theoretical and proof-of-concept stuff

Improving HDDs

Perpendicular Magnetic Recording (PMR)

"Ring" writing element

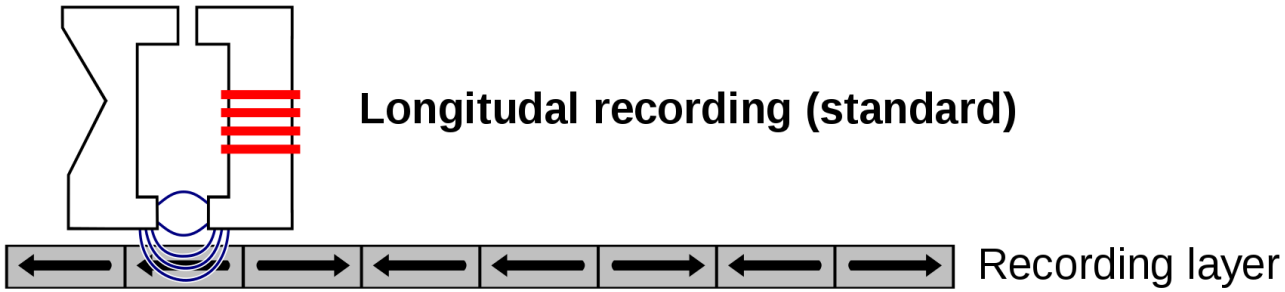

"Monopole" writing element

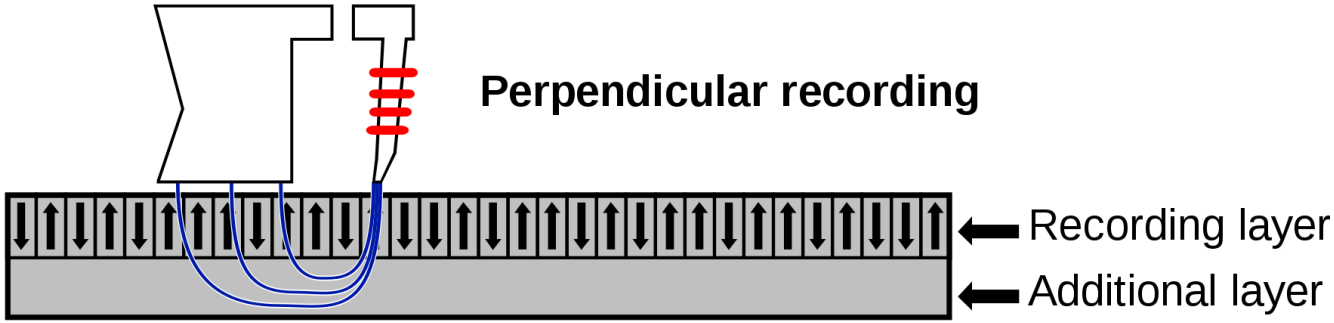

src

Feasible? **Yes.** Most modern large drives do this, deployed since mid-2000s.

Shingled magnetic recording (SMR)

- Due to physics reasons, the write head is always bigger than the read head
  - This means that we write a track of X width, but we just read the middle X/2 of it back.
  - Tracks aren't allowed to overlap, so this leads to waste
- Solution: let them overlap, and deal with resulting destruction

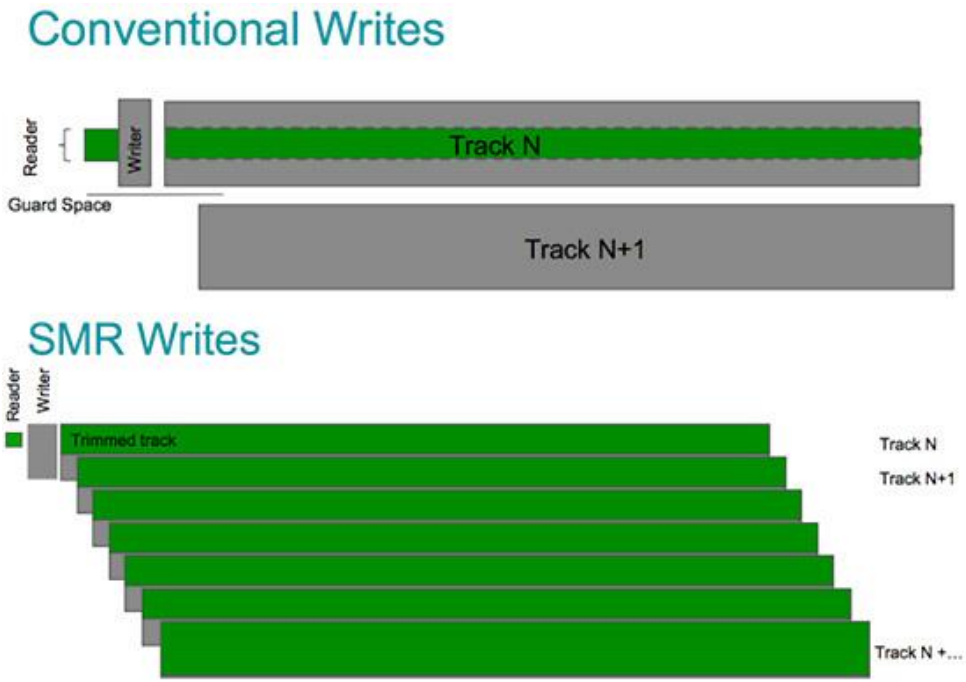

Feasible? Yes. Seagate started shipping in 2013. Typical in "archival" drives.

Diagram source.

Shingled magnetic recording (SMR)

- Dealing with overlap
  - Drive reads neighboring data under threat from a pending write; restores it afterward.
  - If we blindly do that to whole drive, then single write means rewriting whole drive...
  - Solution: Do SMR on track groups.
  - Wow! HDD now like SSD: Small read sectors, big erasure blocks!
  - Lots of cache and optimization opportunities...
- Real-world result in 2020: Western Digital SMR drives had catastrophic performance problems for certain workloads...
  - Made worse because they didn't disclose which were SMR-based drives ☹

Feasible? Yes. Seagate started shipping in 2013. Typical in "archival" drives.

Seal the HDD and fill with helium

- Reduces mechanical power dissipated in air shear
- Allows platters to be placed closer together enabling more capacity

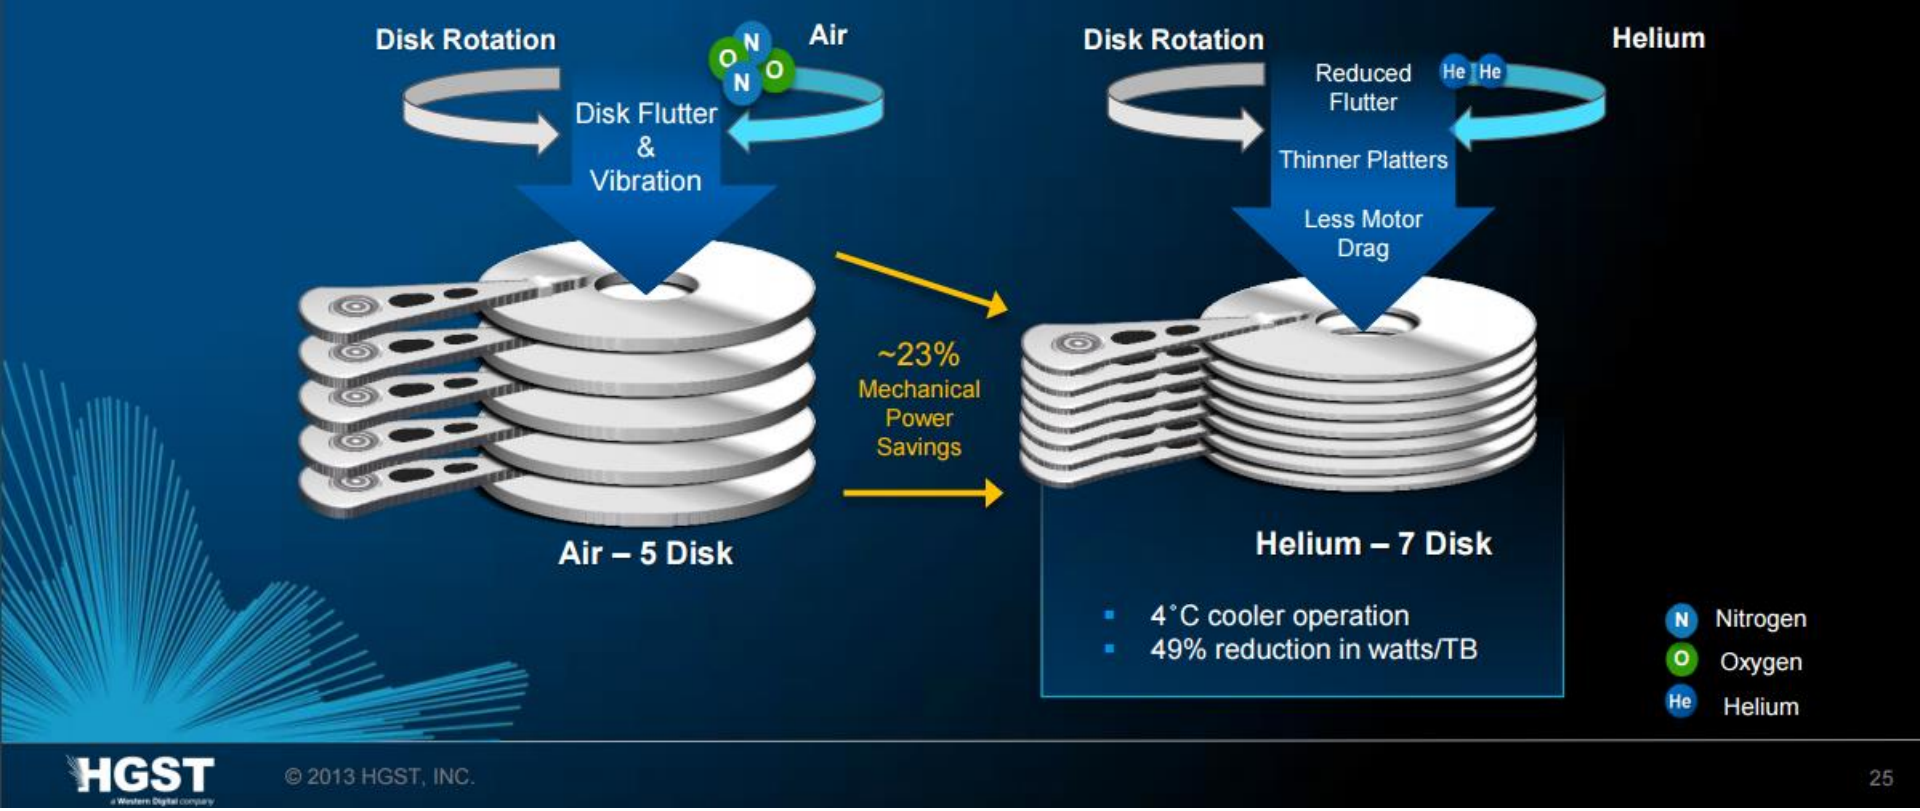

From "Navigating Storage in a Cloudy Environment" by Steve Campbell, HGST.

Feasible? Yes. HGST started shipping in 2013.

Heat-Assisted Magnetic Recording (HAMR)

HAMR : A Whole New Recording System

- Density growth limited by ability to make smaller bits thermally stable
- HAMR combines laser and magnetic field to write the media
- Allows for use of much higher coercivity media and hence enables higher densities

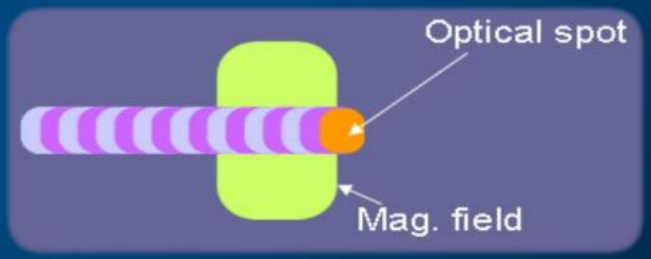

$$\frac{dH_{eff}}{dx} = \frac{dH_k}{dT} \cdot \frac{dT}{dx}$$

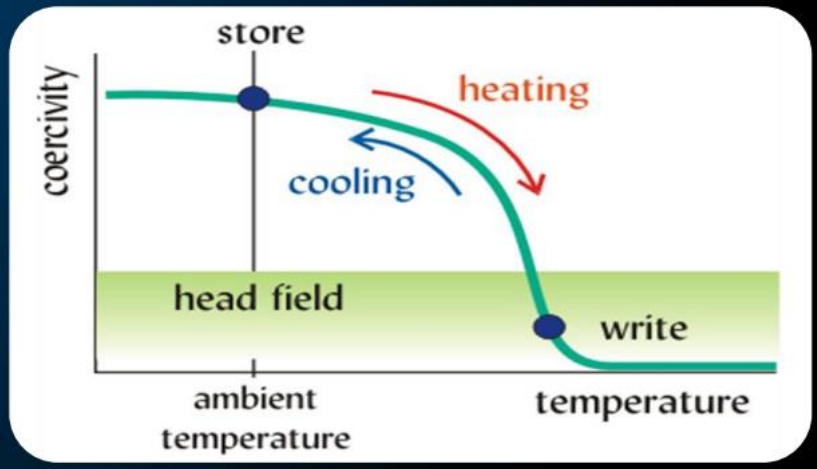

Industry projecting the introduction of HAMR technology in 2016-2017

From "Navigating Storage in a Cloudy Environment" by Steve Campbell, HGST.

See also: Energy- and Microwave-Assisted Magnetic Recording (EAMR/MAMR), used by Western Digital

Feasible? Yes. Seagate shipped a 32TB HAMR drive in June 2023 (src)!

Bit Patterned Media

Granular Media versus Bit Patterned Media

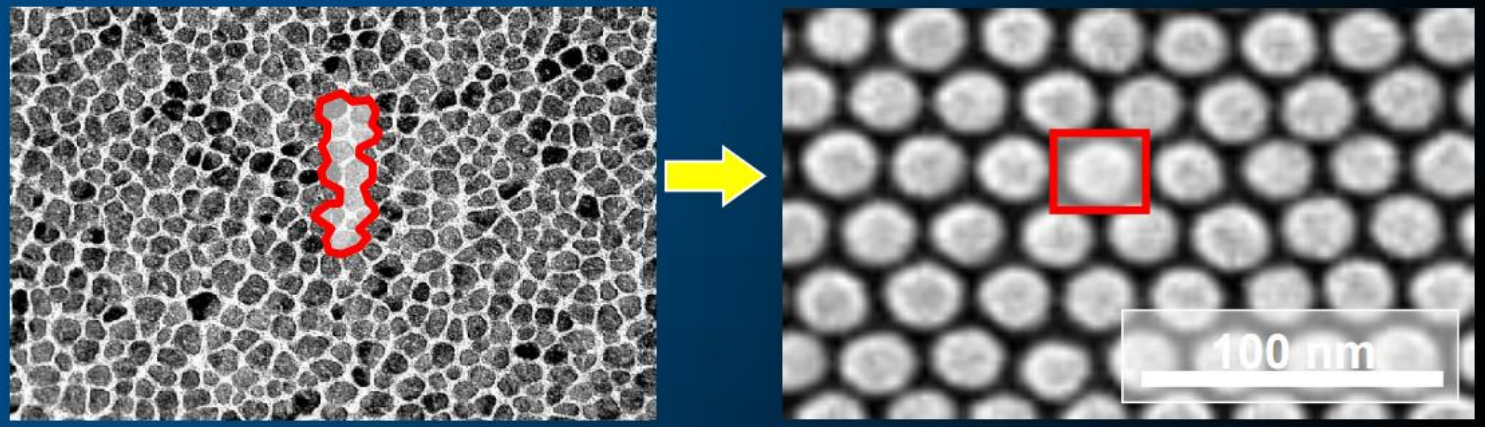

- Extend density by replacing randomly sputtered grains with very uniform, lithographically-defined magnetic islands
- The challenge for bit patterned media is how to fabricate these very small islands precisely and cost-effectively
- Feature sizes will need to be smaller than semiconductor

Have already demonstrated all the steps necessary for 13nm half pitch

From "Navigating Storage in a Cloudy Environment" by Steve Campbell, HGST.

Feasible? Somewhat likely. HGST has proved the lithography, but there are lots of problems still left (src). No drives as of 2025.

Longer term

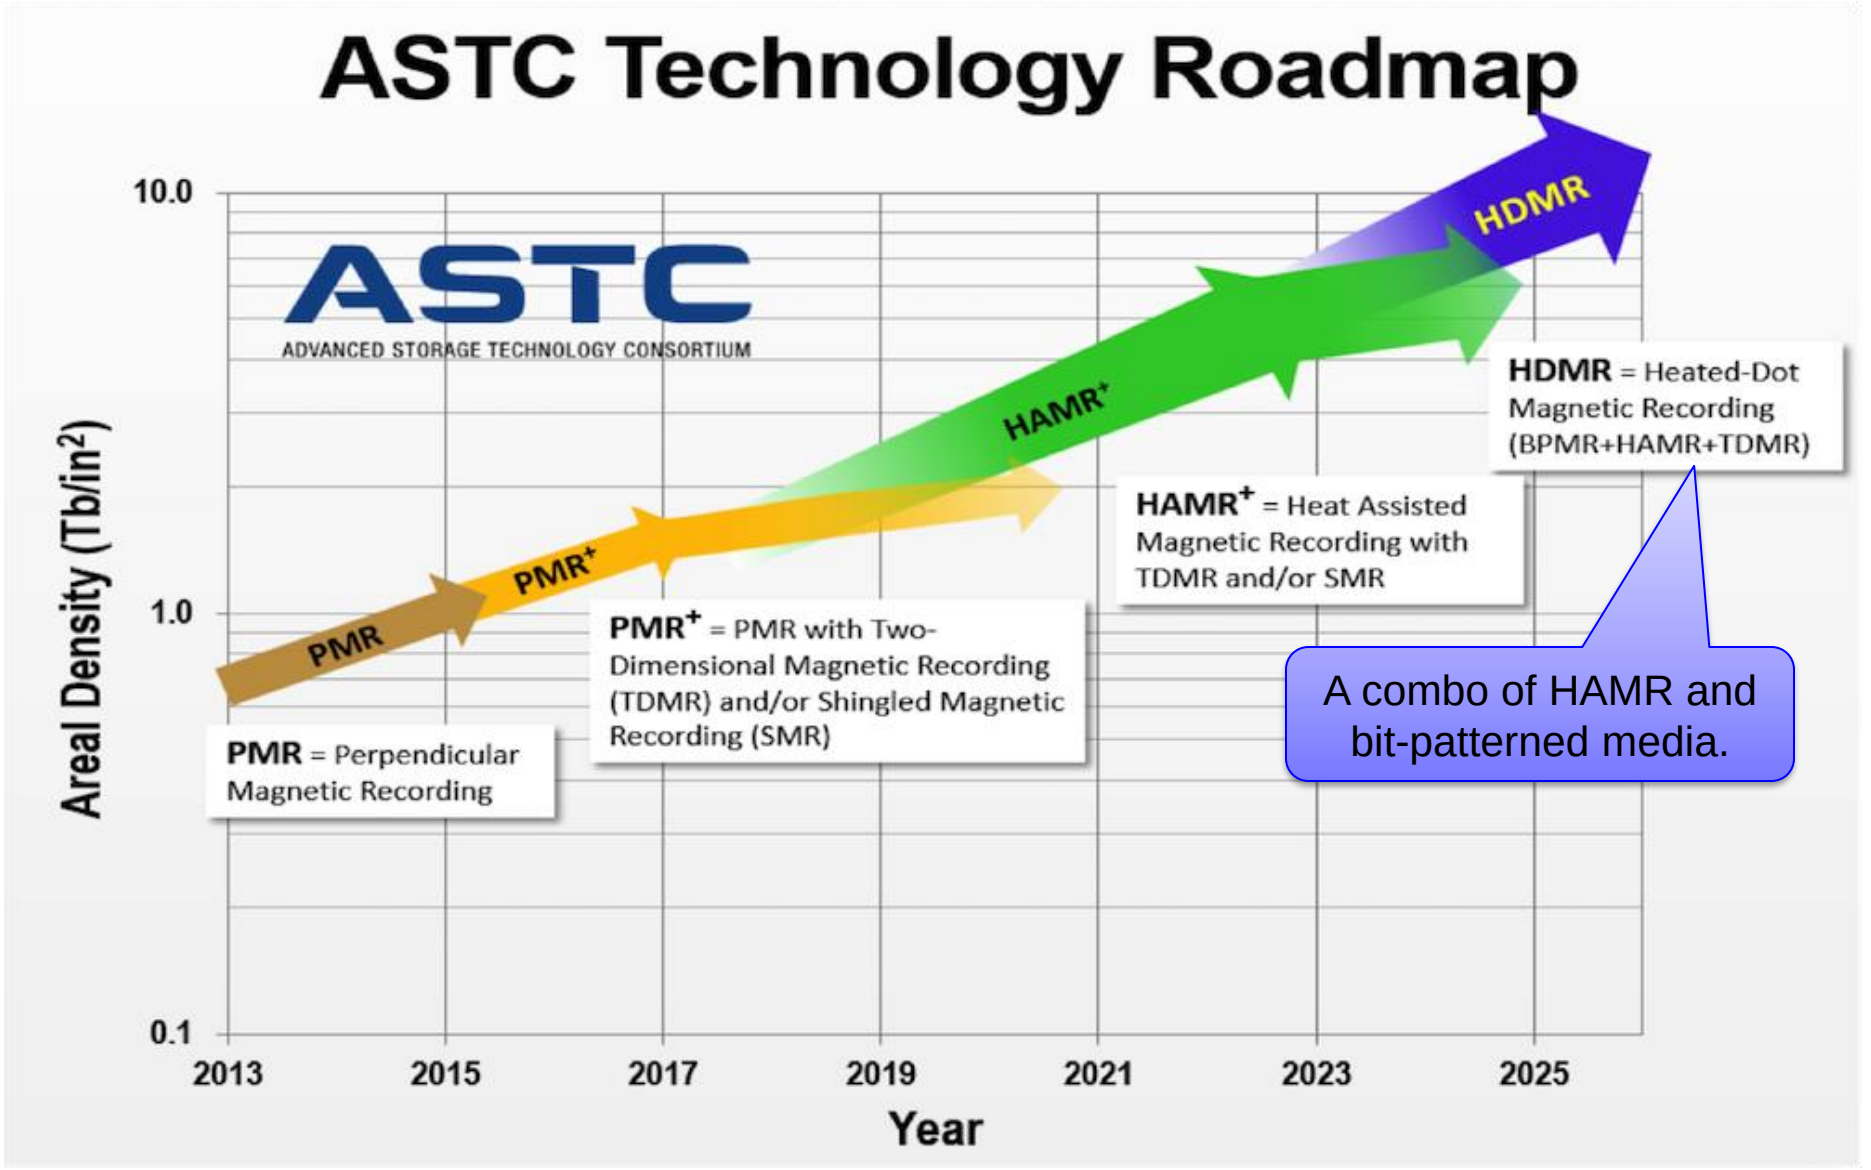
ASTC Technology Roadmap
ASTC
ADVANCED STORAGE TECHNOLOGY CONSORTIUM
Areal Density (Tb/in²)
10.0
1.0
0.1
2013
2015
2017
2019
2021
2023
2025
Year
PMR
PMR+
HAMR+
HDMR
PMR = Perpendicular Magnetic Recording
PMR+ = PMR with Two-Dimensional Magnetic Recording (TDMR) and/or Shingled Magnetic Recording (SMR)
HAMR+ = Heat Assisted Magnetic Recording with TDMR and/or SMR
HDMR = Heated-Dot Magnetic Recording (BPMR+HAMR+TDMR)
A combo of HAMR and bit-patterned media.

Improving SSDs

3D NAND structures

- Current SSD/flash design: NAND gates laid out in 2D

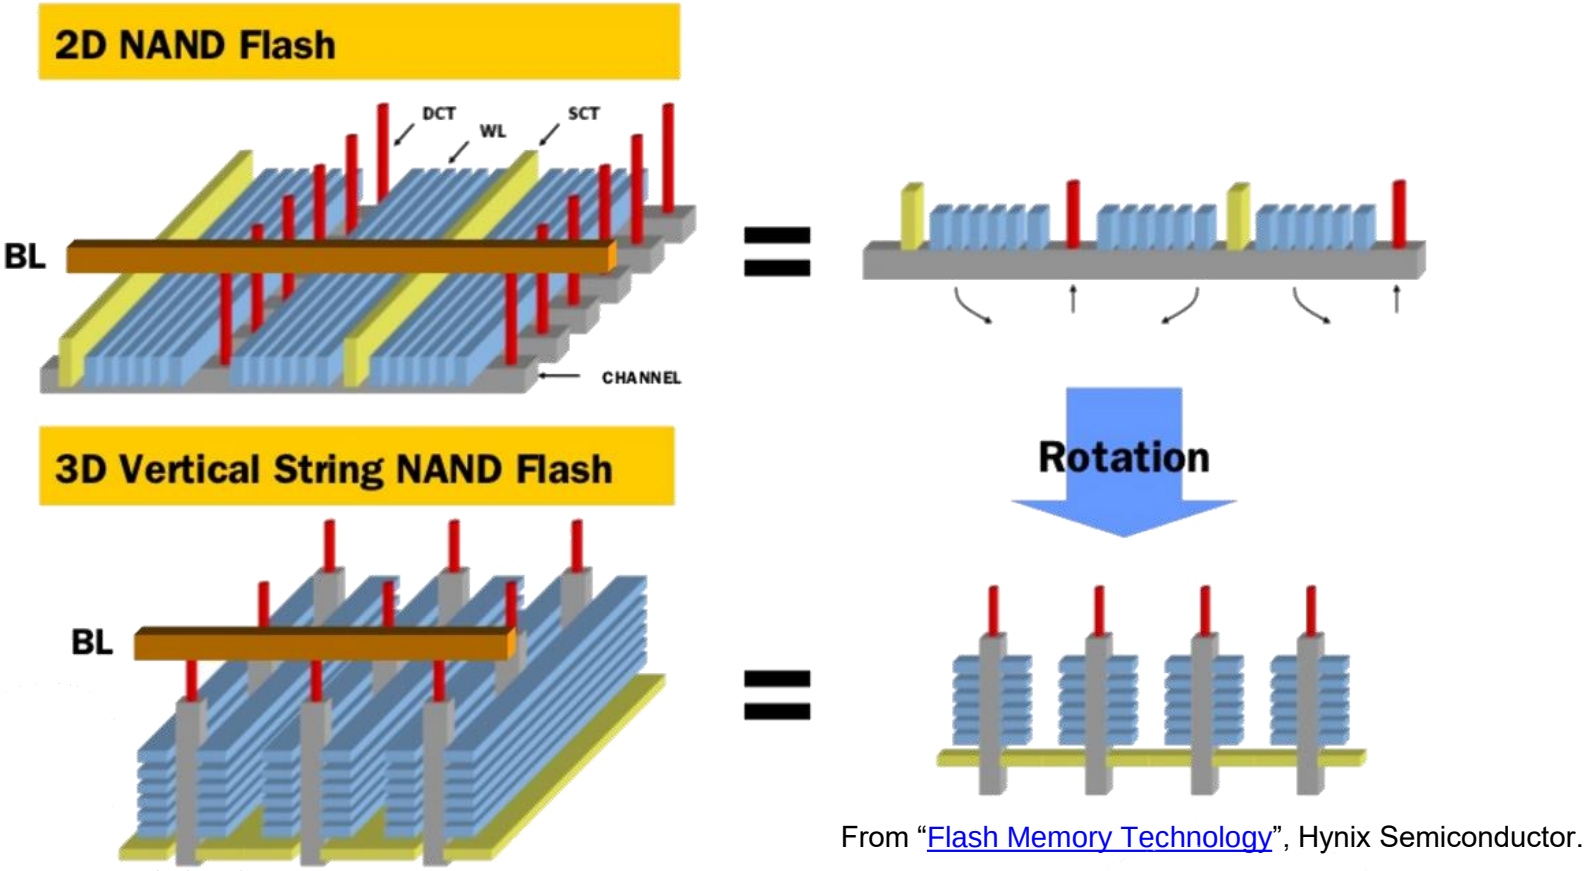

From "Flash Memory Technology", Hynix Semiconductor.

- Novel idea: Make it 3D. Lots of ways to do this...

Feasible? Yes. Intel/Micron have chips shipping, over 200 layers is common. (src)

3D NAND structures

- Lots of ways to do this...

	Toshiba/P-BICS	Hynix DC-SF	Samsung/TCAT
Type of 3D NAND	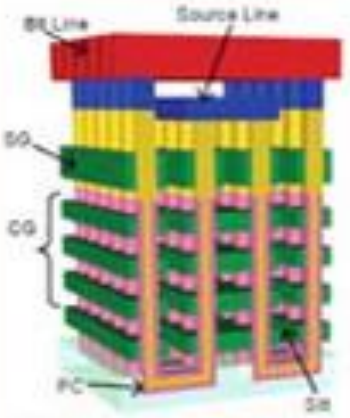 R. Katsumata et al, VLSI 2009	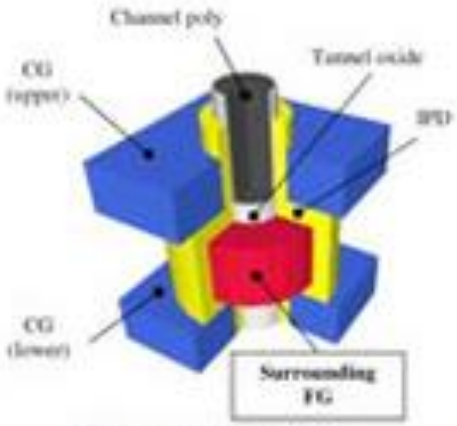 S Whang et al, IEDM 2010	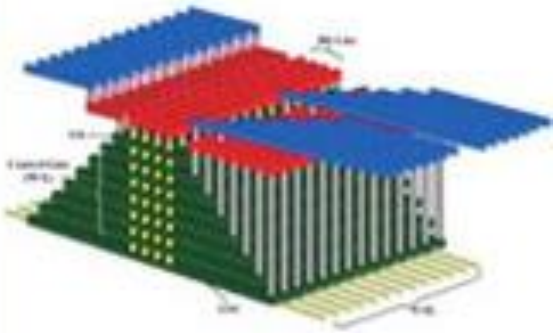 J. Jang et al, VLSI 2009
Transistor	Gate all around; Salicided Poly Si gate	Gate all around; Salicided Poly Si gate	Gate all around; Damascene metal gate
Storage	Charge trap	Floating gate	Charge trap

From “3D NAND Approaches”, IMW 2011. Figure from here.

Feasible? Yes. Intel/Micron have chips shipping, over 200 layers is common. (src)

New solid state memories

Phase-change memory (PCM)

- Fundamental enabler: Chalcogenide glass
  - A glass compound with sulfur, selenium, or other additive
  - Rate of heating/cooling can produce **amorphous** or **crystalline** structure
  - Two structures behave very differently optically and electrically

- This is what makes re-writable CD/DVDs possible
- To "write":
  - Melt with brief, hot pulse of heat; rapid cooling gives amorphous state
  - Melt with long, low-intensity heat; slow cooling gives crystalline state
- To "read":
  - Crystalline is low resistance, amorphous is high resistance
  - Measure resistance with circuit, decide which one means "1"

Feasible? Technically, yes (limited deployment); economically, maybe…

Phase-change memory (PCM)

- Array these elements in a grid like any other RAM
- Use electricity to heat cells (write) and to determine their resistance (read)

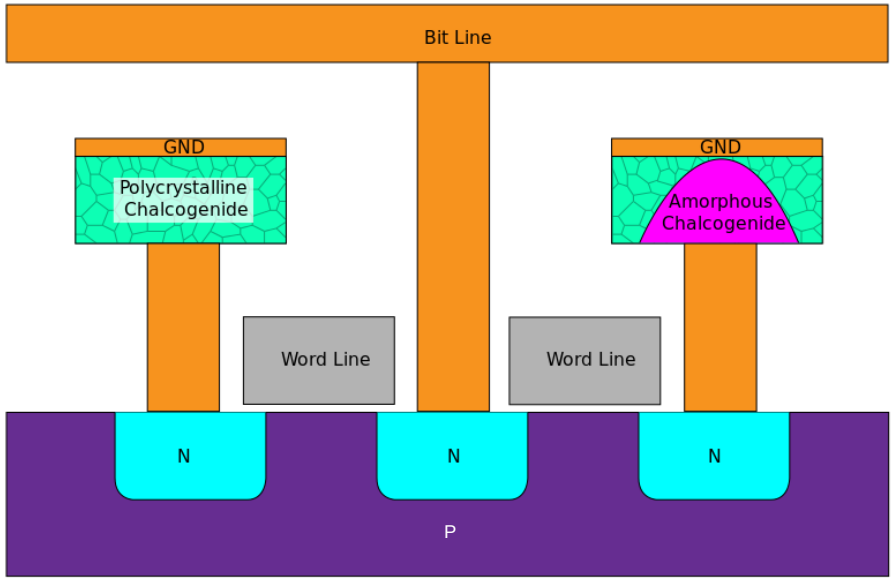

"A cross-section of two PRAM memory cells. One cell is in low resistance crystalline state, the other in high resistance amorphous state."
From Wikipedia, "Phase-change memory"

Feasible? Technically, yes; economically, maybe.

- Shipping memory chips available from many vendors
- Large-scale adoption hasn't happened; flash still wins for most use cases when you factor in cost
- Roller-coaster development history:
  - In 2012, Micron announced PCM for mobile devices (src)
  - In 2014, flash had gotten better (e.g. 3D NAND), and Micron ditched PCM! (src)
  - In 2015, PCM appeared dead, but then Western Digital showed a PCM prototype with 3 million IOPS (src)
  - Intel/Micron's "3D Xpoint memory" is a PCM released in 2016 (src)
  - STMicroelectronics demoed a 28nm device in 2018.
  - Not much lately…

Ferroelectric RAM (FRAM)

- Like DRAM, but uses a "**ferroelectric**" layer instead of the DRAM capacitors' dielectric.
- Ferroelectric material: Material that has an electric polarization which can be flipped
  - Material consists of polarized molecules (one side positive, other side negative)
  - If you flip one molecule, attraction/repulsion resets it
    - Stable, self-correcting
  - Apply enough voltage, flip *all* molecules
    - Settable!

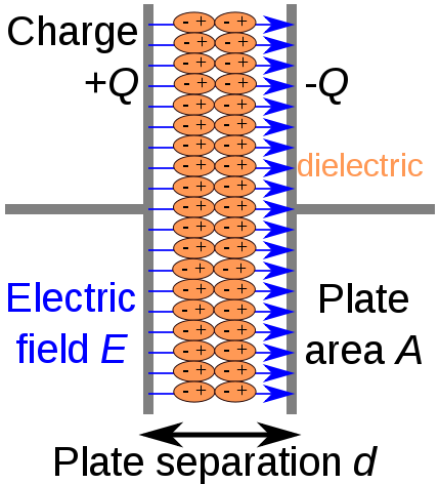

From Wikipedia, "Ferroelectric capacitor"

Feasible? Technically, yes; economically, maybe.

- Shipping memory chips available from vendors
- Large-scale adoption hasn't happened; seems unlikely under current trends
- Density isn't great (130nm), but lower power than flash
- Current niche: storage for very-low-power embedded systems

Magnetoresistive RAM (MRAM)

- Uses a "**ferromagnetic**" material
  - Metal that can change magnetic field to match an external field (e.g., normal iron)
- Exploits "**tunnel magnetoresistance**"
  - Due to wacky probabilistic quantum physics, an electron in the top layer can "tunnel" (randomly transposition to) the bottom layer
  - If both magnets have same polarity, this tunneling is much more likely (src)
  - Macroscopic effect: resistance is lower
- Can flip magnetic polarity with electrically-created field (write), determine polarity by measuring resistance (read)

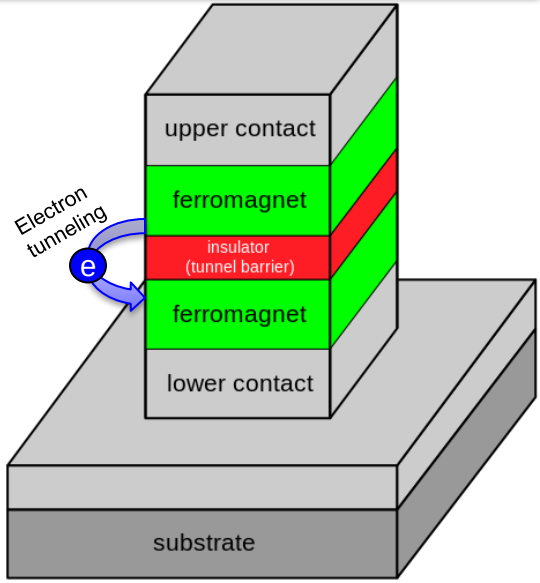

From Wikipedia, "Tunnel magnetoresistance"

Feasible? Technically, yes; economically, maybe.

- Only one shipping commercial part (a 4Mbit chip from Everspin)
- Large-scale adoption hasn't happened; seems uncertain
- Density is lousy (180nm), but great performance and lower power than FRAM
- Current niche: storage for very-low-power embedded systems
- A start-up has announced a microcontroller that includes MRAM (src)
- Other companies are developing MRAM manufacturing capacity (src)

Others

- Conductive bridging RAM (CBRAM): Electrochemical reaction changes resistivity of cells.
  - Development startup Adesto holds the intellectual property, limited products have been realized. One company wants to use it in space.

Feasible? Technically, yes; economically, unlikely?

- Resistive RAM (RRAM): Create/fill electron "vacancies" in a thin oxide layer; changes resistivity of cells.
  - Various small commercial chips exist in the kB range. Adesto's here too.

Feasible? Technically, yes; economically, unlikely?

- "Millipede memory": Create and fill microscopic holes in a thin polymer.
  - In 2005, IBM was aiming to have this out within 2 years, but other forms of storage advanced faster and wrecked it

Feasible? Technically, ???; economically, dead.

Memristors: are they a thing?

- Memristor: A theoretical circuit element that changes resistance based on past current
  - Existence was proposed by taxonomy in 1971:
    "If we have components that relate charge, voltage, current, and magnetic flux, shouldn't this thingy exist"? (src)

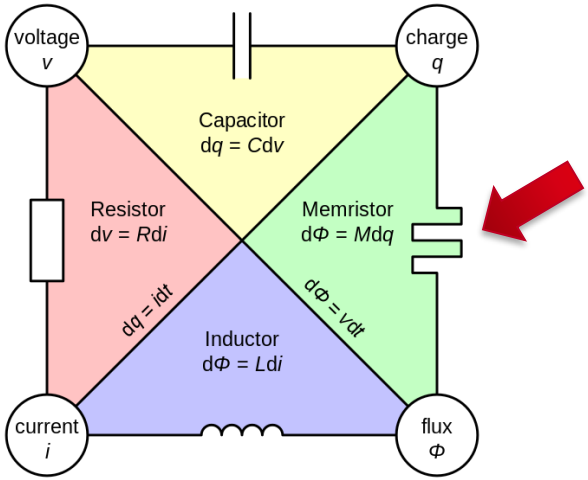

From Wikipedia, "Memristor"

  - By 2011 we didn't a good one, but we liked the name, so it changed to:
    "Any 2-terminal thing that changes resistance" (src)

Memristors: are they a thing?

- Problem: We just changed the definition so that it matches most of the proposed non-volatile RAMs we've discussed!
- Result: LOTS OF CONFUSION.
  - Technology press: ***"Memristors are the next big thing!"***
  - Actual semiconductor engineers working on this: "wtf are you talking about?"
- My opinion: "memristor" isn't a useful concept. Either:
  - It doesn't exist (original definition), or
  - It is achieved through a dozen different unrelated physical processes (new definition).
- The following shows that it's not a real thing:

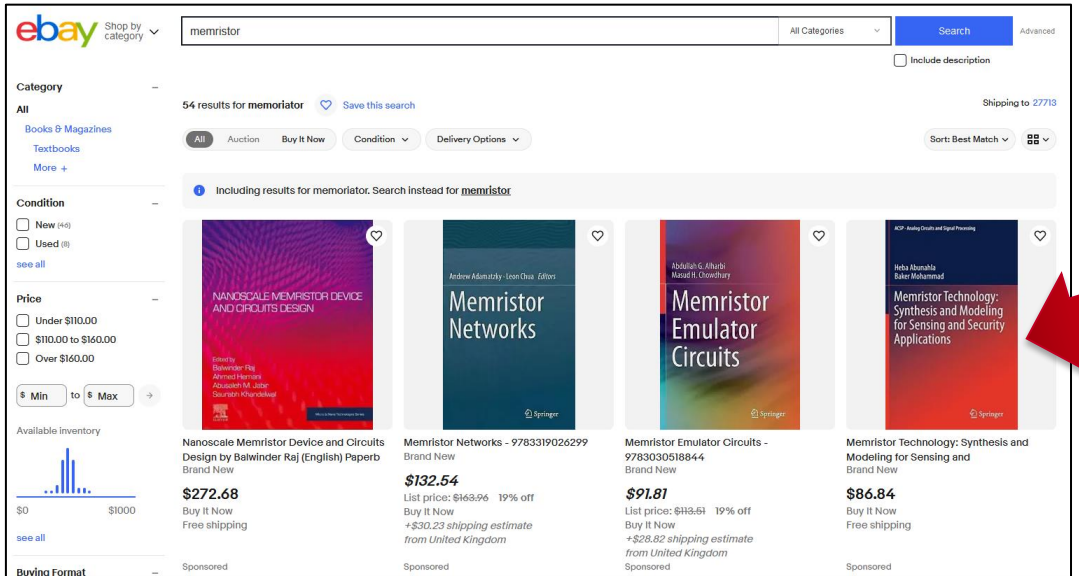

Nothing but a few journals; no actual components to buy!

Speculative future stuff

AKA “A list of things that almost never pan out, except when they do”

Cold data storage in glass

- "Project Silica":
  - Use Lasik-eye-surgery style lasers to etch shapes into glass
  - Used to store "cold data" (e.g. movie prints)
  - 75GB in one 75mm square disc

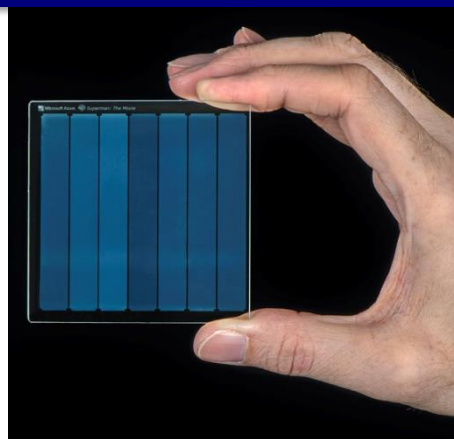

- Experimental stages as of 2025

- Sponsored by Microsoft Research and Warner Bros Studios

Source

Other “3D” or “Holographic” optical storage

- Various attempts to store data in depth of optical media
- Focus is often cold data archival (like Project Silica)
- Not much news since 2010…

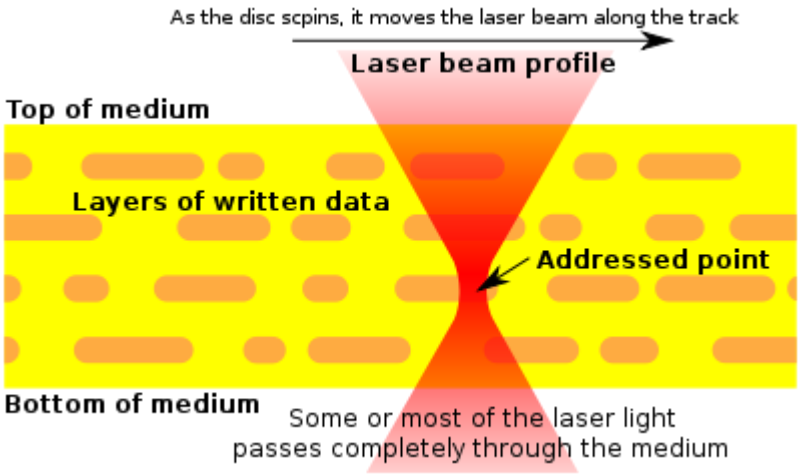

src

Feasible? Technically, yes; economically, looking less likely…

Theoretical and proof-of-concept stuff

- Spintronics: Trying to do stuff with the quantum "spin" of electrons
  - Partially commercialized

- "Nano-RAM": Storing data based on position of carbon nanotubes on a chip substrate
  - IBM, MIT, others have lab demos, but nothing commercially viable

- Skyrmion: A hypothetical quantum particle related to magnetism
  - This is a literal sentence used to describe these:
    "A two-dimensional magnetic skyrmion, as a topological object, is formed, e.g., from a 3D effective-spin "hedgehog" (in the field of micromagnetics: out of a so-called "Bloch point" singularity of homotopy degree +1) by a stereographic projection, whereby the positive north-pole spin is mapped onto a far-off edge circle of a 2D-disk, while the negative south-pole spin is mapped onto the center of the disk."
  - If that makes sense to you, invest in Skyrmion companies I guess?

Any of these feasible?
WHO KNOWS??

Theoretical and proof-of-concept stuff

- DNA: Yeah, we'll just encode data in DNA!
  - Ignore the fact that existing life doesn't do arbitrary write operations on DNA (cells copy, viruses splice, meiosis mixes, and epigenetics alters attached methyl groups, but **nothing** makes arbitrary in-place changes)
  - Ignore that every aspect of DNA's evolution has been focused on protein synthesis, not specific "reads" of DNA locations
  - Ignore that even the fastest and most common IO pattern for DNA, copying during mitosis, takes on the order of hours, giving a data rate of around 60 kBps (similar to a dial-up modem)
  - No, DNA is the definitely the perfect way to store my pirated movies
  - Current research: DNA writing at 1.6 MB/s, no efficient read solution (src)

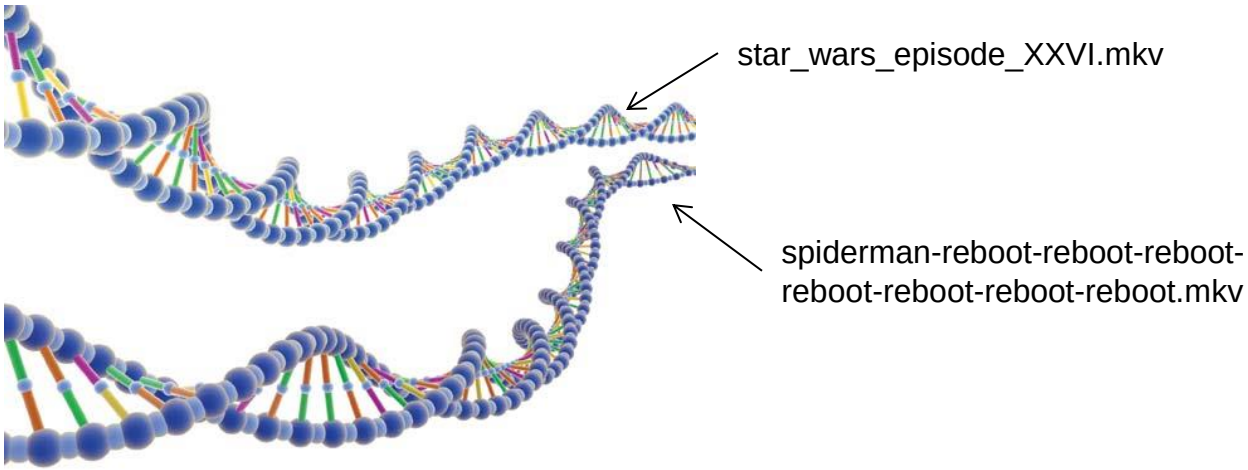

Feasible? LOL