



The Life of a Data Byte

BE KIND AND REWIND

JESSIE FRAZELLE

A byte of data has been stored in a number of different ways through the years as newer, better, and faster storage media are introduced. A byte is a unit of digital information that most commonly refers to eight bits. A bit is a unit of information that can be expressed as 0 or 1, representing a logical state. Let's take a brief walk down memory lane to learn about the origins of bits and bytes.

Going back in time to Babbage's Analytical Engine, you can see that a bit was stored as the position of a mechanical gear or lever. In the case of paper cards, a bit was stored as the presence or absence of a hole in the card at a specific place. For magnetic storage devices, such as tapes and disks, a bit is represented by the polarity of a certain area of the magnetic film. In modern DRAM (dynamic random-access memory), a bit is often represented as two levels of electrical charge stored in a capacitor, a device that stores electrical energy in an electric field. (In the early 1960s, the paper cards used to input programs for IBM mainframes were known as *Hollerith cards*, named after their inventor, Herman Hollerith from the Tabulating Machines Company—which through numerous mergers is what is now known as IBM.)

In June 1956, Werner Buchholz ([archive](#).
[computerhistory.org](#)) coined the word *byte* ([archive](#).
[org](#)) to refer to a group of bits used to encode a single

character of text (bobbemer.com). Let's address character encoding, starting with ASCII (American Standard Code for Information Interchange). ASCII was based on the English alphabet; therefore, every letter, digit, and symbol (a-z, A-Z, 0-9, +, -, /, ", !, etc.) was represented as a seven-bit integer between 32 and 127. This wasn't very friendly to other languages. To support other languages, Unicode extended ASCII so that each character is represented as a code-point. For example, a lowercase `j` is `U+006A`, where `U` stands for Unicode followed by a hexadecimal number.

UTF-8 is the standard for representing characters as eight bits, allowing every code-point from 0 to 127 to be stored in a single byte. This is fine for English characters, but other languages often have characters that are expressed as two or more bytes. UTF-16 is the standard for representing characters as 16 bits, and UTF-32 is the standard for 32 bits. In ASCII every character is a byte, but in Unicode, that's not true—a character can be one, two, three, or more bytes. Groups of characters might also be referred to as *words*, as in this linked Univac ad (archive.org) calling out “1 kiloword or 12,000 characters.” This article refers throughout to different-sized groupings of bits—the number of bits in a byte varying according to the design of the storage medium in the past.

This article also travels in time through various storage media, diving into how data has been stored throughout history. By no means does this include every single storage medium ever manufactured, sold, or distributed. This article is meant to be fun and informative but not encyclopedic. It wraps up with a look at the current and future technologies for storage.

At this point in history, you could measure the speed of a storage algorithm by the distance the tape traveled.

To get started, let's assume we have a byte of data to be stored: the letter `j`, or as an encoded byte `6a`, or in binary `01001010`. As we travel through time, this data byte will come into play in some of the storage technologies covered here.

1951

The story begins in 1951 with the Uniservo tape drive for the Univac 1 computer, the first tape drive made for a commercial computer. The tape was three pounds of a thin strip (half-inch) of nickel-plated phosphor bronze, called vicalloy, which was 1,200 feet long. Our data byte could be stored at a rate of 7,200 characters per second (computerhistory.org) on tape moving at 100 inches per second. At this point in history, you could measure the speed of a storage algorithm by the distance the tape traveled.

1952

Let's fast forward a year to May 21, 1952, when IBM announced its first magnetic tape unit, the IBM 726. Our data byte could now be moved off Uniservo metal tape onto IBM's magnetic tape. This new home would be roomy for our very small data byte since the tape could store up to 2 million digits. This magnetic seven-track tape moved at 75 inches per second with a transfer rate of 12,500 digits (ibm.com) or 7,500 characters (ibm.com)—called *copy groups* at the time—per second. For reference, this article has 35,123 characters.

Seven-track tapes had six tracks for data and one to

maintain parity by ensuring that the total number of 1-bits in the string was even or odd. Data was recorded at 100 bits per linear inch. This system used a *vacuum channel* method of keeping a loop of tape circulating between two points. This allowed the tape drive to start and stop the tape in a split second. This was done by placing long vacuum columns between the tape reels and the read/write heads to absorb sudden increases in tension in the tape, without which the tape would have typically broken. A removable plastic ring on the back of the tape reel provided write protection. About 1.1 MB could be stored on one reel of tape [spectrum.ieee.org].

Think back to VHS tapes: What was required before returning a movie to Blockbuster? Rewinding the tape! The same could be said for the tape used for computers. Programs could not hop around a tape, or randomly access data—they had to read and write in sequential order.

1956

The era of magnetic-disk storage began in 1956 with IBM's completion of the 305 RAMAC computer delivered to Zellerbach Paper in San Francisco [ibm.com]. This computer was the first to use a moving-head HDD (hard disk drive). The RAMAC disk drive consisted of 50 magnetically coated 24-inch-diameter metal platters capable of storing about 5 million characters of data, at seven bits per character, and spinning at 1,200 revolutions per minute. The storage capacity was about 3.75 MB.

RAMAC allowed realtime random access memory to large amounts of data, unlike magnetic tape or punch cards. IBM advertised the RAMAC as being able to store

the equivalent of 64,000 punched cards (youtube.com). Previously, transactions were held until a group of data was accumulated and batch processed. The RAMAC introduced the concept of continuously processing transactions as they occurred so data could be retrieved immediately when it was fresh. Our data byte could now be accessed in the RAMAC at 100,000 bits per second (ibm.com). Prior to this, with tapes, people had to write and read sequential data and could not randomly jump to various parts of the tape. Realtime random access of data was truly revolutionary at this time.

1963

DECtape was introduced in 1963. Its namesake was the Digital Equipment Corporation, known as DEC for short. DECtape was inexpensive and reliable, so it was used in many generations of DEC computers. This $\frac{3}{4}$ -inch tape was laminated and sandwiched between two layers of Mylar on a four-inch reel.

DECtape could be carried by hand, as opposed to its large weighty predecessors, making it great for personal computers. In contrast to seven-track tape, DECtape had six data tracks, two mark tracks, and two clock tracks. Data was recorded at 350 bits per inch. Our data byte, which is eight bits but could be expanded to 12, could be transferred to DECtape at 8,325 12-bit words per second with a tape speed of 93 +/-12 inches per second (pdp8.net). This is eight percent more digits per second than the Uniservo metal tape from 1952.

1967


Four years later, in 1967, a small team at IBM started working on the IBM floppy-disk drive, code-named Minnow (ibm.com). At the time, the team was tasked with developing a reliable and inexpensive way to load microcode into IBM System/370 mainframes (archive.org). The project then got reassigned and repurposed to load microcode into the controller for the IBM 3330 Direct Access Storage Facility, code-named Merlin.

Our data byte could now be stored on read-only eight-inch flexible Mylar disks coated with magnetic material, known as *floppy disks*. At the time of release, the result of the project was named the IBM 23FD Floppy Disk Drive System. The disks could hold 80 KB of data. Unlike hard drives, a user could easily transfer a floppy in its protective jacket from one drive to another. Later, in 1973, IBM released a read/write floppy-disk drive, which then became an industry standard (web.archive.org).

1969

In 1969, the AGC (Apollo Guidance Computer) read-only rope memory was launched into space aboard Apollo 11, which carried American astronauts to the moon and back. This rope memory was made by hand and could hold 72KB of data. Manufacturing rope memory was laborious, slow, and required skills analogous to textile work; it could take months to weave a program into the rope memory (authors.library.caltech.edu), shown in figure 1. But it was the right tool for the job at the time to resist the harsh rigors of space. When a wire went through one of the circular cores, it represented a 1. Wires that went around a

FIGURE 1: ROPE MEMORY



core represented a 0. Our data byte would take a human a few minutes to weave into the rope.

1977

The Commodore PET, the first (successful) mass-market personal computer, was released in 1977. Built into the PET was a Commodore 1530 Datasette (a portmanteau of *data* plus *cassette*). The PET converted data into analog sound signals that were then stored on cassettes (wavprg.sourceforge.net). This made for a cost-effective and reliable storage solution, albeit very slow. Our small data byte could be transferred at a rate of around 60-70 bytes per second (c64-wiki.com). The cassettes could hold about 100 KB per 30-minute side, with two sides per tape. For example, you could fit about two (“rickroll” warning) 55KB images (blog.jessfraz.com) on one side of the cassette. The Datasette also appeared in the Commodore VIC-20 and Commodore 64.

1978

Let’s jump ahead a year to 1978 when the LaserDisc was introduced as Discovision by MCA and Philips. *Jaws* was the first film sold on a LaserDisc in North America. The audio and video quality on a LaserDisc were far better than the competitors’, but too expensive for most consumers. As opposed to VHS tape, which consumers could use to record TV programs, the LaserDisc could not be written to. LaserDiscs used analog video with analog FM stereo sound and PCM (pulse-code modulation) digital audio (tools.ietf.org). The disks were 12 inches in diameter and composed of two single-sided aluminum disks layered in plastic. The

LaserDisc is remembered today as being the foundation CDs that DVDs were built upon.

1979

A year later, in 1979, Alan Shugart and Finis Conner founded Seagate Technology with the idea of scaling down a hard disk drive to be the same size as a 5¼-inch floppy disk, which at the time was the standard. Their first product, in 1980, was the Seagate ST506, the first HDD for microcomputers. The 5¼-inch disk held 5 MB of data, which at the time was five times more than the standard floppy disk. It was a rigid, metallic platter coated on both sides with a thin layer of magnetic material to store data. Our data byte could be transferred at a speed of 625 KB/s (pcmag.com) onto the disk. That's about one (second and final "rickroll" warning) 625KB animated GIF (blog.jessfraz.com) per second.

1981

A couple of years later Sony introduced the first 3½-inch floppy drive. Hewlett-Packard was the first adopter of the technology in 1982 with its HP-150. This put the 3½-inch floppy disk on the map and gave it wide distribution in the industry (jstor.org). The disks were single-sided with a formatted capacity of 161.2 KB and an unformatted capacity of 218.8 KB. In 1982 the double-sided version was made available, and the Microfloppy Industry Committee, a consortium of 23 media companies, based a spec for a 3½-inch floppy on Sony's original designs, cementing the format into history (americanradiohistory.com). Our data byte could now be stored on the early version of one of the most widely

distributed storage media: the 3½-inch floppy disk.

1984

In 1984 the CD-ROM (compact disk read-only memory), holding 550 MB of prerecorded data, was announced by Sony and Philips. This format grew out of CD-DA (compact disk digital audio), developed by the two companies in 1982. The CD-DA, which was used for distributing music, had a capacity of 74 minutes. When Sony and Philips were negotiating the standard for a CD-DA, legend has it that one of the four people involved insisted it be able to hold all of Beethoven's *Ninth Symphony* (wired.com). The first product released on CD-ROM was Grolier's Electronic Encyclopedia, which came out in 1985. The encyclopedia contained 9 million words, occupying only 12 percent of the available space, which was 553 mebibytes (books.google.co.uk). There would be more than enough room for the encyclopedia and our data byte. Shortly thereafter in 1985, computer and electronics companies worked together to create a standard for the disks so any computer would be able to access the information.

1984

Also in 1984, Fujio Masuoka published his work on a new type of floating-gate memory, called *flash memory*, which was capable of being erased and reprogrammed multiple times.

Let's first review how floating-gate memory works. It uses transistors, which are electrical gates that can be switched on and off individually. Since each transistor can be in two distinct states (on or off), it can store two

different numbers: 0 and 1. *Floating gate* refers to the second gate added to the middle transistor. This second gate is insulated by a thin oxide layer. These transistors use a small voltage applied to the gate of the transistor to denote whether it is on or off, which in turn translates to a 0 or 1.

With a floating gate, when a suitable voltage is applied across the oxide layer, the electrons tunnel through it and get stuck on the floating gate. Therefore, even if the power is disconnected, the electrons remain present on the floating gate. When no electrons are on the floating gate, it represents a 1; and when electrons are trapped on the floating gate, it represents a 0. Reversing this process and applying a suitable voltage across the oxide layer in the opposite direction causes the electrons to tunnel off the floating gate and restores the transistor to its original state. Thus, the cells are made programmable and nonvolatile (economist.com). Our data byte could be programmed into the transistors as `01001010`, with electrons trapped in the floating gates to represent the zeros.

Masuoka's design was a bit more affordable but less flexible than EEPROM (electrically erasable programmable read-only memory), since it required multiple groups of cells to be erased together, but this also accounted for its speed. At the time, Masuoka was working for Toshiba but quit the company shortly after to become a professor at Tohoku University. He was displeased with Toshiba for not rewarding him for his work and sued the company, demanding compensation for his work. The case settled in 2006 with a one-time payment of ¥87m,

equivalent to \$758,000. This still seems low, given how impactful flash memory has been on the industry.

While on the topic of flash memory, let's look at the difference between NOR and NAND flash. Flash stores information in memory cells made up of floating-gate transistors. The names of the technologies are tied directly to the way the memory cells are organized.

In NOR flash, individual memory cells are connected in parallel, allowing random access. This architecture enables the short read times required for the random access of microprocessor instructions. NOR flash is ideal for lower-density applications that are mostly read-only. This is why most CPUs typically load their firmware from NOR flash. Masuoka and colleagues presented the invention of NOR flash in 1984 and NAND flash in 1987 (ieeexplore.ieee.org).

In contrast, NAND flash designers gave up the ability for random access in a tradeoff to gain a smaller memory cell size. This also has the benefits of a smaller chip size and lower cost-per-bit. NAND flash's architecture consists of an array of eight memory transistors connected in a series. This leads to high storage density, smaller memory-cell size, and faster write and erase since it can program blocks of data at a time. This comes at the cost of having to overwrite data when it is not sequentially written and data already exists in a block (aturing.umcs.maine.edu).

1991

Let's jump ahead to 1991, when a prototype SSD (solid-state disk) module was made for evaluation by IBM from SanDisk, at the time known as SunDisk (computerhistory.org). This

SunDisk was 20MB in a 2½-inch form factor and sold for around \$1,000.

design combined a flash storage array and nonvolatile memory chips with an intelligent controller to detect and correct defective cells automatically. The disk was 20 MB in a 2½-inch form factor and sold for around \$1,000 [meseec.ce.rit.edu]. IBM wound up using it in the ThinkPad pen computer [westerndigital.com].

1994

In 1994 Iomega released the Zip disk, a 100MB cartridge in a 3½-inch form factor, a bit thicker than a standard 3½-inch disk. Later versions of the Zip disk could store up to 2 GB. These disks had the convenience of being as small as a floppy disk but with the ability to hold a larger amount of data, which made them compelling. Our data byte could be written onto a Zip disk at 1.4 MB/s. At the time, a 1.44MB 3½-inch floppy would write at about 16 kB/s. In a Zip drive, heads are noncontact read/write and fly above the surface, which is similar to a hard drive but unlike other floppies. Because of reliability problems and the affordability of CDs, Zip disks eventually became obsolete.

1994

Also in 1994, SanDisk introduced CompactFlash, which was widely adopted into consumer devices such as digital and video cameras. Like CD-ROMs, CompactFlash speed is based on x-ratings (8x, 20x, 133x, etc.). The maximum transfer rate is calculated based on the original audio CD transfer rate of 150 kB/s. This winds up looking like $R = K \times 150 \text{ kB/s}$, where R is the transfer rate and K is the speed rating. For 133x CompactFlash, our data byte would be written at $133 \times 150 \text{ kB/s}$, or around 19,950 kB/s or

19.95 MB/s. The CompactFlash Association was founded in 1995 to create an industry standard for flash-based memory cards.

1997

A few years later in 1997 the CD-RW (compact disc rewritable) was introduced. These optical discs were used for data storage, as well as for backing up and transferring files to various devices. CD-RWs can be rewritten only about 1,000 times, which at the time was not a limiting factor since users rarely overwrote data on one disk.

CD-RWs are based on phase-change technology. During a phase change of a given medium, certain properties of the medium change. In the case of CD-RWs, phase shifts in a special compound, composed of silver, tellurium, and indium, cause *reflecting lands* and *non-reflecting bumps*, each representing a 0 or 1. When the compound is in a crystalline state, it is translucent, which indicates a 1. When the compound is melted into an amorphous state, it becomes opaque and nonreflective, which indicates a 0 [computer.howstuffworks.com]. We could write our data byte `01001010` as non-reflecting bumps and reflecting lands this way.

CD-RWs eventually lost much of their market share to DVDs.

1999

In 1999 IBM introduced the smallest hard drives in the world at the time: the IBM microdrive in 170MB and 340MB capacities. These were small hard disks, one inch in size, designed to fit into CompactFlash Type II slots. The intent

was to create a device to be used like CompactFlash but with more storage capacity. These were soon replaced by USB flash drives, however, and larger CompactFlash cards once they became available. Like other hard drives, microdrives were mechanical and contained small, spinning disk platters.

2000

USB flash drives were introduced in 2000. These consisted of flash memory encased in a small form factor with a USB interface. Depending on the version of the USB interface used, the speed varies: USB 1.1 is limited to 1.5 Mbps, whereas USB 2.0 can handle 35 Mbps, and USB 3.0 can handle 625 Mbps [diffen.com]. The first USB 3.1 type C drives were announced in March 2015 and have read/write speeds of 530 Mbps [web.archive.org]. Unlike floppy and optical disks, USB devices are harder to scratch but still deliver the same use cases of data storage and transferring and backing up files. Because of this, drives for floppy and optical disks have since faded in popularity, replaced by USB ports.

2005

HDD manufacturers started shipping products using PMR (perpendicular magnetic recording) [youtube] in 2005. Interestingly, this happened at the same time that Apple announced the iPod Nano, which used flash as opposed to the one-inch hard drives in the iPod Mini, causing a bit of an industry hoohaw [eetimes.com].

A typical hard drive contains one or more rigid disks coated with a magnetically sensitive film consisting of tiny

magnetic grains. Data is recorded when a magnetic write-head flies just above the spinning disk, much like a record player and a record, except a needle is in physical contact with the record. As the platters spin, the air in contact with them creates a slight breeze. Just as air on an airplane wing generates lift, the air generates lift on the head's airfoil (books.google.com). The write-head rapidly flips the magnetization of one magnetic region of grains so that its magnetic pole points up or down to denote a 1 or a 0.

The predecessor to PMR was LMR (longitudinal magnetic recording). PMR can deliver more than three times the storage density of LMR. The key difference between the two is that the grain structure and the magnetic orientation of the stored data of PMR media is columnar instead of longitudinal. PMR has better thermal stability and improved SNR (signal-to-noise ratio) as a result of better grain separation and uniformity. It also benefits from better writability because of stronger head fields and better magnetic alignment of the media. Like LMR, PMR's fundamental limitations are based on the thermal stability of magnetically written bits of data and the need to have sufficient SNR to read back the information.

2007

Hitachi Global Storage Technologies announced the first 1TB HDD in 2007. The Hitachi Deskstar 7K1000 used five 3.5-inch 200GB platters and rotated at 7,200 RPM. This is in stark contrast to the world's first HDD, the IBM RAMAC 350, which had a storage capacity of approximately 3.75 MB. How far we have come in 51 years! But wait, there's more.

2009

In 2009 technical work was beginning on NVMe (nonvolatile memory express) [flashmemorysummit.com]. NVM is a type of memory that has persistence, in contrast to volatile memory, which needs constant power to retain data. NVMe filled a need for a scalable host controller interface for PCIe (peripheral component interconnect express)-based SSDs (nvmexpress.org). More than 90 companies were part of the working group that developed the design. This was all based on prior work to define the NVMHCIS (nonvolatile memory host controller interface specification). Opening up a modern server would likely result in finding some NVMe drives. The best NVMe drives today can do about a 3,500 MB/s read and 3,300 MB/s write (pcgamer.com). For the data byte we started with, the character `j`, that is extremely fast compared with a couple of minutes to handweave rope memory for the Apollo Guidance Computer.

TODAY AND THE FUTURE

Now that we have traveled through time a bit, let's take a look at the state of the art for SCM (storage class memory). Like NVM, SCM is persistent but goes further by also providing performance better than or comparable to primary memory, as well as byte addressability [iee.org]. SCM aims to address some of the problems faced by caches today, such as the low density of SRAM (static random access memory). DRAM [arxiv.org] provides better density, but this comes at a cost of slower access times. DRAM also suffers from requiring constant power to refresh memory.

Let's break this down a bit. Power is required since the electric charge on the capacitors leaks off little by little; this means that without intervention, the data on the chip would soon be lost. To prevent this leakage, DRAM requires an external memory-refresh circuit that periodically rewrites the data in the capacitors, restoring them to their original charge.

To solve the problems with density and power leakage, a few SCM technologies are in development: PCM (phase-change memory), STT-RAM (spin-transfer torque random access memory), and ReRAM (resistive random access memory). One nice aspect of all these technologies is their ability to function as MLCs (multilevel cells). This means they can store more than one bit of information, compared with SLCs (single-level cells), which can store only one bit per memory cell, or element. Typically, a memory cell consists of one MOSFET (metal-oxide-semiconductor field-effect transistor). MLCs reduce the number of MOSFETs required to store the same amount of data as SLCs, making them denser or smaller to deliver the same amount of storage as technologies using SLCs. Let's go over how each of these SCM technologies work.

Phase-change memory

PCM is similar to phase change for CD-RWs, described earlier. Its phase-change material is typically GST, or GeSbTe (germanium-antimony-tellurium), which can exist in two different states: amorphous and crystalline. The amorphous state has a higher resistance, denoting a 0, than the crystalline state, denoting a 1. By assigning data values to intermediate resistances, PCM can be used to

store multiple states as an MLC (ieeexplore.ieee.org).

Spin-transfer torque random access memory

STT-RAM consists of two ferromagnetic, permanent magnetic layers separated by a dielectric, an insulator that can transmit electric force without conduction. It stores bits of data based on differences in magnetic directions. One magnetic layer, called the reference layer, has a fixed magnetic direction, while the other magnetic layer, called the free layer, has a magnetic direction that is controlled by passing current. For a 1, the magnetization directions of the two layers are aligned. For a 0, the two layers have opposing magnetic directions.

Resistive random access memory

A ReRAM cell consists of two metal electrodes separated by a metal-oxide layer. This is similar to Masuoka's original flash-memory design, where electrons would tunnel through the oxide layer and get stuck in the floating gate or vice versa. With ReRAM, however, the state of the cell is determined by the concentration of oxygen vacancy in the metal-oxide layer.

SCM downsides and upsides

While these SCM technologies are promising, they still have downsides. PCM and STT-RAM have high write latencies. PCM's latencies are 10 times that of DRAM, while STT-RAM has 10 times the latencies of SRAM. PCM and ReRAM have a limit on write endurance before a hard error occurs, meaning a memory element gets stuck at a particular value (arxiv.org).

In August 2015 Intel announced Optane, a product built

on 3DXPoint, pronounced 3D cross-point (anandtech.com). Optane claims performance 1,000 times faster than NAND SSDs with 1,000 times the performance, while being four to five times the price of flash memory. Optane is proof that SCM is not just experimental. It will be interesting to watch how these technologies evolve.

Helium hard disk drive (HHDD)

An HHDD (helium hard disk drive) is a high-capacity HDD filled with helium and hermetically sealed during manufacturing. Like other hard disks, covered earlier, it is much like a record player with a rotating magnetic-coated platter. Typical HDDs would just have air inside the cavity, but that air causes an amount of drag on the spin of the platters.

Helium balloons float, so we know helium is lighter than air. Helium is, in fact, one-seventh the density of air, therefore reducing the amount of drag on the spin of the platters, causing a reduction in the amount of energy required for the disks to spin. This is actually a secondary feature; the primary advantage of helium is that it allows for packing seven platters in the same form factor that would typically hold only five. Attempting this with air-filled drives would cause turbulence. If you recall the airplane-wing analogy from earlier, this ties in perfectly. Helium reduces drag, thus eliminating turbulence.

As everyone knows, however, helium-filled balloons start to sink after a few days because the gas is escaping the balloons. The same could be said for these drives. It took years before manufacturers had created a container that prevented the helium from escaping the form factor

SMR results in a much more complex writing process, since writing to one track winds up overwriting an adjacent track.

for the life of the drive. Backblaze found that HHDDs had a lower annualized error rate of 1.03 percent, while standard hard drives resulted in 1.06 percent. Of course, that is so small a difference—it is hard to conclude much from it [backblaze.com].

A helium-filled form factor can have an HDD encapsulated that uses PMR, or it could contain an MAMR (microwave-assisted magnetic recording) or HAMR (heat-assisted magnetic recording) drive. Any magnetic storage technology can be paired with helium instead of air. In 2014 HGST (Hitachi Global Storage Technologies) combined two cutting-edge technologies into its 10TB HHDD that used *host-managed* SMR (shingled magnetic recording).

Shingled magnetic recording

PMR, covered earlier, was SMR's predecessor. In contrast to PMR, SMR writes new tracks that overlap part of the previously written magnetic track, which in turn makes the previous track narrower, allowing for higher track density. The technology's name stems from the fact that the overlapping tracks resemble roof shingles.

SMR results in a much more complex writing process, since writing to one track winds up overwriting an adjacent track. This doesn't come into play when a disk platter is empty and data is sequential. Once you are writing to a series of tracks that already contain data, however, this process is destructive to existing adjacent data. If an adjacent track contains valid data, then it must be rewritten. This is quite similar to NAND flash, as covered earlier.

Device-managed SMR devices hide this complexity by having the device firmware manage it, resulting in an

interface like any other hard disk you might encounter. *Host-managed* SMR devices, on the other hand, rely on the operating system to know how to handle the complexity of the drive.

Seagate started shipping SMR drives in 2013, claiming a 25 percent greater density than PMR (anandtech.com).

Microwave-assisted magnetic recording

MAMR is an energy-assisted magnetic storage technology—like HAMR, which is covered next—that uses 20- to 40-GHz frequencies to bombard the disk platter with a circular microwave field. This lowers its coercivity, meaning that the magnetic material of the platter has a lower resistance to changes in magnetization. As already discussed, changes in magnetization of a region of the platter are used to denote a 0 or a 1, so since the platter has a lower resistance to changes in magnetization, the data can be written more densely. The core of this new technology is the spin torque oscillator used to generate the microwave field without sacrificing reliability.

Western Digital, also known as WD, unveiled this technology in 2017 (storagereview.com). Toshiba followed shortly after in 2018 (theregister.co.uk). While WD and Toshiba are busy pursuing MAMR, Seagate is betting on HAMR.

Heat-assisted magnetic recording

HAMR is an energy-assisted magnetic storage technology for greatly increasing the amount of data that can be stored on a magnetic device, such as an HDD, by using heat delivered by a laser to help write data onto the surface of

a platter. The heat causes the data bits to be much closer together on the platter, which allows greater data density and capacity.

This technology is quite difficult to achieve. A 200-milliwatt laser heats a teensy area of the region to 750 degrees F (400 degrees C) quickly before writing the data, while not interfering with or corrupting the rest of the data on the disk (fstoppers.com). The process of heating, writing the data, and cooling must be completed in less than a nanosecond. These challenges required the development of nano-scale surface plasmons, also known as a surface-guided laser, instead of direct laser-based heating, as well as new types of glass platters and heat-control coatings to tolerate rapid spot-heating without damaging the recording head or any nearby data, and various other technical challenges that needed to be overcome (seagate.com).

Seagate first demonstrated this technology, despite many skeptics, in 2013 (computerworld.com). It started shipping the first drives in 2018 (blog.seagate.com).

END OF TAPE, REWIND

This article started off with the state of the art in storage media in 1951 and concludes after looking at the future of storage technology. Storage has changed a lot over time—from paper tape to metal tape, magnetic tape, rope memory, spinning disks, optical disks, flash, and others. Progress has led to faster, smaller, and more performant devices for storing data.

Comparing NVMe to the 1951 Uniservo metal tape shows that NVMe can read 486,111 percent more digits per

Related articles

➡ The Most Expensive One-byte Mistake Did Ken, Dennis, and Brian choose wrong with NUL-terminated text strings?

Poul-Henning Kamp

<https://queue.acm.org/detail.cfm?id=2010365>

➡ Should You Upload or Ship Big Data to the Cloud?

The accepted wisdom does not always hold true.

Sachin Date

<https://queue.acm.org/detail.cfm?id=2933408>

➡ Injecting Errors for Fun and Profit Error-detection and correction features are only as good as our ability to test them.

Steve Chessin

<https://queue.acm.org/detail.cfm?id=1839574>

second. Comparing NVMe to the Zip disks of 1994, you see that NVMe can read 213,623 percent more digits per second.

One thing that remains true is the storing of 0s and 1s. The means by which that is done vary greatly. I hope the next time you burn a CD-RW with a mix of songs for a friend, or store home videos in an optical disc archive (yup, you heard that right (pro. sony), you think about how the nonreflective bumps translate to a 0 and the reflective lands of the disk translate to a 1. If you are creating a mixtape on a cassette, remember that those are closely related to the Datasette used in the Commodore PET. Lastly,

remember to be kind and rewind (this is a tribute to Blockbuster, but there are still open formats for using tape today (wikipedia.org)).

Acknowledgments

Thank you to Robert Mustacchi (<https://twitter.com/rmustacc>), Trammel Hudson (<https://twitter.com/qrs>), and Rick Altherr (<https://twitter.com/kc8apf>) for tidbits (I can't help myself) throughout this article.

Jessie Frazelle is the cofounder and chief product officer of

the Oxide Computer Company. Before that, she worked on various parts of Linux, including containers as well as the Go programming language.

Copyright © 2020 held by owner/author. Publication rights licensed to ACM.

