Focus on Computer Database Storage

Russell Stokes

FOCUS ON COMPUTER DATABASE STORAGE

FOCUS ON COMPUTER DATABASE STORAGE

Russell Stokes

BIBLIOTEX
Digital Library

Focus on Computer Database Storage
by Russell Stokes

Copyright© 2022 BIBLIOTEX

www.bibliotex.com

All rights reserved. No part of this book may be reproduced or used in any manner without the prior written permission of the copyright owner, except for the use brief quotations in a book review.

To request permissions, contact the publisher at info@bibliotex.com

Ebook ISBN: 9781984664303

Published by:
Bibliotex
Canada
Website: www.bibliotex.com

Contents

1

Memory System and Storage Device

Computer memory can refer to many types of memory within a computer, but, typically, it refers to random access memory (RAM). It is physically found on computer chips that are inserted onto the computer's motherboard. RAM is electronic, rather than mechanical; that is, it does not have moving parts and therefore data access to it is very fast. Modern computers often have somewhere between 256 MB (megabytes) and 2 GB (gigabytes) of RAM, although there are, of course, computers with more or less RAM.

RAM is also volatile, meaning that it gets lost when the computer is switched off. The expensive nature of RAM spurred the creation of another type of computer memory called virtual memory. With virtual memory, a slow down in performance is observed only when you try to operate a programme whose files are in the virtual memory. In essence, this slow down is only observed when shifting between programmes. In this way, virtual memory often provides a cheaper alternative to RAM.

A third type of computer memory is cache. There are two types of cache. Primary cache, or level 1 cache, is built right into the central processing unit (CPU) and ensures instant availability of data that the CPU frequently needs. Secondary cache, or level 2 cache, is usually built on a memory chip, is located very close to the CPU, and has a direct connection to the CPU through a dedicated circuit. Secondary cache is bigger in capacity than primary cache.

Cache basically speeds up the rate at which data moves from the main memory to the CPU. The registers form a fourth type of computer memory. These are units within the CPU that contain specific types of data, especially for the Arithmetic and Logic Unit (ALU). A final group of computer memory is called flash. This is a solid-state, rewritable type of memory. Examples of flash memory include BIOS and memory cards. Just like the RAM, they are electronic and not mechanical. They are also non-volatile, and are therefore suitable for digital cameras, mobile phones and other miniaturized computers

OVERVIEW

The system memory is the place where the computer holds current programmes and data that are in use. There are various levels of computer memory (memory), including ROM, RAM, cache, page and graphics, each with specific objectives for system operation. This section focusses on the role of computer memory, and the technology behind it.

Although memory is used in many different forms around modern PC systems, it can be divided into two essential types: RAM and ROM. ROM, or Read Only Memory, is relatively small, but essential to how a computer works. ROM is always found on motherboards, but is increasingly found on graphics cards and some other expansion cards and peripherals. Generally speaking, ROM does not change. It forms the basic instruction set for operating the hardware in the system, and the data within remains intact even when the computer is shut down. It is possible to update ROM, but it's only done rarely, and at need. If ROM is damaged, the computer system simply cannot function.

RAM, or Random Access Memory, is "volatile." This means that it only holds data while power is present. RAM changes constantly as the system operates, providing the storage for all data required by the operating system and software. Because of the demands made by increasingly powerful operating systems and software, system RAM requirements have accelerated dramatically over time.

For instance, at the turn of the millennium a typical computer may have only 128Mb of RAM in total, but in 2007 computers commonly ship with 2Gb of RAM installed, and may include graphics cards with their own additional 512Mb of RAM and more. Clearly, modern computers have significantly more memory than the first PCs of the early 1980s, and this has had an effect on development of the PC's architecture. The trouble is, storing and retrieving data from a large block of memory is more time-consuming than from a small block. With a large amount of memory, the difference in time between a register access and a memory access is very great, and this has resulted in extra layers of cache in the storage hierarchy.

When accessing memory, a fast processor will demand a great deal from RAM. At worst, the CPU may have to waste clock cycles while it waits for data to be retrieved. Faster memory designs and motherboard buses can help, but since the 1990s "cache memory" has been employed as standard between the main memory and the processor.

Not only this, CPU architecture has also evolved to include ever larger internal caches. The organization of data this way is immensely complex, and the system uses ingenious electronic controls to ensure that the data the processor needs next is already in cache, physically closer to the processor and ready for fast retrieval and manipulation. Read on for a closer look at the technology behind computer memory, and how developments in RAM and ROM have enabled systems to function with seemingly exponentially increasing power.

MEMORY SYSTEM

Computer memory is used to store two things: i) instructions to execute a programme and ii) data. When the computer is doing any job, the data that have to be processed are stored in the primary memory. This data may come from an input device like keyboard or from a secondary storage device like a floppy disk.

As programme or the set of instructions is kept in primary memory, the computer is able to follow instantly the set of instructions. For example, when you book ticket from railway reservation counter, the computer has to follow the same steps: take the request, check the availability of seats, calculate fare, wait for money to be paid, store the reservation and get the ticket printed out. The programme containing these steps is kept in memory of the computer and is followed for each request.

But inside the computer, the steps followed are quite different from what we see on the monitor or screen. In computer's memory both programmes and data are stored in the binary form. You have already been introduced with decimal number system, that is the numbers 1 to 9 and 0. The binary system has only two values 0 and 1. These are called *bits*. As human beings we all understand decimal system but the computer can only understand binary system.

It is because a large number of integrated circuits inside the computer can be considered as switches, which can be made ON, or OFF. If a switch is ON it is considered 1 and if it is OFF it is 0. A number of switches in different states will give you a message like this: 110101….10. So the computer takes input in the form of 0 and 1 and gives output in the form 0 and 1 only. Is it not absurd if the computer gives outputs as 0's and 1's only? But you do not have to worry about.

Every number in binary system can be converted to decimal system and vice versa; for example, 1010 meaning decimal 10. Therefore it is the computer that takes information or data in decimal form from you, convert it in to binary form, process it producing output in binary form and again convert the output to decimal form.

The primary memory as you know in the computer is in the form of IC's (Integrated Circuits). These circuits are called Random Access Memory (RAM). Each of RAM's locations stores one *byte* of information. (One *byte* is equal to *8 bits*). A bit is an acronym for *binary digit*, which stands for one binary piece of information. This can be either 0 or 1. You will know more about RAM later. The Primary or internal storage section is made up of several small storage locations (ICs) called cells. Each of these cells can store a fixed number of bits called *word length*.

Each cell has a unique number assigned to it called the address of the cell and it is used to identify the cells. The address starts at 0 and goes up to (N-1). You should know that the memory is like a large cabinet containing as many drawers as there are addresses on memory. Each drawer contains a word and the address is written on outside of the drawer.

CLASSIFICATION OF MEMORY: RAM

RAM is perhaps the most important of the input/output devices. When we talk about computer memory, we are mainly talking about RAM. In this class, when we think about the banks of light switches that can be manipulated, we are thinking of RAM memory. The term Random Access is pretty unfortunate. There is nothing random about how memory is accessed. Random access memory or RAM most commonly refers to computer chips that temporarily store dynamic data to enhance computer performance.

By storing frequently used or active files in random access memory, the computer can access the data faster than if it to retrieve it from the far-larger hard drive. Random access memory is also used in printers and other devices. Random access memory is volatile memory, meaning it loses its contents once power is cut. This is different from non-volatile memory such as hard disks and flash memory, which do not require a power source to retain data.

When a computer shuts down properly, all data located in random access memory is committed to permanent storage on the hard drive or flash drive. At the next boot-up, RAM begins to fill with programmes automatically loaded at startup, and with files opened by the user.

There are several different types of random access memory chips which come several to a "stick." A stick of RAM is a small circuit board shaped like a large stick of gum. Sticks of RAM fit into "banks" on the motherboard.

Adding one or more sticks increases RAM storage and performance. Random access memory is categorized by architecture and speed. As technology progresses, RAM chips become faster and employ new standards so that RAM must be matched to a compatible motherboard. The motherboard will only support certain types of random access memory, and it will also have a limit as to the amount of RAM it can support.

For example, one motherboard may support dual-channel Synchronous Dynamic Random Access Memory (SDRAM), while an older motherboard might only support Single In-line Memory Modules (SIMMS) or Dual In-line Memory Modules (DIMMS). Since random access memory can improve performance, the type and amount of RAM a motherboard will support becomes a major factor when considering a new computer.

If there is a faster, better random access memory chip on the market, the buyer will want to consider purchasing a motherboard capable of using it. A year down the road, that 'new' RAM might be standard, while the buyer may be stuck with an old style motherboard. A new variety of non-volatile random access memory made with nanotubes or other technologies will likely be forthcoming in the near future.

These RAM chips would retain data when powered down. Memory, which is commonly referred to as RAM (Random Access Memory), is a temporary (Volatile) storage area utilized by the CPU. Before a programme can be run, the programme is loaded into the memory which allows the CPU direct access to the programme.

Memory is a necessity for any computer and it is recommend that you have at least 64MB of memory for your IBM or Macintosh. Memory is commonly confused with Hard Drive Space. There are two types of memory; the first type of memory is the memory explained in the above paragraph, this memory is available in computer chips; the other type of memory is actually Hard Drive Space which is stored on the computer Hard Disk Drive.

The Hard drive is actually a physical drive which contains several parts and is generally larger than the amount of memory found in your computer. The below explanation wil help in describing the most commonly found types of RAM in computers.

SINGLE IN-LINE MEMORY MODULE (SIMM)

A SIMM, or single in-line memory module, is a type of memory module containing random access memory used in computers from the early 1980s to the late 1990s. It differs from a dual in-line memory module (DIMM), the most predominant form of memory module today, in that the contacts on a SIMM are redundant on both sides of the module. SIMMs were standardised under the JEDEC JESD-21C standard.

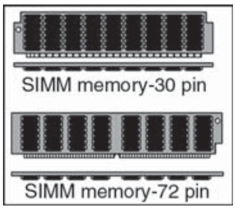

A slender circuit board dedicated to storing memory chips. Each chip is capable of holding 8 to 9 chips per board, the ninth chip usually an error-checking chip (parity/ non parity). The typical bus from the chip to the motherboard is 32-bits wide. When upgrading a Pentium motherboard you will be required to upgrade 2 of the same type of chips at the same time to accommodate the Pentium processor.

SPEED

Memory to a computer that was designed to run slower memory. However, your system will operate at the speed of the slowest memory module.

One thing to keep in mind is that the memory does need to be the same type - for example, SDRAM cannot be mixed with DDR, and DDR cannot be mixed with DDR2 and DDR2 cannot work in a DDR3 system.

We recommend that you use the Crucial Memory Advisor™ or System Scanner tools to find the right memory for your computer. Rather than give memory modules catchy names, the industry refers to modules by their specifications. But if you don't know a lot about memory, the numbers can be confusing. Here's a short summary of the most popular types of memory and what the numbers refer to. Can determine the size amount of the chip by looking at the part number of each chip on the SIMM board. For 2-, 8- and 9- chip SIMMs, all the chips should have the same part numbers. Look at the number that ends with a dash and a digit such as "-7". This is the rate speed or nanoseconds of the chip. With "-7" this would indicate that the memory is 70ns.

SIZE

Look at the four digits to the left of this number; these often carry information about the number of bits in the chip. A 4256 indicates 256K bits arranged in sets of four, for a total of 1Mb. "1000" indicates 1MB of bits arranged in one set. With some types of memory, the last one or two digits may be changed to indicate different kinds of memory; there are 1MB chips that end with 4256, 4257, and 4258. In this case, round the last digits to an even 256 or thousand.

Three-chip SIMMs will typically have two larger chips that are four times the capacity of the third chip (because 4 plus 4 plus 1 makes 9, which is the number of bits needed per byte including parity).

PARITY/ NON- PARITY

To determine if the SIMM is Parity/ Non- Parity, look for x36/ x9 which indicate that the chip is parity (Error checking). x36 is used with 72-pin SIMMs and x9 is for 30-pin SIMMs. If x32/ x8 this would be an indication that the chip in Non- Parity (Non- Error checking) x32 is used with 72-pin SIMMs and x8 is used with 30-pin SIMMs. Another method: count the chips. If you see three or nine discrete chips, the SIMM probably includes a parity. If there are two or eight chips, the SIMM probably does not include parity bit. In this case, divide the number of bits by 8 to determine bytes.

SIMM MEMORY AND PENTIUM COMPUTERS

When updating the computer's memory with a Pentium processor, ensure that you purchase two SIMMs rather than one. Such as if planning to upgrade to 32 MB of RAM that you use two 16 MB SIMMs rather than one. This must be done to allow the memory to work properly with the Pentium processor.

DUAL IN-LINE MEMORY MODULE (DIMM)

A DIMM or dual in-line memory module comprises a series of dynamic random-access memory integrated circuits. These modules are mounted on a printed circuit board and designed for use in personal computers, workstations and servers. DIMMs began to replace SIMMs (single in-line memory modules) as the predominant type of memory module as Intel P5-based Pentium processors began to gain market share.

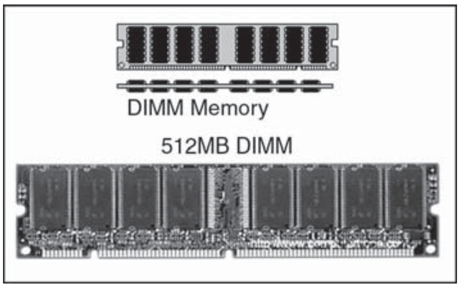

Type of Circuit board that holds memory chips. DIMMS have a 64-bit path because of the Pentium Processor requirements. Because of the new bit path, DIMMS can be installed one at a time un-like SIMMs which on a Pentium would require two be added.

SOME OF THE ADVANTAGES DIMMS HAVE OVER SIMMS

DIMMs have separate contacts on each side of the board, thereby providing twice as much data as a single SIMM. The command address and control signals are buffered on the DIMMs. With heavy memory requirements this will reduce the loading effort of the memory.

STATIC RAM (SRAM)

Static random-access memory (SRAM) is a type of semiconductor memory that uses bistable latching circuitry to store each bit. The term static differentiates it from dynamic RAM (DRAM) which must be periodically refreshed. SRAM exhibits data remanence, but it is still volatile in the conventional sense that data is eventually lost when the memory is not powered.

Static RAM is a type of RAM that holds its data without external refresh, for as long as power is supplied to the circuit. This is contrasted to dynamic RAM (DRAM), which must be refreshed many times per second in order to hold its data contents. SRAMs are used for specific applications within the PC, where their strengths outweigh their weaknesses compared to DRAM:

- *Simplicity*: SRAMs don't require external refresh circuitry or other work in order for them to keep their data intact.
- *Speed*: SRAM is faster than DRAM.

In contrast, SRAMs have the following weaknesses, compared to DRAMs:

- *Cost*: SRAM is, byte for byte, several times more expensive than DRAM.

- *Size*: SRAMs take up much more space than DRAMs (which is part of why the cost is higher).

These advantages and disadvantages taken together obviously show that performance-wise, SRAM is superior to DRAM, and we would use it exclusively if only we could do so economically.

Unfortunately, 32 MB of SRAM would be prohibitively large and costly, which is why DRAM is used for system memory. SRAMs are used instead for level 1 cache and level 2 cache memory, for which it is perfectly suited; cache memory needs to be very fast, and not very large. SRAM is manufactured in a way rather similar to how processors are: highly-integrated transistor patterns photo-etched into silicon. Each SRAM bit is comprised of between four and six transistors, which is why SRAM takes up much more space compared to DRAM, which uses only one (plus a capacitor). Because an SRAM chip is comprised of thousands or millions of identical cells, it is much easier to make than a CPU, which is a large die with a non-repetitive structure.

DYNAMIC RAM (DRAM)

Dynamic random-access memory (DRAM) is a type of random-access memory that stores each bit of data in a separate capacitor within an integrated circuit. The capacitor can be either charged or discharged; these two states are taken to represent the two values of a bit, conventionally called 0 and 1. Since capacitors leak charge, the information eventually fades unless the capacitor charge is refreshed periodically. Because of this refresh requirement, it is a dynamic memory as opposed to SRAM and other static memory. The main memory (the "RAM") in personal computers is dynamic RAM (DRAM). It is the RAM in desktops, laptops and workstation computers as well as some of the RAM of video game consoles.

The advantage of DRAM is its structural simplicity: only one transistor and a capacitor are required per bit, compared to four or six transistors in SRAM. This allows DRAM to reach very high densities. Unlike flash memory, DRAM is volatile memory (vs. non-volatile memory), since it loses its data quickly when power is removed. The transistors and capacitors used are extremely small; billions can fit on a single memory chip.

Dynamic RAM is a type of RAM that only holds its data if it is continuously accessed by special logic called a refresh circuit. Many hundreds of times each second, this circuitry reads the contents of each memory cell, whether the memory cell is being used at that time by the computers or not.

Due to the way in which the cells are constructed, the reading action itself refreshes the contents of the memory. If this is not done regularly, then the DRAM will lose its contents, even if it continues to have power supplied to it. This refreshing action is why the memory is called dynamic. All PCs use DRAM for their main system memory, instead of SRAM, even though DRAMs are slower than SRAMs and require the overhead of the refresh circuitry. It may seem weird to want to make the computer's memory out of something that can only hold a value for a fraction of a second. In fact, DRAMs are both more complicated and slower than SRAMs. The reason that DRAMs are used is simple: they are much cheaper and take up much less space, typically 1/4 the silicon area of SRAMs or less. To build a 64 MB core memory from SRAMs would be very expensive. The overhead of the refresh circuit is tolerated in order to allow the use of large amounts of inexpensive, compact memory.

The refresh circuitry itself is almost never a problem; many years of using DRAM has caused the design of these circuits to be all but perfected. DRAMs are smaller and less expensive than SRAMs because SRAMs are made from four to six transistors (or more) per bit, DRAMs use only one, plus a capacitor.

The capacitor, when energized, holds an electrical charge if the bit contains a "1" or no charge if it contains a "0". The transistor is used to read the contents of the capacitor. The problem with capacitors is that they only hold a charge for a short period of time, and then it fades away. These capacitors are tiny, so their charges fade particularly quickly.

This is why the refresh circuitry is needed: to read the contents of every cell and refresh them with a fresh "charge" before the contents fade away and are lost. Refreshing is done by reading every "row" in the memory chip one row at a time; the process of reading the contents of each capacitor re-establishes the charge.

DRAM is manufactured using a similar process to how processors are: a silicon substrate is etched with the patterns that make the transistors and capacitors (and support structures) that comprise each bit. DRAM costs much less than a processor because it is a series of simple, repeated structures, so there isn't the complexity of making a single chip with several million individually-located transistors.

DDR-SDRAM

Double data rate synchronous dynamic random-access memory (DDR SDRAM) is a class of memory integrated circuits used in computers. DDR SDRAM, also called DDR1 SDRAM, has been superseded by DDR2 SDRAM and DDR3 SDRAM, neither of which is either forward or backward compatible with DDR1 SDRAM -meaning that DDR2 or DDR3 memory modules will not work in DDR1-equipped motherboards, and vice versa.

Compared to single data rate (SDR) SDRAM, the DDR SDRAM interface makes higher transfer rates possible by more strict control of the timing of the electrical data and clock signals. Implementations often have to use schemes such as phase-locked loops and self-calibration to reach the required timing accuracy. The interface uses double pumping (transferring data on both the rising and falling edges of the clock signal) to lower the clock frequency.

One advantage of keeping the clock frequency down is that it reduces the signal integrity requirements on the circuit board connecting the memory to the controller. The name "double data rate" refers to the fact that a DDR SDRAM with a certain clock frequency achieves nearly twice the bandwidth of a SDR SDRAM running at the same clock frequency, due to this double pumping.

With data being transferred 64 bits at a time, DDR SDRAM gives a transfer rate of (memory bus clock rate) × 2 (for dual rate) × 64 (number of bits transferred)/ 8 (number of bits/ byte). Thus, with a bus frequency of 100 MHz, DDR SDRAM gives a maximum transfer rate of 1600 MB/s.

"Beginning in 1996 and concluding in June 2000, JEDEC developed the DDR (Double Data Rate) SDRAM specification

(JESD79)." JEDEC has set standards for data rates of DDR SDRAM, divided into two parts. The first specification is for memory chips, and the second is for memory modules.

The DDR-SDRAM (Double Data Rate SDRAM) is a memory, based on the SDRAM technology, which doubles the transfer rate of the SDRAM using the same frequency. Data are read or written into memory based on a clock.

Standard DRAM memories use a method known as SDR (Single Data Rate) involving reading or writing a piece of data at each leading edge.

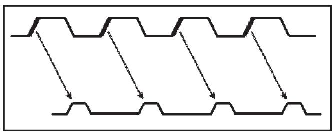

The DDR doubles the frequency of reading/writing, with a clock at the same frequency, by sending data to each leading edge and to each trailing edge.

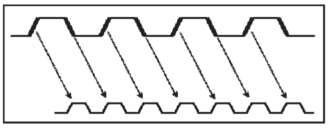

DDR memories generally have a product name such as PCXXXX where "XXXX" represents the speed in Mo/s.

DDR2-SDRAM

DDR2 SDRAM is a double data rate synchronous dynamic random-access memory interface. It superseded the original DDR SDRAM specification, and has since been superseded by DDR3 SDRAM. DDR2 DIMMs are neither forward compatible with DDR3 nor backward compatible with DDR.

In addition to double pumping the data bus as in DDR SDRAM (transferring data on the rising and falling edges of the bus clock signal), DDR2 allows higher bus speed and requires lower power by running the internal clock at half the speed of the data bus. The two factors combine to require a total of four data transfers per internal clock cycle. With data being transferred 64 bits at a time, DDR2 SDRAM gives a transfer rate of (memory clock rate) × 2 (for bus clock multiplier) × 2 (for dual rate) × 64 (number of bits transferred)/ 8 (number of bits/byte). Thus with a memory clock frequency of 100 MHz, DDR2 SDRAM gives a maximum transfer rate of 3200 MB/s.

Since the DDR2 internal clock runs at half the DDR external clock rate, DDR2 memory operating at the same external data bus clock rate as DDR results in DDR2 being able to provide the same bandwidth but with higher latency. Alternatively, DDR2 memory operating at twice the external data bus clock rate as DDR may provide twice the bandwidth with the same latency. The best-rated DDR2 memory modules are at least twice as fast as the best-rated DDR memory modules.

DDR2 (or DDR-II) memory achieves speeds that are twice as high as those of the DDR with the same external frequency. QDR (Quadruple Data Rate or quad-pumped) designates the reading and writing method used. DDR2 memory in fact uses two separate channels for reading and writing, so that it is able to send or receive twice as much data as the DDR.

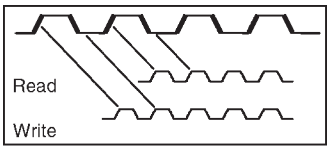

DDR2 also has more connectors than the classic DDR (240 for DDR2 compared with 184 for DDR).

OPERATION OF RANDOM ACESS MEMORY

Alternatively referred to as main memory, primary memory, or system memory, Random Access Memory (RAM) is a computer storage location that allows information to be stored and accessed quickly from random locations within DRAM on a memory module. Because information is accessed randomly instead of sequentially like a CD or hard drive the computer is able to access the data much faster than it would if it was only reading the hard drive.

However, unlike ROM and the hard drive RAM is a volatile memory and requires power in order to keep the data accessible, if power is lost all data contained in memory lost.

As the computer loads parts of the operating system and drivers are loaded into memory, which allows the CPU to process the instructions much faster and your computer to load faster. After the operating system has loaded, each programme you open such as the browser you're using to view this page is loaded into memory while it is running. If too many programmes are open the computer will swap the data in the memory between the RAM and the hard disk drive.

The random access memory comprises hundreds of thousands of small capacitors that store loads. When loaded, the logical state of the capacitor is equal to 1, otherwise it is 0, meaning that each capacitor represents one memory bit.

Given that the capacitors become discharged they must be constantly recharged (the exact term is refresh) at regular intervals, known as the refresh cycle. DRAM memories for example require refresh cycles of around 15 nanoseconds (ns). Each capacitor is coupled with a transistor (MOS-type) enabling "recovery" or amendment of the status of the capacitor. These transistors are arranged in the form of a table (matrix) thus we access a memory box (also called memory point) via a line and a column.

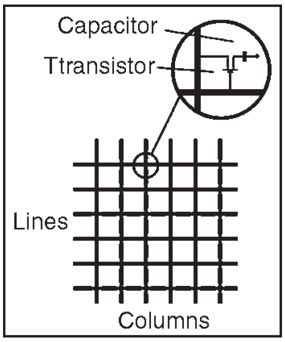

Each memory point is thus characterised by an address which corresponds to a row number and a column number. This access is not instant and the access time period is known as latency time.

Consequently, time required for access to data in the memory is equal to cycle time plus latency time. Thus, for a DRAM memory, access time is 60 nanoseconds (35ns cycle time and 25ns latency time). On a computer, the cycle time corresponds to the opposite of the clock frequency; for example, for a computer with frequency of 200 MHz, cycle time is 5 ns (1/200*106)).

Consequently a computer with high frequency using memories with access time much longer than the processor cycle time must perform wait states to access the memory. For a computer with frequency of 200 MHz using DRAM memories (and access time of 60ns), there are 11 wait states for a transfer cycle. The computer's performance decreases as the number of wait states increases, therefore we recommend the use of faster memories.

EXTERNAL MEMORY

MAGNETIC DISKS

A hard drive is a special disk that is usually mounted permanently inside your computer's cabinet. You rarely see the hard drive, and almost never take it out. Hard drives are made of different material than floppies, and they are physically hard (although if you touched the actual hard surface, you would destroy it!) They spin much more quickly than floppies, and require much more precision. They are sealed inside a special case, and that is sealed inside the computer case. A hard drive has a much larger capacity than a floppy, and is much faster at saving and retrieving information. Modern computers frequently have hard drives with 500 MB or more of capacity. As this capacity grows, people are beginning to measure it in terms of gigabytes. Software programmes are always becoming larger and taking more room on hard drives. It never takes long to completely fill up the capacity of a drive.

READ ONLY MEMORY (ROM)

There is another memory in computer, which is called Read Only Memory (ROM). Again it is the ICs inside the PC that form the ROM. The storage of programme and data in the ROM is permanent. The ROM stores some standard processing programmes supplied by the manufacturers to operate the personal computer.

The ROM can only be read by the CPU but it cannot be changed. The basic input/output programme is stored in the ROM that examines and initializes various equipment attached to the PC when the switch is made ON. The memories, which do not loose their content on failure of power supply, are known as non-volatile memories. ROM is non-volatile memory.

Memory is like the banks of switches we have been thinking about that simply contain 1/0 binary patterns. Read Only Memory (abbreviated ROM) is a special kind of memory that cannot be changed. The most basic instructions for the CPU are built in to the ROM at the factory. To the end user, there really is very little to worry about regarding ROM. The amount of ROM in your computer doesn't really matter to you. All your programmes and information will go into another kind of memory that we will learn about soon.

DEFINITION

ROM is an integrated-circuit memory chip that contains configuration data. ROM is commonly called firmware because its programming is fully embedded into the ROM chip. As such, ROM is a hardware and software in one.

Because data is fully incorporated at the ROM chip's manufacture, data stored can neither be erased nor replaced. This means permanent and secure data storage. However, if a mistake is made in manufacture, a ROM chip becomes unusable. The most expensive stage of ROM manufacture, therefore, is creating the template. If a template is readily available, duplicating the ROM chip is very easy and affordable.

A ROM chip is also non-volatile so data stored in it is not lost when power is turned off. There is a type of memory that stores data without electrical current; it is the ROM (Read Only Memory) or is sometimes called non-volatile memory as it is not erased when the system is switched off.

This type of memory lets you stored the data needed to start up the computer. Indeed, this information cannot be stored on the hard disk since the disk parameters (vital for its initialisation) are part of these data which are essential for booting.

Different ROM-type memories contain these essential start-up data, that is:

- The BIOS is a programme for controlling the system's main input-output interfaces, hence the name BIOS ROM which is sometimes given to the read-only memory chip of the mother board which hosts it.
- *The bootstrap loader*: a programme for loading (random access) memory into the operating system and launching it. This generally seeks the operating system on the floppy drive then on the hard disk, which allows the operating system to be launched from a system floppy disk in the event of malfunction of the system installed on the hard disk.
- The CMOS Setup is the screen displayed when the computer starts up and which is used to amend the system parameters (often wrongly referred to as BIOS).
- The Power-On Self Test (POST), a programme that runs automatically when the system is booted, thus allowing the system to be tested (this is why the system "counts" the RAM at start-up).

Given that ROM are much slower than RAM memories (access time for a ROM is around 150 ns whereas for SDRAM it is around 10 ns), the instructions given in the ROM are sometimes copied to the RAM at start-up; this is known as shadowing, though is usually referred to as shadow memory). One major type of memory that is used in PCs is called read-only memory, or ROM for short. ROM is a type of memory that normally can only be read, as opposed to RAM which can be both read and written.

There are two main reasons that read-only memory is used for certain functions within the PC:

- *Permanence*: The values stored in ROM are always there, whether the power is on or not. A ROM can be removed from the PC, stored for an indefinite period of time, and then replaced, and the data it contains will still be there. For this reason, it is called non-volatile storage. A hard disk is also non-volatile, for the same reason, but regular RAM is not.
- *Security*: The fact that ROM cannot easily be modified provides a measure of security against accidental (or malicious) changes to its contents. You are not going to find viruses infecting true ROMs, for example; it's just not possible. (It's technically possible with erasable EPROMs, though in practice never seen.)

Read-only memory is most commonly used to store system-level programmes that we want to have available to the PC at all times. The most common example is the system BIOS programme, which is stored in a ROM called (amazingly enough) the system BIOS ROM.

Having this in a permanent ROM means it is available when the power is turned on so that the PC can use it to boot up the system. Remember that when you first turn on the PC the system memory is empty, so there has to be something for the PC to use when it starts up.

While the whole point of a ROM is supposed to be that the contents cannot be changed, there are times when being able to change the contents of a ROM can be very useful.

There are several ROM variants that can be changed under certain circumstances; these can be thought of as "mostly read-only memory":^)

TYPES OF ROMS

There are types of ROMs with a description of their relative modifiability.

ROM

Read-only memory (ROM) is a class of storage medium used in computers and other electronic devices. Data stored in ROM cannot be modified, or can be modified only slowly or with difficulty, so it is mainly used to distribute firmware (software that is very closely tied to specific hardware, and unlikely to need frequent updates).

A regular ROM is constructed from hard-wired logic, encoded in the silicon itself, much the way that a processor is. It is designed to perform a specific function and cannot be changed. This is inflexible and so regular ROMs are only used generally for programmes that are static (not changing often) and mass-produced. This product is analagous to a commercial software CD-ROM that you purchase in a store.

PROGRAMMABLE ROM (PROM)

A programmable read-only memory (PROM) or field programmable read-only memory (FPROM) or one-time programmable non-volatile memory (OTP NVM) is a form of digital memory where the setting of each bit is locked by a fuse or antifuse. Such PROMs are used to store programmes permanently. The key difference from a strict ROM is that the programming is applied after the device is constructed.

There is another type of primary memory in computer, which is called Programmable Read Only Memory (PROM). You know that it is not possible to modify or erase programmes stored in ROM, but it is possible for you to store your programme in PROM chip. Once the programmes are written it cannot be changed and remain intact even if power is switched off. Therefore programmes or instructions written in PROM or ROM cannot be erased or changed.

This is a type of ROM that can be programmed using special equipment; it can be written to, but only once. This is useful for companies that make their own ROMs from software they write, because when they change their code they can create new PROMs without requiring expensive equipment. This is similar to the way a CD-ROM recorder works by letting you "burn" programmes onto blanks once and then letting you read from them many times. In fact, programming a PROM is also called burning, just like burning a CD-R, and it is comparable in terms of its flexibility.

ERASABLE PROGRAMMABLE ROM (EPROM)

This stands for Erasable Programmable Read Only Memory, which over come the problem of PROM and ROM. EPROM chip can be programmed time and again by erasing the information stored earlier in it. Information stored in EPROM exposing the chip for some time ultraviolet light and it erases chip is reprogrammed using a special programming facility.

When the EPROM is in use information can only be read.

An EPROM (rarely EROM), or erasable programmable read only memory, is a type of memory chip that retains its data when its power supply is switched off. In other words, it is non-volatile. It is an array of floating-gate transistors individually programmed by an electronic device that supplies higher voltages than those normally used in digital circuits. Once programmed, an EPROM can be erased by exposing it to strong ultraviolet light source (such as from a mercury-vapor light). EPROMs are easily recognizable by the transparent fused quartz window in the top of the package, through which the silicon chip is visible, and which permits exposure to UV light during erasing.

An EPROM is a ROM that can be erased and reprogrammed. A little glass window is installed in the top of the ROM package, through which you can actually see the chip that holds the memory. Ultraviolet light of a specific frequency can be shined through this window for a specified period of time, which will erase the EPROM and allow it to be reprogrammed again. Obviously this is much more useful than a regular PROM, but it does require the erasing light. Continuing the "CD" analogy, this technology is analogous to a reusable CD-RW.

ELECTRICALLY ERASABLE PROGRAMMABLE ROM (EEPROM)

Short for Electrically Eras able Programmable Read-Only Memory, EEPROM is a PROM that can be erased and reprogrammed using an electrical charge that was developed by George Perlegos while at Intel in 1978. Unlike most memory inside a computer, this memory remembers data when the power is turned off. EEPROM was a replacement for PROM and EPROM chips and is used for later computer's BIOS that were built after 1994. Having a computer with an EEPROM allows a computer user to update the BIOS in their computer without having to open the computer or remove any chips.

The next level of erasability is the EEPROM, which can be erased under software control. This is the most flexible type of ROM, and is now commonly used for holding BIOS programmes.

When you hear reference to a "flash BIOS" or doing a BIOS upgrade by "flashing", this refers to reprogramming the BIOS EEPROM with a special software programme. Here we are blurring the line a bit between what "read-only" really means, but remember that this rewriting is done maybe once a year or so, compared to real read-write memory (RAM) where rewriting is done often many times per second!

CACHE MEMORY

The speed of CPU is extremely high compared to the access time of main memory. Therefore the performance of CPU decreases due to the slow speed of main memory. To decrease the mismatch in operating speed, a small memory chip is attached between CPU and Main memory whose access time is very close to the processing speed of CPU.

It is called CACHE memory. CACHE memories are accessed much faster than conventional RAM. It is used to store programmes or data currently being executed or temporary data frequently used by the CPU. So each memory makes main memory to be faster and larger than it really is. It is also very expensive to have bigger size of cache memory and its size is normally kept small.

SECONDARY STORAGE DEVICES

Alternatively referred to as external memory and auxiliary storage, secondary storage is a storage medium that holds information until it is deleted or overwritten regardless if the computer has power. For example, a floppy disk drive and hard drive are both good examples of secondary storage devices. As can be seen by the below picture there are three different storage on a computer, although primary storage is accessed much faster than secondary storage because of the price and size limitations secondary storage is used with today's computers to store all your programmes and your personal data.

DEFINITION

Storage Devices are the data storage devices that are used in the computers to store the data. The computer has many types of data storage devices. Some of them can be classified as the removable data Storage Devices and the others as the non removable data Storage Devices. The data Storage Devices come in many sizes and shapes. And more over the technology used for the storage of the data over them is also altogether different.

The storage devices are one of the most important components of the computer system. The memory is of two types; one is the primary memory and the other one is the secondary memory. The primary memory is the volatile memory and the secondary memory is the non- volatile memory. The volatile memory is the kind of the memory that is erasable and the non- volatile memory is the one where in the contents cannot be erased.

Basically when we talk about the data storage devices it is generally assumed to be the secondary memory. The secondary memory is used to store the data permanently in the computer. The secondary storage devices are usually as follows: hard disk drives - this is the most common type of storage device that is used in almost all the computer systems. The other ones include the floppy disk drives, the CD ROM, and the DVD ROM. The flash memory, the USB data card etc.

The storage devices are used to record the data over any storage surface. The memories may also be of different types depending upon the architecture and the design like the optical data storage memory, magnetic media storage and the mechanical storage media etc and also the flash memory devices etc. The storage devices are actually defined as the peripheral unit which holds the data like the tape, disk, or flash memory card etc.

The most of the drives that are used for the purpose of data storage are fragile and the data can be easily corrupted in them. The data storage devices are the ones that are also used for the backup and the archiving of the data. The data storage devices were at a time in the past used to be too costly and expensive. But these days the data storage devices are becoming cheap day by day. Hence the data storage devices price is falling. So, we are in a position to get a storage device for a comparatively cheaper price than the earlier drive. The technology is improving a lot and now the memory storage capacity has gone up TB. The data in the storage devices can be in the form of the files, data bases, digital video and the audio etc. The storage devices that are called as the non-volatile can store the data permanently until otherwise erased purposely. This is in the case of the hard disk drives or the floppy disk drives. The other kinds of the storage media like for example the CD and the DVD can even have again two types of the storage; the first one is that in which the data once written cannot be erased.

It is stored permanently over it. While the second type of the CD's or the DVD's are called as the rewritable; where in the data that is once written can be erased completely and the same storage device can be used again for storing the different data. The storage devices are used to record the data over any storage surface.

The memories may also be of different types depending upon the architecture and the design like the optical data storage memory, magnetic media storage and the mechanical storage media etc and also the flash memory devices etc. The storage devices are actually defined as the peripheral unit which holds the data like the tape, disk, or flash memory card etc.

The most of the drives that are used for the purpose of data storage are fragile and the data can be easily corrupted in them. The data storage devices are the ones that are also used for the backup and the archiving of the data. The data storage devices were at a time in the past used to be too costly and expensive.

But these days the data storage devices are becoming cheap day by day. Hence the data storage devices price is falling.

So, we are in a position to get a storage device for a comparatively cheaper price than the earlier drive. The technology is improving a lot and now the memory storage capacity has gone up TB. The data in the storage devices can be in the form of the files, databases, digital video and the audio etc. The storage devices that are called as the non-volatile can store the data permanently until otherwise erased purposely. This is in the case of the hard disk drives or the floppy disk drives.

The other kinds of the storage media like for example the CD and the DVD can even have again two types of the storage; the first one is that in which the data once written cannot be erased. It is stored permanently over it. While the second type of the CD's or the DVD's are called as the rewritable; where in the data that is once written can be erased completely and the same storage device can be used again for storing the different data.

TYPES OF STORAGE DEVICES

Physical components or materials on which data is stored are called storage media. Hardware components that read/write to storage media are called storage devices.

REMOVABLE DATA STORAGE DEVICES

The Removable Storage Media are the types of storage devices that are used for the storage of the data. The main reason for using the removable storage media is the ease of portability. That is the ease with which the data can be transferred. The data is of importance and is necessary for the movement in some cases. Here the case of portability of the storage device comes into the picture. The removable storage media serves the same purpose. More over there are many application of the storage of the data over the removable media. The first and the foremost is the portability that is already discussed above. The next is the issue of the backup and archiving of the data that is important and also sensitive in nature. The backup of data is as important as the storage of data is important.

The CD and the DVD are key players in the removable storage types that are used for the purpose of backup and archiving of the data. The main reason for storing the data over the removable devices is that the data can be easily moved physically from one place to another at the time of any disaster or any other casualty.

The floppy disk drive was the removable storage media that was used for the purpose of storing the data over the disk for the backup. This was the case in the 80's; but as the pace has increased in the technology the CD's and the DVD's had gained momentum for the purpose of providing backup to the data. The floppy disk drives that are available even today do not provide enough memory capacity. And now the concept of the flash memory devices is doing rounds for providing data transfer from place to another. There are some situations where in the normal data transfer through the networking is not possible.

But as in the case of the removal storage media the data transfer can be easily done. Hence the main advantage regarding the removal storage media is the data transfer from the computers to the other computers where there is no existence of the network. The data transfer facility is important even in case; when the network actually exists but there is a failure of the network. In such situations the external data transfer is quite useful.

Also in some cases the data from one system is required to be used through the other system. Here the network connections are not at all possible. Hence the data has to be physically moved from one place to the other. The example for this may be the regular video or the audio. If it is so desired that the video or the audio is to be used through the mode of the large panel television; but the digital video or the audio is in the computer. Hence in such a situation it is not possible to connect the two different types of systems that are the computer and the television. In such a situation the removable data storage devices come into the play.

The solution to this problem is that the digital video or the audio should be burned over the CD or the DVD in the format that is supported by the VCD player associated with the television. And after that is done the video could be used through the television. Similar is the case with the audio also. The removable storage media has one more example of the Zip drives.

NON- REMOVABLE DATA STORAGE DEVICES

The storage media is generally referred to the secondary memory devices such as the hard disk and the floppy disk drive, CD drive, DVD disk drive, USB drives and the flash memory devices and also to the Zip drives. But apart from the hard disk, none of them are the devices that are non-removable in nature. It means that the hard disk is the only mentioned device that is considered as the Non-Removable

Storage Media. All the rest are the removable storage media.

Another example that can be used for the Non-Removable Storage Media is the RAM. But it is a volatile memory that cannot be used for permanently storing the data. It is just the primary memory that is of utmost importance to the system but at the same time cannot be used for the backup or the archiving of the data.

The non removable storage media is fixed permanently in to the computer case and can also be used for the backup or the archiving of the data. The hard disk drives are the best suitable example for this type of storage media.

It also important to mention here that the hard disk drive can also either be internal or the external disk drive. In the internal hard disk; the drive is fixed permanently to the computer case. Where as, in the case of the external disk drive the hard drive can be moved from one place to another. Such a drive has a hard casing made over the normal disk drive. It is made to with stand the shocks in case of any physical damage.

The external hard disk drive's case acts as shield and a shock absorber in case of any physical damage. The external hard drive is provided with a USB cable that is plugged into the computer case for the ease of usage. Hence this type of the external hard disk becomes the removal type of data storage. The discussion matters here only with the non-removable storage media. Hence the internal hard disks and its types can be discussed elaborately.

The internal hard disk drive is usually fixed in the computer case. The type of hard disk controller gives the name for the different types of the hard disks. The IDE hard disk is the normal hard disk that is used in the regular computers. The IDE controller is used for the purpose of providing connections to this type of the hard disk drive.

The SCSI controllers are used to provide the connections for the SCSI hard disk. The other technology that can be implemented by the internal hard disk drives is the RAID. The RAID is the Redundant Array of Independent (or Inexpensive) disks. In the regular cases the hard disk or mostly the internal hard disk is connected to the motherboard. The motherboard has the IDE connector where in the bus from the hard disk drive is to be plugged in.

If two or more than two hard disk drives are used, then in such case the hard disk drive should clearly be specified as which one is used as the master drive and the other one should also clearly be specified as the slave. The hard disk drive has the jumper settings, by which it can be determined or specified. The internal hard disk is specifically used as the secondary memory. And most importantly the secondary memory that is used in the computer is the non- removable storage media. The removal storage media is not used as the secondary memory.

SECONDARY STORAGE DEVICES

A secondary storage device is another memory device for a computer that will hold and store more information. Examples of secondary storage devices are floppy disks and external hard drives. These devices will hold information regardless of if the computer itself has power.

Secondary Storage refers to non-volatile data storage which is not directly accessible by the CPU and only accessible via primary storage devices using I/O (Input/Output) channels or device drivers. Typical examples of Secondary Storage are hard disks, floppy disks, optical storage devices (such as CD, DVD drives), magnetic tape data storage devices, RAM drives, and flash memory (such as USB sticks/keys.)

The secondary storage helps in securing the data on media types and storage categories, as data is vulnerable to network attacks, administrative access and media theft. There are many organizations who are working with third parties and disaster recovery efforts. Often data goes offsite and is in the hands of employees that are not authorized to see critical company data. And storage consolidation opens the door to greater administrative access. All of these trends drive the need to ensure the data at rest is secure. Secondary storage devices hold files that are not currently being used.

For a file to be used it must first be copied to main memory first. After any modifications files must be saved to secondary storage. It is advisable to save your data files at the regular intervals as you work on them as data can be lost unexpectedly because of various reasons like interruption in power supply, memory management problems, freezing keyboard, etc. Compared to main memory, secondary memory is slower, bigger, and cheaper (per unit storage.)

Any method for secondary storage must involve two physical parts: a peripheral device (the component of the computer which 'reads' in or 'writes' out the information to/from the system unit,) and an input/output medium, on which the information is actually stored. Diskettes and cassettes are types of media. The medium has to be 'in' a corresponding peripheral device for information transfer to take place. In most methods of secondary storage, this transfer is realized by passing the medium by a read/write head, which is capable of sensing/writing information from/ to that type of medium. Secondary storage media may be removable (easily separable from the computer) like your diskettes, or fixed, like hard disks, which, for example, can be found in relatively expensive PCís.

A fixed medium generally has much greater capacity (and is faster) compared to a removable one of the same type. On the other hand, by replacing your removable medium by another one when it is full, you can attain a virtually unlimited storage capacity. Another nice feature of removable media is that they allow backup copies of important material to be taken and kept away from the computer, so that they do not get corrupted if some disaster strikes the computer. The act of retrieving pieces of stored information is called access.

There may be two kinds of access to data stored as a sequence of items: Sequential access, where the items are traversed one by one from the beginning of the sequence to the desired one, and direct access (or random access), where any item can be accessed relatively independently of its location in the sequence. Different types of secondary memory media support one or more of these different methods, depending on their nature. Cassette tapes, for example, are sequential access media: You have to 'go through' the first 10 minutes if you want to listen to a piece which begins at the 11th minute of the tape. The most common type of secondary storage medium is the magnetic disk.

Diskettes (small disks made of flexible plastic) and hard disks (made of rigid metal) are the two different kinds of magnetic disks. Magnetic disks are coated with a magnetizable substance, the spots on which mean '0' or '1', depending on whether they are magnetized or not. Each surface of the disk is subdivided into concentric rings called tracks. Disks with bigger capacity have more tracks than others. In larger computers, one stores the same amount of data in each track and keeps several disks mounted on a shaft on top of each other as a disk pack. The group of tracks which are at the same position in their respective disks on a disk pack is called the disk cylinder for that position. The associated peripheral device is called the disk unit. At least one read-write head is assigned for each surface.

The read-write heads are mounted on a device called the access mechanism, which positions them on the cylinder in which the appropriate data item is to be located. The time for the read-write head to move from its initial location to the appropriate track is called seek time. In larger computers, the shaft on which the disk is mounted continuously rotates at a great speed, and the time required for the relevant portion of the track to come near the head is called rotational delay.

The time needed for the actual information transfer between head and disk is called data movement time. Disk access time is the sum of these three components. In larger computers, the read/write head never touches the disk. Diskettes, on the other hand, rotate only when a read/write is taking place.

Furthermore, the head actually touches the diskette. Diskettes usually come in one of two diameters: 3.5 inches or 5.25 inches. 3.5-inch diskettes are products of a newer technology and can store more than 5.25-inch diskettes.

Different types of computers employ different methods of data organization on a diskette. For this reason, a new diskette has to be formatted before it is used on a particular computer. Formatting a disk involves designating parts of each track as individually addressable (directly accessible) sectors. Each diskette contains a file directory, in which the names, lengths and starting addresses of the files (collections of related data) stored on it are listed.

People with many files to store in the same disk are encouraged to use a tree-like directory structure in which related files are kept in separate subdirectories. If a computer system has enough main memory and a programme designed to read/write data from/to a disk is to be executed, one can transfer the contents of the disk to main memory and then use a method called disk emulation (or RAM disk) to 'trick' the computer into thinking that it is accessing the disk, while it actually deals with the much faster main memory. An older type of secondary storage medium is magnetic tape.

Characters of data are represented in byte form across the tracks which go along the length of the tape. Other types of secondary storage media include mass storage units, which are combinations of spools of magnetic tape that can be automatically retrieved from a huge 'library' and read to disk. These store enormous amounts of information and do not require the intervention of a human operator, note that some disk packs and tapes need this kind of intervention. Optical disks employ a relatively new technology.

They use laser beams for the read/write operations. Most optical disks are of the CD-ROM (Compact Disk/Read-Only

Memory) type; whose vendor-determined contents cannot be changed by the user. Because of their great capacity, these are used for storing multimedia (*i.e.* mixed text, graphics, sound, etc.) applications of great size. Their speed is low, compared to magnetic hard disks. As secondary storage media can be damaged and files on them become corrupted, it is suggested to make back-up copies of valuable files on a regular basis. Lots of people skip the last but very important step in the backup procedure - check that the backup copy of files is not damaged.

TYPES OF SECONDARY STORAGE DEVICES

It is important to know the difference between secondary storage and a computer's main memory. Secondary storage is also called auxiliary storage and is used to store data and programmes when they are not being processed. Secondary storage is more permanent than main memory, as data and programmes are retained when the power is turned off. The needs of secondary storage can vary greatly between users. A personal computer might only require 20,000 bytes of secondary storage but large companies, such as banks, may require secondary storage devices that can store billions of characters. Because of such a variety of needs, a variety of storage devices are available. The two most common types of secondary storage are magnetic tapes and magnetic disks.

FLOPPY DISK

They are plastic square disks, usually with a silver or black sliding piece going across the top. These disks come in a variety of colours and they hold about 144 million bytes. (Bytes are characters, symbols and letters).

MAGNETIC TAPE STORAGE

Magnetic tape data storage uses digital recording on magnetic tape to store digital information. Modern magnetic tape is most commonly packaged in cartridges and cassettes. The device that performs actual writing or reading of data is a tape drive. Autoloaders and tape libraries automate cartridge handling.

Magnetic tape is a one-half inch or one-quarter inch ribbon of plastic material on which data is recorded. The tape drive is an input/output device that reads, writes and erases data on tapes. Magnetic tapes are erasable, reusable and durable. They are made to store large quantities of data inexpensively and therefore are often used for backup. Magnetic tape is not suitable for data files that are revised or updated often because it stores data sequentially.

DISKETTES

A diskette is a random access, removable data storage medium that can be used with personal computers. The term usually refers to the magnetic medium housed in a rigid plastic cartridge measuring 3.5 inches square and about 2millimeters thick. Also called a "3.5-inch diskette," it can store up to 1.44 megabytes (MB) of data.

Although many personal computers today come with a 3.5-inch diskette drive pre-installed, some notebook computers and centrally-administered desktop computers omit them.

Made of flexible Mylar, a diskette can record data as magnetized spots on tracks on its surface. Diskettes became popular along with the personal computer. The older diskette, 5-1/4 inches in diameter, is still in use, but newer computers use the 3-1/2 inch diskette.

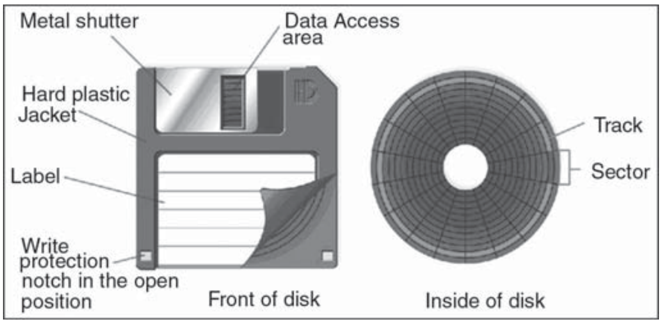

The 3-1/2 inch diskette has the protection of a hard plastic jacket, a size to fit conveniently in a shirt pocket or purse, and the capacity to hold significantly more data than a 5-1/4 inch diskette. Diskettes offer particular advantages, which, as you will see, are not readily available with hard disk:

PORTABILITY.

Diskettes easily transport data from one computer to another. Workers, for example, carry their files from office computer to home computer and back on a diskette instead of in a briefcase. Students use the campus computers but keep their files on their own diskettes.

BACKUP

It is convenient to place an extra copy of a hard disk file on a diskette.

NEW SOFTWARE

Although, for convenience, software packages are kept on hard disk, new software out of the box may come on diskettes (new software also may come on CD-ROM disks, which we will discuss shortly). The end of the diskettes useful lifetime may be upon us. In 1998 Macintosh introduced its new computer, the IMAC, without a floppy disk drive.

Alternatives such as Zip disks (discussed later), or transferring data via networks are making the low-capacity diskette become obsolete. The diskette was introduced in the early 1970s by IBM as a new type of secondary storage. Originally they were eight inches in diameter and were thin and flexible which gave them the name floppy disks, or floppies. Diskettes are used as the principle medium of secondary storage for personal computers. They are available in two different sizes: 3 1/2 inch and 5 1/4 inch.

STORAGE CAPACITY

Before you can store data on your diskette, it must be formatted (G).

The amount of data you can store on a diskette depends on the recording density and the number of tracks on the diskette. The recording density is the number of bits (G) that can be recorded on one inch of track on the diskette, or bits per inch (bpi).

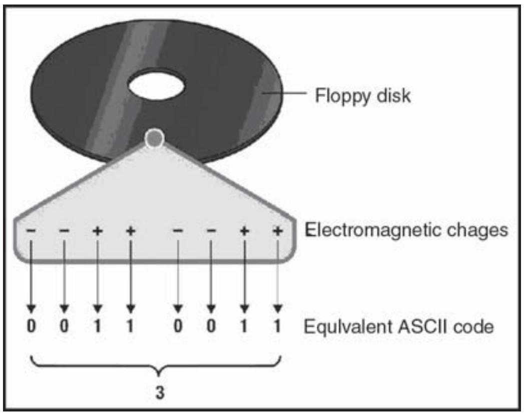

The second factor that influences the amount of data stored on a diskette is the number of tracks on which the data can be stored or tracks per inch (tpi). Commonly used diskettes are referred to as either double-density or high-density (single-density diskettes are no longer used). Double-density diskettes (DD) can store 360K for a 5 1/4 inch diskette and 720K for a 3 1/2 inch diskette. High-density diskettes (HD) can store 1.2 megabytes (G) on a 5 1/4 inch diskette and 1.44 megabytes on a 3 1/2 inch diskette.

CARE OF DISKETTES

You should keep diskettes away from heat, cold, magnetic fields (including telephones) and contaminated environments such as dust, smoke, or salt air. Also keep them away from food and do not touch the disk surface.

MAGNETIC DISK STORAGE

The most common physical device for storing files is the magnetic disk. Actually, a disk typically contains several rotating disks, or platters. The surfaces of the platters are covered in metal oxide, and read/written by electromagnetic recording heads, rather like those on an audio cassette recorder.

There is one head for each surface, and all the heads move together. The disk rotates at around 3600 rpm (or approx 90mph), with the heads floating microscopic distances above the surfaces. Modern disks for workstations typically hold 500MB - 9GB, and cost of the order of £200 - 3000; prices are currently dropping rapidly.

Magnetic disks are the most widely used storage medium for computers. A magnetic disk offers high storage capacity, reliability, and the capacity to directly access stored data. Magnetic disks hold more data in a small place and attain faster data access speeds. Types of magnetic disks include diskettes, hard disks, and removable disk cartridges.

Diskettes and hard disks are magnetic media; that is, they are based on a technology of representing data as magnetized spots on the disk with a magnetized spot representing a 1 bit and the absence of such a spot representing a 0 bit.

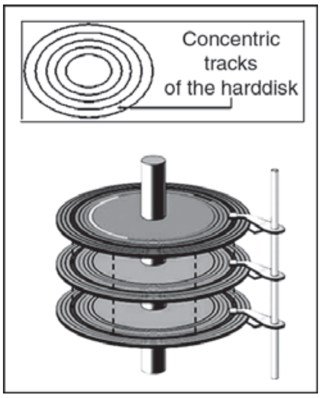

Reading data from the disk means converting the magnetized data to electrical impulses that can be sent to the processor. Writing data to disk is the opposite: sending electrical impulses from the processor to be converted to magnetized spots on the disk. The surface of each disk has concentric tracks on it. The number of tracks per surface varies with the particular type of disk.

HARD DISKS

A hard disk drive (HDD)is a data storage device used for storing and retrieving digital information using rapidly rotating disks (platters) coated with magnetic material. An HDD retains its data even when powered off. Data is read in a random-access manner, meaning individual blocks of data can be stored or retrieved in any order rather than sequentially. An HDD consists of one or more rigid ("hard") rapidly rotating disks (platters) with magnetic heads arranged on a moving actuator arm to read and write data to the surfaces.

A hard disk is a metal platter coated with magnetic oxide that can be magnetized to represent data. Hard disks come in a variety of sizes.

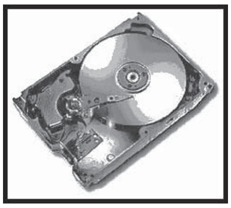

Hard disk for mainframes and minicomputers may be as large as 14 inches in diameter. Several disks can be assembled into a disk pack. There are different types of disk packs, with the number of platters varying by model.

Each disk in the pack has top and bottom surfaces on which to record data. Many disk devices, however, do not record data on the top of the top platter or on the bottom of the bottom platter.

A disk drive is a machine that allows data to be read from a disk or written on a disk. A disk pack is mounted on a disk drive that is a separate unit connected to the computer. Large computers have dozens or ever hundreds of disk drives. In a disk pack all disks rotate at the same time although only one disk is being read or written on at any one time.

The mechanism for reading or writing data on a disk is an access arm; it moves a read/write head into position over a particular track. The read/write head on the end of the access arm hovers just above the track but does not actually touch the surface. When a read/write head does accidentally touch the disk surface, this is called a head crash and all data is destroyed. Data can also be destroyed if a read/write head encounters even minuscule foreign matter on the disk surface. A disk pack has a series of access arms that slip in between the disks in the pack.

Two read/write heads are on each arm, one facing up for the surface above it and one facing down for the surface below it. However, only one read/write head can operate at any one time. In some disk drives the access arms can be retracted; then the disk pack can be removed from the drive. Most disk packs, however, combine the disks, access arms, and read/write heads in a sealed module called a Winchester disk. Winchester disk assemblies are put together in clean rooms so even microscopic dust particles do not get on the disk surface.

Hard disks for personal computers are 5-1/4 inch or 3-1/2 inch disks in sealed modules and even gigabytes are not unusual. Hard disk capacity for personal computers has soared in recent years; capacities of hundreds of megabytes are common and gigabytes are not unusual. Although an individual probably cannot imagine generating enough output-letters, budgets, reports, and so forth-to fill a hard disk, software packages take up a lot of space and can make a dent rather quickly.

Furthermore, graphics images and audio and video files require large file capacities. Perhaps more important than capacity, however, is the convenience of speed. Personal computer users find accessing files on a hard disk is significantly faster and thus more convenient than accessing files on a diskette. Hard disks provide larger and faster secondary storage capabilities than diskettes. Usually hard disks are permanently mounted inside the computer and are not removable like diskettes.

On minicomputers (G) and mainframes (G), hard disks are often called fixed disks. They are also called direct-access storage devices (DASD). Most personal computers have two to four disk drives. The input/output device that transfers data to and from the hard disk is the hard disk drive.

HARD DISKS IN GROUPS

A concept of using several small disks that work together as a unit is called a redundant array of inexpensive disks, or simply RAID. The group of connected disks operates as if it were just one large disk, but it speeds up reading and writing by having multiple access paths. The data file for, say, aircraft factory tools, may be spread across several disks; thus, if the computer is used to look up tools for several workers, the computer need not read the data in turn but instead read them at the same time in parallel.

Furthermore, data security is improved because if a disk fails, the disk system can reconstruct data on an extra disk; thus, computer operations can continue uninterrupted. This is significant data insurance.

HOW DATA IS ORGANIZED ON A DISK

There is more than one way of physically organizing data on a disk. The methods we will consider here are the sector method and the cylinder method.

THE SECTOR METHOD

In the sector method each track is divided into sectors that hold a specific number of characters. Data on the track is accessed by referring to the surface number, track number, and sector number where the data is stored. The sector method is used for diskettes as well as disk packs.

ZONE RECORDING

The fact that a disk is circular presents a problem: The distances around the tracks on the outside of the disk are greater than that of the tracks or the inside.

A given amount of data that takes up 1 inch of a track on the inside of a disk might be spread over several inches on a track near the outside of a disk. This means that the tracks on the outside are not storing data as efficiently. Zone recording involves dividing a disk into zones to take advantage of the storage available on all tracks, by assigning more sectors to tracks in outer zones than to those in inner zones. Since each sector on the disk holds the same amount of data, more sectors mean more data storage than if all tracks had the same number of sectors.

THE CYLINDER METHOD

A way to organize data on a disk pack is the cylinder method. The organization in this case is vertical. The purpose is to reduce the time it takes to move the access arms of a disk pack into position. Once the access arms are in position, they are in the same vertical position on all disk surfaces. To appreciate this, suppose you had an empty disk pack on which you wished to record data.

You might be tempted to record the data horizontally-to start with the first surface, fill track 000, then fill track 001, track 002, and so on, and then move to the second surface and again fill tracks 000, 001, 002, and so forth.

Each new track and new surface, however, would require movement of the access arms, a relatively slow mechanical process. Recording the data vertically, on the other hand, substantially reduces access arm movement. The data is recorded on the tracks that can be accessed by one positioning of the access arms-that is, on one cylinder.

To visualize cylinder organization, pretend a cylindrically shaped item, such as a tin can, were figuratively dropped straight down through all the disks in the disk pack.

All the tracks thus encountered, in the same position on each disk surface, comprise a cylinder. The cylinder method, then, means all tracks of a certain cylinder on a disk pack are lined up one beneath the other, and all the vertical tracks of one cylinder are accessible by the read/write heads with one positioning of the access arms mechanism.

Tracks within a cylinder are numbered according to this vertical perspective: A 20-surface disk pack contains cylinder tracks numbered 0 through 19, top to bottom.

HARD DISK STORAGE CAPACITY

Like diskettes, hard disks must be formatted before they can store information. The storage capacity for hard drives is measured in megabytes. Common sizes for personal computers range from 100MB to 500MB of storage. Each 10MB of storage is equivalent to approximately 5,000 printed pages (with approximately 2,000 characters per page).

REMOVABLE STORAGE: ZIP DISKS

The Zip drive is a medium-capacity removable disk storage system that was introduced by Iomega in late 1994. Originally, Zip disks launched with capacities of 100 MB, but later versions increased this to first 250 MB and then 750 MB. Personal computer users, who never seem to have enough hard disk storage space, may turn to a removable hard disk cartridge. Once full, a removable hard disk cartridge can be replaced with a fresh one.

In effect, a removable cartridge is as portable as a diskette, but the disk cartridge holds much more data. Removable units also are important to businesses concerned with security, because the units can be used during business hours but hidden away during off hours. A disadvantage of a removable hard disk is that it takes longer to access data than a built-in hard drive.

Fig. Zip Disk

The most popular removable disk media is the Zip drive as shown in Figure. Over 100's of millions have been sold, making it the de facto standard. The disk cartridges look like a floppy disk, but are slightly bigger in all dimensions. Older Zip disks hold 100MB, newer ones hold 250MB and cost $8-$10 a piece (Floppies hold 1.4MB and cost around $2). The drive sells for around $80- $125. Many new PCs come with Zip drives built in addition to floppy drives. Zip disks are a great way to store large files and software programmes.

BACKUP

Backup means creating a copy of important programmes and data. To backup diskettes, copy the data from one to the other. Diskettes are frequently used to backup important files on hard drives of personal computers.

DISK CARTRIDGES

Removable disk cartridges are another form of disk storage for personal computers. They offer the storage and fast access of hard disks and the portability of diskettes. They are often used when security is an issue since, when you are done using the computer, the disk cartridge can be removed and locked up leaving no data on the computer.

DATA RECOVERY

Data Recovery is the process of salvaging data from damaged, failed, corrupted, or inaccessible secondary storage media when it cannot be accessed normally. Often the data are being salvaged from storage media such as internal or external hard disk drives, solid-state drives (SSD), USB flash drive, storage tapes, CDs, DVDs, RAID, and other electronics. Recovery may be required due to physical damage to the storage device or logical damage to the file system that prevents it from being mounted by the host operating system.

There does exist certain software programmes that, under special circumstances, can recover data and programmes that have been "lost." Often your operating system (G) will be able to do this, or applications such as Norton Utilities.

OPTICAL DISK STORAGE

In computing and optical disc recording technologies, an optical disc (OD) is a flat, usually circular disc which encodes binary data (bits) in the form of pits (binary value of 0 or off, due to lack of reflection when read) and lands (binary value of 1 or on, due to a reflection when read) on a special material (often aluminium) on one of its flat surfaces.

The explosive growth in storage needs has driven the computer industry to provide cheaper, more compact, and more versatile storage devices with greater capacity. This demanding shopping list is a description of the optical disk, like a CD. The technology works like this: A laser hits a layer of metallic material spread over the surface of a disk.

When data is being entered, heat from the laser produces tiny spots on the disk surface. To read the data, the laser scans the disk, and a lens picks up different light reflections from the various spots. Optical storage technology is categorized according to its read/write capability. Read-only media are recorded on by the manufacturer and can be read from but not written to by the user. Such a disk cannot, obviously, be used for your files, but manufacturers can use it to supply software. Applications software packages sometimes include a dozen diskettes or more; all these could fit on one optical disk with plenty of room to spare. The most prominent optical technology is the CD-ROM, for compact disk read-only memory. CD-ROM has a major advantage over other optical disk designs: The disk format is identical to that of audio compact disks, so the same dust-free manufacturing plants that are now stamping out digital versions of Mozart or Mary Chapin Carpenter can easily convert to producing anything from software to an encyclopaedia. Furthermore, CD-ROM storage is large -up to 660 megabytes per disk, the equivalent of over 400 3-1/2 inch diskettes.

When buying a computer the speed of the CD-ROM drive is advertised using an "X" factor, like 12X, or 24X. This indicates the speed at which the CD can transfer data to the CPU - the higher the X factor, the faster the CD. Modern computers now offer a write CD drive or, CD-RW as an option. CD-RW is a write-once, read-many media. With a CD-RW drive, you can create your own CDs. This offers an inexpensive, convenient, safe way to store large volumes of data such as favourite songs, photographs, etc.

DVDS

DVD (sometimes explained as "digital video disc" or "digital versatile disc") is a digital optical disc storage format, invented and developed by Philips, Sony, Toshiba, and Panasonic in 1995. DVDs offer higher storage capacity than compact discs while having the same dimensions.

Digital Versatile Disk (DVD) drives are now widely available in computers as well as home entertainment centres. DVD-ROM drives can read data, such as stored commercial videos for playing. DVD-RW allow DVDs to be created on a computer.

The DVD is a flat disk, the size of a CD - 4.7 inches diameter and.05 inches thick. Data are stored in a small indentation in a spiral track, just like in the CD. DVD disks are read by a laser beam of shorter wave-length than used by the CD

ROM drives. This allows for smaller indentations and increased storage capacity.

The data layer is only half as thick as in the CD-ROM. This opens the possibility to write data in two layers. The outer gold layer is semi transparent, to allow reading of the underlying silver layer. The laser beam is set to two different intensities, strongest for reading the underlying silver layer.

A 4.7 GB side of a DVD can hold 135 minutes top quality video with 6 track stereo. This requires a transmission rate of 4692 bits per second. The 17 GB disk holds 200 hours top quality music recording.

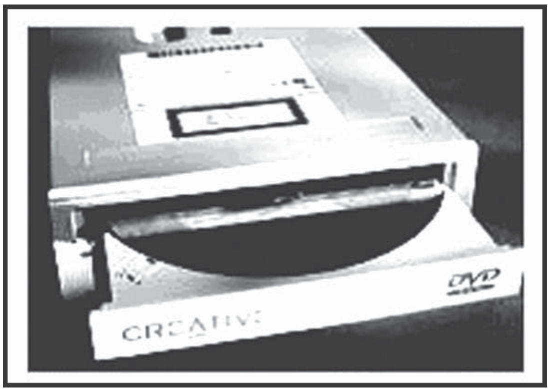

Fig. DVD Disk and Drive

DVD movies are made in two "codes." Region one is USA and Canada, while Europe and Asia is region two. When you play movies, your hardware (MPEG decoder. MGEG is the data coding for movies similar to JPEG for pictures.) must match the DVD region. The movies are made in two formats, each with their own coding. The DVD drives come in 2X, 4X, etc. versions, like the CD-ROM's.

The DVD drives will not replace the magnetic hard disks. The hard disks are being improved as rapidly as DVD, and they definitely offer the fastest seek time and transmission rate (currently 5-10 MB/second). No optic media can keep up with this. But the DVD will undoubtedly gain a place as the successor to the CD ROM and is playing an important role in the blending of computers and entertainment centres.

MAGNETIC TAPE STORAGE

We saved magnetic tape storage for last because it has taken a subordinate role in storage technology. Magnetic tape looks like the tape used in music cassettes plastic tape with a magnetic coating. As in other magnetic media, data is stored as extremely small magnetic spots. Tapes come in a number of forms, including 1/2-inch-wide tape wound on a reel, 1/4-inch- wide tape in data cartridges and cassettes, and tapes that look like ordinary music cassettes but are designed to store data instead of music. The amount of data on a tape is expressed in terms of density, which is the number of characters per inch (cpi) or bytes per inch (bpi) that can be stored on the tape.

The highest-capacity tape is the digital audio tape, or DAT, which uses a different method of recording data. Using a method called helical scan recording, DAT wraps around a rotating read/write head that spins vertically as it moves. This places the data in diagonal bands that run across the tape rather than down its length.

This method produces high density and faster access to data. Two reels are used, a supply reel and a take-up reel. The supply reel, which has the tape with data on it or on which data will be recorded, is the reel that is changed. The take-up reel always stays with the magnetic tape unit. Many cartridges and cassettes have the supply and take-up reels built into the same case. Tape now has a limited role because disk has proved the superior storage medium. Disk data is quite reliable, especially within a sealed module. Furthermore, as we will see, disk data can be accessed directly, as opposed to data on tape, which can be accessed only by passing by all the data ahead of it on the tape. Consequently, the primary role of tape today is as an inexpensive backup medium.

BACKUP SYSTEMS

Although a hard disk is an extremely reliable device, a hard disk drive is subject to electromechanical failures that cause loss of data. Furthermore, data files, particularly those accessed by several users, are subject to errors introduced by users.

There is also the possibility of errors introduced by software. With any method of data storage, a backup system a way of storing data in more than one place to protect it from damage and errors is vital. As we have already noted, magnetic tape is used primarily for backup purposes.

For personal computer users, an easy and inexpensive way to back up a hard disk file is to simply copy it to a diskette whenever it is updated. But this is not practical for a system with many files or many users. Personal computer users have the option of purchasing their own tape backup system, to be used on a regular basis for copying all data from hard disk to a high-capacity tape. Data thus saved can be restored to the hard disk later if needed. A key advantage of a tape backup system is that it can copy the entire hard disk in minutes, saving you the trouble of swapping diskettes in and out of the machine.

A rule of thumb among computer professionals is to estimate disk needs generously and then double that amount. But estimating future needs is rarely easy. Many users, therefore, make later adjustments like adding a removable hard disk cartridge to accommodate expanding storage needs.

BENEFITS OF SECONDARY STORAGE DEVICES

Secondary storage, sometimes called auxiliary storage, is storage separate from the computer itself, where you can store software and data on a semi permanent basis. Secondary storage is necessary because memory, or primary storage, can be used only temporarily.

If you are sharing your computer, you must yield memory to someone else after your programme runs; if you are not sharing your computer, your programmes and data will disappear from memory when you turn off the computer. Howeve, you probably want to store the data you have used or the information you have derived from processing; that is why secondary storage is needed.

Furthermore, memory is limited in size, whereas secondary storage media can store as much data as necessary.

- *Capacity*: Organizations may store the equivalent of a roomful of data on sets of disks that take up less space than a breadbox. A simple diskette for a personal computer holds the equivalent of 500 printed pages, or one book. An optical disk can hold the equivalent of approximately 400 books.
- *Reliability*: Data in secondary storage is basically safe, since secondary storage is physically reliable. Also, it is more difficult for unscrupulous people to tamper

with data on disk than data stored on paper in a file cabinet. Convenience. With the help of a computer, authorized people can locate and access data quickly. Cost. Together the three previous benefits indicate significant savings in storage costs. It is less expensive to store data on tape or disk (the principal means of secondary storage) than to buy and house filing cabinets. Data that is reliable and safe is less expensive to maintain than data subject to errors. But the greatest savings can be found in the speed and convenience of filing and retrieving data.

These benefits apply to all the various secondary storage devices but, as you will see, some devices are better than others. We begin with a look at the various storage media, including those used for personal computers, and then consider what it takes to get data organized and processed.

2

Data Storage Tag

A data storage tag (DST), also sometimes known as an archival tag, is a data logger that uses sensors to record data at predetermined intervals. Data storage tags usually have a large memory size and a long lifetime. Most archival tags are supported by batteries that allow the tag to record positions for several years. Alternatively some tags are solar powered and allow the scientist to set their own interval; this then allows data to be recorded for significantly longer than battery-only powered tags.

OPERATION

Data storage tags can have a variety of sensors; temperature, depth, light, salinity, pressure, pitch and roll, GPS, magnetic and compass. They can be used internally or externally in fish, marine animals or research animals. They are also used in other industries such as the food and beverage industry. At the end of the monitoring period, the loggers can be connected to a computer and the data uploaded for analysis.

DEPLOYMENT

Archival tags archive data to internal memory. Once they are recovered the data is then extracted by the researcher. The tag is generally mounted to the animal either by cutting a slit into the animal, inserting the tag, and sewing the opening closed. Alternatively researchers externally attach tags to animals by running anchor lines through the tag and into the dorsal fin for most "fish" species. For turtles the tag is epoxied to the shell of the turtle.

INFORMATION REPOSITORY

An information repository is an easy way to deploy a secondary tier of data storage that can comprise multiple, networked data storage technologies running on diverse operating systems, where data that no longer needs to be in primary storage is protected, classified according to captured metadata, processed, de-duplicated, and then purged, automatically, based on data service level objectives and requirements. In information repositories, data storage resources are virtualized as composite storage sets and operate as a federated environment. Information repositories were developed to mitigate problems arising from data proliferation and eliminate the need for separately deployed data storage solutions because of the concurrent deployment of diverse storage technologies running diverse operating systems. They feature centralized management for all deployed data storage resources. They are self-contained, support heterogeneous storage resources, support resource management to add, maintain, recycle, and terminate media, track of off-line media, and operate autonomously.

AUTOMATED DATA MANAGEMENT

Since one of the main reasons for the implementation of an Information repository is to reduce the maintenance workload placed on IT staff by traditional data storage systems, information repositories are automated. Automation is accomplished via polices that can process data based on time, events, data age, and data content. Policies manage the following:

- File system space management
- Irrelevant data elimination (mp3, games, etc.)
- Secondary storage resource management

Data is processed according to media type, storage pool, and storage technology. Because information repositories are intended to reduce IT staff workload, they are designed to be easy to deploy and offer configuration flexibility, virtually limitless extensibility, redundancy, and reliable failover.

DATA RECOVERY

Information repositories feature robust, client based data search and recovery capabilities that, based on permissions, enable end users to search the information repository, view information repository contents, including data on off-line media, and recover individual files or multiple files to either their original network computer or another network computer.

FILE SYSTEM

A file system (sometimes written as filesystem) is a method of storing and organizing computer files and their data. Essentially, it organizes these files into a database for the storage, organization, manipulation, and retrieval by the computer's operating system. File systems are used on data storage devices such as hard disks or CD-ROMs to maintain the physical location of the files. Beyond this, they might provide access to data on a file server by acting as clients for a network protocol (e.g., NFS, SMB, or 9P clients), or they may be virtual and exist only as an access method for virtual data (e.g., procfs). It is distinguished from a directory service and registry.

ASPECTS OF FILE SYSTEMS

Most file systems make use of an underlying data storage device that offers access to an array of fixed-size physical sectors, generally a power of 2 in size (512 bytes or 1, 2, or 4 KiB are most common). The file system is responsible for organizing these sectors into files and directories, and keeping track of which sectors belong to which file and which are not being used. Most file systems address data in fixed-sized units called "clusters" or "blocks" which contain a certain number of disk sectors (usually 1-64). This is the smallest amount of disk space that can be allocated to hold a file. However, file systems need not make use of a storage device at all. A file system can be used to organize and represent access to any data, whether it is stored or dynamically generated (e.g., procfs).

FILE NAMES

A file name (or filename) is a name assigned to a file in order to secure storage location in the computer memory. Whether the file system has an underlying storage device or not, file systems typically have directories which associate file names with files, usually by connecting the file name to an index in a file allocation table of some sort, such as the FAT in a DOS file system, or an inode in a Unix-like file system. Directory structures may be flat, or allow hierarchies where directories may contain subdirectories. In some file systems, file names are structured, with special syntax for filename extensions and version numbers. In others, file names are simple strings, and per-file metadata is stored elsewhere.

METADATA

Other bookkeeping information is typically associated with each file within a file system. The length of the data contained in a file may be stored as the number of blocks allocated for the file or as an exact byte count. The time that the file was last modified may be stored as the file's timestamp. Some file systems also store the file creation time, the time it was last accessed, and the time the file's meta-data was changed. (Note that many early PC operating systems did not keep track of file times.) Other information can include the file's device type (e.g., block, character, socket, subdirectory, etc.), its owner user ID and group ID, and its access permission settings (e.g., whether the file is read-only, executable, etc.). Arbitrary attributes can be associated on advanced file systems, such as NTFS, XFS, ext2/ext3, some versions of UFS, and HFS+, using extended file attributes. This feature is implemented in the kernels of Linux, FreeBSD and Mac OS X operating systems, and allows metadata to be associated with the file at the *file system* level. This, for example, could be the author of a document, the character encoding of a plain-text document, or a checksum.

HIERARCHICAL FILE SYSTEMS

The hierarchical file system (not to be confused with Apple's HFS) was an early research interest of Dennis Ritchie of Unix fame; previous implementations were restricted to only a few levels, notably the IBM implementations, even of their early databases like IMS. After the success of Unix, Ritchie extended the file system concept to every object in his later operating system developments, such as Plan 9 and Inferno.

FACILITIES

Traditional file systems offer facilities to create, move and delete both files and directories. They lack facilities to create additional links to a directory (hard links in Unix), rename parent links ("..” in Unix-like OS), and create bidirectional links to files. Traditional file systems also offer facilities to truncate, append to, create, move, delete and in-place modify files. They do not offer facilities to prepend to or truncate from the beginning of a file, let alone arbitrary insertion into or deletion from a file. The operations provided are highly asymmetric and lack the generality to be useful in unexpected contexts. For example, interprocess pipes in

Unix have to be implemented outside of the file system because the pipes concept does not offer truncation from the beginning of files.

SECURE ACCESS

Secure access to basic file system operations can be based on a scheme of access control lists or capabilities. Research has shown access control lists to be difficult to secure properly, which is why research operating systems tend to use capabilities. Commercial file systems still use access control lists.

TYPES OF FILE SYSTEMS

File system types can be classified into disk file systems, network file systems and special purpose file systems.

DISK FILE SYSTEMS

A *disk file system* is a file system designed for the storage of files on a data storage device, most commonly a disk drive, which might be directly or indirectly connected to the computer. Examples of disk file systems include FAT (FAT12, FAT16, FAT32, exFAT), NTFS, HFS and HFS+, HPFS, UFS, ext2, ext3, ext4, btrfs, ISO 9660, ODS-5, Veritas File System, VMFS, ZFS, ReiserFS and UDF. Some disk file systems are journaling file systems or versioning file systems.

OPTICAL DISCS

ISO 9660 and Universal Disk Format (UDF) are the two most common formats that target Compact Discs, DVDs and Blu-ray discs. Mount Rainier is a newer extension to

UDF supported by Linux 2.6 series and Windows Vista that facilitates rewriting to DVDs in the same fashion as has been possible with floppy disks.

FLASH FILE SYSTEMS

A *flash file system* is a file system designed for storing files on flash memory devices. These are becoming more prevalent as the number of mobile devices is increasing, and the capacity of flash memories increase. While a disk file system can be used on a flash device, this is suboptimal for several reasons:

- Erasing blocks: Flash memory blocks have to be explicitly erased before they can be rewritten. The time taken to erase blocks can be significant, thus it is beneficial to erase unused blocks while the device is idle.
- Random access: Disk file systems are optimized to avoid disk seeks whenever possible, due to the high cost of seeking. Flash memory devices impose no seek latency.
- Wear levelling: Flash memory devices tend to wear out when a single block is repeatedly overwritten; flash file systems are designed to spread out writes evenly.

Log-structured file systems have many of the desirable properties for a flash file system. Such file systems include JFFS2 and YAFFS.

NO NEED FOR DEFRAGMENTING

Because flash media effectively has zero seek time, defragmentation of flash media is unnecessary and has no performance benefit. Instead defragmentation is detrimental to the life of the media since it wears out the data storage cells with unnecessary writing.

MEDIA IS TYPICALLY PPRE-FORMATTED

Due to the problems of limited write cycles per data cell, it is generally not recommended to format flash storage devices since it reduces the life of all cells on the media. Flash devices are typically sold with a common type of file system already created on the media, to remove the need for formatting.

TAPE FILE SYSTEMS

A *tape file system* is a file system and tape format designed to store files on tape in a self-describing form. Magnetic tapes are sequential storage media with significantly longer random data access times than disks, posing challenges to the creation and efficient management of a general-purpose file system. In a disk file system there is typically a master file directory, and a map of used and free data regions. Any file additions, changes, or removals require updating the directory and the used/free maps. Random access to data regions is measured in milliseconds so this system works well for disks. However, tape requires linear motion to wind and unwind potentially very long reels of media, and this tape motion may take several seconds to several minutes to move the read/write head from one end of the tape to the other. Consequently, a master file directory and usage map can extremely slow and inefficient with tape. Writing typically involves reading the block usage map to find free blocks for writing, updating the usage map and directory to add the data, and then advancing the tape to write the data in the correct spot. Each additional file write requires updating the map and directory and writing the data, which may take several seconds to occur for each file.

Tape file systems instead typically allow for the file directory to be spread across the tape intermixed with the data, referred to as *streaming*, so that time-consuming and repeated tape motions are not required to write new data.

However a side effect of this design is that reading the file directory of a tape usually requires scanning the entire tape to read all the scattered directory entries. Most data archiving software that works with tape storage will store a local copy of the tape catalog on a disk file system, so that adding files to a tape can be done quickly without having to rescan the tape media. The local tape catalog copy is usually discarded if not used for a specified period of time, at which point the tape must be re-scanned if it is to be used in the future.

IBM has recently announced a new file system for tape called the Linear Tape File System. The IBM implementation of this file system has been released as the open-source IBM Long Term File System product.

TAPE FORMATTING

Writing data to a tape is often a significantly time-consuming process that may take several hours. Similarly, completely erasing or formatting a tape can also take several hours. With many data tape technologies it is not necessary to format the tape before over-writing new data to the tape. This is due to the inherently destructive nature of overwriting data on sequential media. Because of the time it can take to format a tape, typically tapes are pre-formatted so that the tape user does not need to spend time preparing each new tape for use. All that is usually necessary is to write an identifying media label to the tape before use, and even this can be automatically written by software when a new tape is used for the first time.

TAPE BOOTING

Tapes can sometimes be used to boot a computer, but the long random data access delays can make this a very slow and tedious process. The boot-file storage can be optimized so that programme data is stored in the order required for use, with minimal need for seeking. Typically tape booting is only used for infrequent special purposes such as operating system installation or for system recovery/restore.

DATABASE FILE SYSTEMS

A recent concept for file management is the idea of a database-based file system. Instead of, or in addition to, hierarchical structured management, files are identified by their characteristics, like type of file, topic, author, or similar metadata.

TRANSACTIONAL FILE SYSTEMS

Some programmes need to update multiple files "all at once." For example, a software installation may write programme binaries, libraries, and configuration files. If the software installation fails, the programme may be unusable. If the installation is upgrading a key system utility, such as the command shell, the entire system may be left in an unusable state. Transaction processing introduces the isolation guarantee, which states that operations within a transaction are hidden from other threads on the system until the transaction commits, and that interfering operations on the system will be properly serialized with the transaction. Transactions also provide the atomicity guarantee, that operations inside of a transaction are either all committed, or the transaction can be aborted and the system discards all of its partial results. This means that if there is a crash or power failure, after recovery, the stored state will be consistent. Either the software will be completely installed or the failed installation will be completely rolled back, but an unusable partial install will not be left on the system.

Windows, beginning with Vista, added transaction support to NTFS, abbreviated TxF. TxF is the only commercial implementation of a transactional file system, as transactional file systems are difficult to implement correctly in practice. There are a number of research prototypes of transactional file systems for UNIX systems, including the Valor file system, Amino, LFS, and a transactional ext3 file system on the TxOS kernel, as well as transactional file systems targeting embedded systems, such as TFFS.

Ensuring consistency across multiple file system operations is difficult, if not impossible, without file system transactions. File locking can be used as a concurrency control mechanism for individual files, but it typically does not protect the directory structure or file metadata. For instance, file locking cannot prevent TOCTTOU race conditions on symbolic links. File locking also cannot automatically roll back a failed operation, such as a software upgrade; this requires atomicity. Journaling file systems are one technique used to introduce transaction-level consistency to file system structures. Journal transactions are not exposed to programmes as part of the OS API; they are only used internally to ensure consistency at the granularity of a single system call.

NETWORK FILE SYSTEMS

A network file system is a file system that acts as a client for a remote file access protocol, providing access to files on a server. Examples of network file systems include clients for the NFS, AFS, SMB protocols, and file-system-like clients for FTP and WebDAV.

SHARED DISK FILE SYSTEMS

A shared disk file system is one in which a number of machines (usually servers) all have access to the same external disk subsystem (usually a SAN). The file system arbitrates access to that subsystem, preventing write collisions. Examples include GFS from Red Hat, GPFS from IBM, and SFS from DataPlow.

SPECIAL PURPOSE FILE SYSTEMS

A special purpose file system is basically any file system that is not a disk file system or network file system. This includes systems where the files are arranged dynamically by software, intended for such purposes as communication between computer processes or temporary file space. Special purpose file systems are most commonly used by file-centric operating systems such as Unix. Examples include the procfs (/proc) file system used by some Unix variants, which grants access to information about processes and other operating system features.

Deep space science exploration craft, like Voyager I and II used digital tape-based special file systems. Most modern space exploration craft like Cassini-Huygens used Real-time operating system file systems or RTOS influenced file systems. The Mars Rovers are one such example of an RTOS file system, important in this case because they are implemented in flash memory.

FILE SYSTEMS AND OPERATING SYSTEMS

Most operating systems provide a file system, as a file system is an integral part of any modern operating system. Early microcomputer operating systems' only real task was file management — a fact reflected in their names. Some early operating systems had a separate component for handling file systems which was called a disk operating system. On some microcomputers, the disk operating system was loaded separately from the rest of the operating system. On early operating systems, there was usually support for only one, native, unnamed file system; for example, CP/M supports only its own file system, which might be called "CP/M file system" if needed, but which didn't bear any official name at all. Because of this, there needs to be an interface provided by the operating system software between the user and the file system. This interface can be textual (such as provided by a command line interface, such as the Unix shell, or OpenVMS DCL) or graphical (such as provided by a graphical user interface, such as file browsers). If graphical, the metaphor of the *folder*, containing documents, other files, and nested folders is often used.

FLAT FILE SYSTEMS

In a flat file system, there are no subdirectories—everything is stored at the same (root) level on the media, be it a hard disk, floppy disk, etc. While simple, this system rapidly becomes inefficient as the number of files grows, and makes it difficult for users to organize data into related groups. Like many small systems before it, the original Apple Macintosh featured a flat file system, called Macintosh File System. Its version of Mac OS was unusual in that the file management software (Macintosh Finder) created the illusion of a partially hierarchical filing system on top of EMFS. This structure meant that every file on a disk had to have a unique name, even if it appeared to be in a separate folder. MFS was quickly replaced with Hierarchical File System, which supported real directories. A recent addition to the flat file system family is Amazon's S3, a remote storage service, which is intentionally simplistic to allow users the ability to customize how their data is stored. The only constructs are buckets (imagine a disk drive of unlimited size) and objects (similar, but not identical to the standard concept of a file). Advanced file management is allowed by being able to use nearly any character (including '/') in the object's name, and the ability to select subsets of the bucket's content based on identical prefixes.

FILE SYSTEMS UNDER UNIX-LIKE OPERATING SYSTEMS

Unix-like operating systems create a virtual file system, which makes all the files on all the devices appear to exist in a single hierarchy. This means,in those systems, there is one root directory, and every file existing on the system is located under it somewhere. Unix-like systems can use a RAM disk or network shared resource as its root directory. Unix-like systems assign a device name to each device, but this is not how the files on that device are accessed. Instead, to gain access to files on another device, the operating system must first be informed where in the directory tree those files should appear. This process is called mounting a file system. For example, to access the files on a CD-ROM, one must tell the operating system "Take the file system from this CD-ROM and make it appear under such-and-such directory". The directory given to the operating system is called the *mount point* – it might, for example, be /media. The /media directory exists on many Unix systems (as specified in the Filesystem Hierarchy Standard) and is intended specifically for use as a mount point for removable media such as CDs, DVDs, USB drives or floppy disks. It may be empty, or it may contain subdirectories for mounting individual devices. Generally, only the administrator (i.e. root user) may authorize the mounting of file systems.

Unix-like operating systems often include software and tools that assist in the mounting process and provide it new functionality. Some of these strategies have been coined "auto-mounting" as a reflection of their purpose.

1. In many situations, file systems other than the root need to be available as soon as the operating system has booted. All Unix-like systems therefore provide a facility for mounting file systems at boot time. System administrators define these file systems in the configuration file fstab (*vfstab* in Solaris), which also indicates options and mount points.
2. In some situations, there is no need to mount certain file systems at boot time, although their use may be desired thereafter. There are some utilities for Unix-like systems that allow the mounting of predefined file systems upon demand.
3. Removable media have become very common with microcomputer platforms. They allow programmes and data to be transferred between machines without a physical connection. Common examples include USB flash drives, CD-ROMs, and DVDs. Utilities have therefore been developed to detect the presence and availability of a medium and then mount that medium without any user intervention.
4. Progressive Unix-like systems have also introduced a concept called supermounting; see, for example, the Linux supermount-ng project. For example, a floppy disk that has been supermounted can be physically removed from the system. Under normal circumstances, the disk should have been synchronized and then unmounted before its removal. Provided synchronization has occurred, a different disk can be inserted into the drive. The system automatically notices that the disk has changed and

updates the mount point contents to reflect the new medium. Similar functionality is found on Windows machines.

5. An automounter will automatically mount a file system when a reference is made to the directory atop which it should be mounted. This is usually used for file systems on network servers, rather than relying on events such as the insertion of media, as would be appropriate for removable media.

FILE SYSTEMS UNDER LINUX

Linux supports many different file systems, but common choices for the system disk include the ext* family (such as ext2, ext3 and ext4), XFS, JFS, ReiserFS and btrfs.

FILE SYSTEMS UNDER SOLARIS

The Sun Microsystems Solaris operating system in earlier releases defaulted to (non-journaled or non-logging) UFS for bootable and supplementary file systems. Solaris defaulted to, supported, and extended UFS. Support for other file systems and significant enhancements were added over time, including Veritas Software Corp. (Journaling) VxFS, Sun Microsystems (Clustering) QFS, Sun Microsystems (Journaling) UFS, and Sun Microsystems (open source, poolable, 128 bit compressible, and error-correcting) ZFS. Kernel extensions were added to Solaris to allow for bootable Veritas VxFS operation. Logging or Journaling was added to UFS in Sun's Solaris 7. Releases of Solaris 10, Solaris Express, OpenSolaris, and other open source variants of the Solaris operating system later supported bootable ZFS.

Logical Volume Management allows for spanning a file system across multiple devices for the purpose of adding redundancy, capacity, and/or throughput. Legacy environments in Solaris may use Solaris Volume Manager (formerly known as Solstice DiskSuite.) Multiple operating systems (including Solaris) may use Veritas Volume Manager. Modern Solaris based operating systems eclipse the need for Volume Management through leveraging virtual storage pools in ZFS.

FILE SYSTEMS UNDER MAC OS X

Mac OS X uses a file system that it inherited from classic Mac OS called HFS Plus, sometimes called *Mac OS Extended*. HFS Plus is a metadata-rich and case preserving file system. Due to the Unix roots of Mac OS X, Unix permissions were added to HFS Plus. Later versions of HFS Plus added journaling to prevent corruption of the file system structure and introduced a number of optimizations to the allocation algorithms in an attempt to defragment files automatically without requiring an external defragmenter. Filenames can be up to 255 characters. HFS Plus uses Unicode to store filenames. On Mac OS X, the filetype can come from the type code, stored in file's metadata, or the filename. HFS Plus has three kinds of links: Unix-style hard links, Unix-style symbolic links and aliases. Aliases are designed to maintain a link to their original file even if they are moved or renamed; they are not interpreted by the file system itself, but by the File Manager code in userland. Mac OS X also supports the UFS file system, derived from the BSD Unix Fast File System via NeXTSTEP. However, as of Mac OS X 10.5 (Leopard), Mac OS X can no longer be installed on a UFS volume, nor can a pre-Leopard system installed on a UFS volume be upgraded to Leopard. Newer versions Mac OS X are capable of reading and writing to the legacy FAT file systems(16 & 32). They are capable of reading, but not writing to the NTFS file system. Third party software is still necessary to write to the NTFS file system under Snow Leopard 10.6.4.

FILE SYSTEMS UNDER PLAN 9 FROM BELL LABS

Plan 9 from Bell Labs was originally designed to extend some of Unix's good points, and to introduce some new ideas of its own while fixing the shortcomings of Unix. With respect to file systems, the Unix system of treating things as files was continued, but in Plan 9, *everything* is treated as a file, and accessed as a file would be (i.e., no ioctl or mmap). Perhaps surprisingly, while the file interface is made universal it is also simplified considerably: symlinks, hard links and suid are made obsolete, and an atomic create/open operation is introduced. More importantly the set of file operations becomes well defined and subversions of this like ioctl are eliminated. Secondly, the underlying 9P protocol was used to remove the difference between local and remote files (except for a possible difference in latency or in throughput). This has the advantage that a device or devices, represented by files, on a remote computer could be used as though it were the local computer's own device(s). This means that under Plan 9, multiple file servers provide access to devices, classing them as file systems. Servers for "synthetic" file systems can also run in user space bringing many of the advantages of micro kernel systems while maintaining the simplicity of the system.

Everything on a Plan 9 system has an abstraction as a file; networking, graphics, debugging, authentication, capabilities, encryption, and other services are accessed via I-O operations on file descriptors. For example, this allows the use of the IP stack of a gateway machine without need of NAT, or provides a network-transparent window system without the need of any extra code. Another example: a Plan-9 application receives FTP service by opening an FTP site. The ftpfs server handles the open by essentially mounting the remote FTP site as part of the local file system. With ftpfs as an intermediary, the application can now use the usual file-system operations to access the FTP site as if it were part of the local file system. A further example is the mail system which uses file servers that synthesize virtual files and directories to represent a user mailbox as /mail/fs/mbox. The wikifs provides a file system interface to a wiki. These file systems are organized with the help of private, per-process namespaces, allowing each process to have a different view of the many file systems that provide resources in a distributed system. The Inferno operating system shares these concepts with Plan 9.

FILE SYSTEMS UNDER MICROSOFT WINDOWS

Windows makes use of the FAT and NTFS file systems. Unlike many other operating systems, Windows uses a *drive letter* abstraction at the user level to distinguish one disk or partition from another. For example, the path C:\WINDOWS represents a directory WINDOWS on the partition represented by the letter C. The C drive is most commonly used for the primary hard disk partition, on which Windows is usually installed and from which it boots. This "tradition" has become so firmly ingrained that bugs came about in older applications which made assumptions that the drive that the operating system was installed on was C. The tradition of using "C" for the drive letter can be traced to MS-DOS, where the letters A and B were reserved for up to two floppy disk drives. This in turn derived from CP/M in the 1970s, and ultimately from IBM's CP/CMS of 1967. Network drives may also be mapped to drive letters.

FAT

The File Allocation Table (FAT) filing system, supported by all versions of Microsoft Windows, was an evolution of that used in Microsoft's earlier operating system (MS-DOS which in turn was based on 86-DOS). FAT ultimately traces its roots back to the short-lived M-DOS project and Standalone disk BASIC before it. Over the years various features have been added to it, inspired by similar features found on file systems used by operating systems such as Unix. Older versions of the FAT file system (FAT12 and FAT16) had file name length limits, a limit on the number of entries in the root directory of the file system and had restrictions on the maximum size of FAT-formatted disks or partitions. Specifically, FAT12 and FAT16 had a limit of 8 characters for the file name, and 3 characters for the extension (such as .exe). This is commonly referred to as the 8.3 filename limit. VFAT, which was an extension to

FAT12 and FAT16 introduced in Windows NT 3.5 and subsequently included in Windows 95, allowed long file names (LFN). FAT32 also addressed many of the limits in FAT12 and FAT16, but remains limited compared to NTFS. exFAT (also known as FAT64) is the newest iteration of FAT, with certain advantages over NTFS with regards to file system overhead. exFAT is only compatible with newer Windows systems, such as Windows 2003, Windows Vista, Windows 2008, Windows 7 and more recently, support has been added for WinXP.

NTFS

NTFS, introduced with the Windows NT operating system, allowed ACL-based permission control. Hard links, multiple file streams, attribute indexing, quota tracking, sparse files, encryption, compression, reparse points (directories working as mount-points for other file systems, symlinks, junctions, remote storage links) are also supported, though not all these features are well-documented.

OTHER FILE SYSTEMS

- The Prospero File System is a file system based on the Virtual System Model. The system was created by Dr. B. Clifford Neuman of the Information Sciences Institute at the University of Southern California.
- RSRE FLEX file system - written in ALGOL 68
- The file system of the Michigan Terminal System (MTS) is interesting because: (i) it provides “line files” where record lengths and line numbers are associated as metadata with each record in the file, lines can

be added, replaced, updated with the same or different length records, and deleted anywhere in the file without the need to read and rewrite the entire file; (ii) using programme keys files may be shared or permitted to commands and programmes in addition to users and groups; and (iii) there is a comprehensive file locking mechanism that protects both the file's data and its metadata.

DATA DEDUPLICATION

In computing, data deduplication is a specialized data compression technique for eliminating coarse-grained redundant data, typically to improve storage utilization. In the deduplication process, duplicate data is deleted, leaving only one copy of the data to be stored, along with references to the unique copy of data. Deduplication is able to reduce the required storage capacity since only the unique data is stored.

Depending on the type of deduplication, redundant files may be reduced, or even portions of files or other data that are similar can also be removed. As a simple example of file based deduplication, a typical email system might contain 100 instances of the same one megabyte (MB) file attachment. If the email platform is backed up or archived, all 100 instances are saved, requiring 100 MB storage space. With data deduplication, only one instance of the attachment is actually stored; each subsequent instance is just referenced back to the one saved copy. In this example, the deduplication ratio is roughly 100 to 1. Different applications and data types naturally have different levels of data redundancy.

Backup applications generally benefit the most from deduplication due to the nature of repeated full backups of an existing file system. Like a traditional stream-based dictionary coder, deduplication identifies identical sections of data and replaces them by references to a single copy of the data. However, whereas standard file compression tools like LZ77 and LZ78 identify short repeated substrings inside single files, the focus of data deduplication is to take a very large volume of data and identify large sections - such as entire files or large sections of files - that are identical, and store only one copy of it. This copy may be additionally compressed by single-file compression techniques.

BENEFITS

Data deduplication reduces the amount of storage needed for a given set of files. It is most effective in applications where many copies of very similar or even identical data are stored on a single disk—a surprisingly common scenario.

One very good application for data deduplication is in backups. Most data in a given backup isn't changed from the previous backup; common backup systems try to exploit this by omitting (or hard linking) files that haven't changed or storing differences between files. Neither approach captures all redundancies, however. Hard linking does not help with large files that have only changed in small ways, such as an email database; differences only find redundancies in adjacent versions of a single file (consider a section that was deleted and later added in again, or a logo image included in many documents). Data deduplication allows a backup programme to essentially just copy files onto a backup disk without trying to omit or difference them; the storage system itself will ensure that only one copy of the data ends up on the disk, no matter how many versions ago the duplicate data occurred or even if the similarity appears in a different file. This can reduce backup storage requirements by 90% or more—making it feasible to retain data for months on a fast, readily accessible backup medium. It also reduces the data that must be sent across a WAN for remote backups, replication, and disaster recovery.

Data deduplication is also especially effective when used with virtual servers, allowing the nominally separate system files for each virtual server to be coalesced into a single storage space. (At the same time, if a given server customizes a file, deduplication will not change the files on the other servers—something that alternatives like hard links or shared disks do not offer.) Backing up or making duplicate copies of virtual environments is similarly improved.

By reducing the amount of storage needed, deduplication can save other resources, too: the energy use, physical volume, cooling needs, and carbon footprint needed to store the data are all reduced. Less hardware needs to be purchased, recycled, and replaced, further lowering costs.

DEDUPLICATION OVERVIEW

Deduplication may occur “in-line”, as data is flowing, or “post-process” after it has been written. With post-process deduplication, new data is first stored on the storage device and then a process at a later time will analyze the data looking for duplication. The benefit is that there is no need to wait for the hash calculations and lookup to be completed before storing the data thereby ensuring that store performance is not degraded. Implementations offering policy-based operation can give users the ability to defer optimization on "active" files, or to process files based on type and location. One potential drawback is that you may unnecessarily store duplicate data for a short time which is an issue if the storage system is near full capacity. This is the process where the deduplication hash calculations are created on the target device as the data enters the device in real time. If the device spots a block that it already stored on the system it does not store the new block, just references to the existing block. The benefit of in-line deduplication over post-process deduplication is that it requires less storage as data is not duplicated. On the negative side, it is frequently argued that because hash calculations and lookups takes so long, it can mean that the data ingestion can be slower thereby reducing the backup throughput of the device. However, certain vendors with in-line deduplication have demonstrated equipment with similar performance to their post-process deduplication counterparts. Post-process and in-line deduplication methods are often heavily debated.

Deduplication can occur close to where data is created, which is often referred to as "source deduplication." It can occur close to where the data is stored, which is commonly called "target deduplication."

When describing deduplication for backup architectures, it is common to hear two terms: source deduplication and target deduplication. Source deduplication ensures that data on the data source is deduplicated. This generally takes place directly within a file-system. The file system will periodically scan new files creating hashes and compare them to hashes of existing files. When files with same hashes are found then the file copy is removed and the new file points to the old file. Unlike hard links however, duplicated files are considered to be separate entities and if one of the duplicated files is later modified, then using a system called Copy-on-write a copy of that file or changed block is created. The deduplication process is transparent to the users and backup applications. Backing up a deduplicated filesystem will often cause duplication to occur resulting in the backups being bigger than the source data.

Target deduplication is the process of removing duplicates of data in the secondary store. Generally this will be a backup store such as a data repository or a virtual tape library. There are three different ways of performing the deduplication process.

One of the most common forms of data deduplication implementations work by comparing chunks of data to detect duplicates. For that to happen, each chunk of data is assigned an identification, calculated by the software, typically using cryptographic hash functions. In many implementations, the assumption is made that if the identification is identical, the data is identical, even though this cannot be true in all cases due to the pigeonhole principle; other implementations do not assume that two blocks of data with the same identifier are identical, but actually verify that data with the same identification is identical. If the software either assumes that a given identification already exists in the deduplication namespace or actually verifies the identity of the two blocks of data, depending on the implementation, then it will replace that duplicate chunk with a link. Once the data has been deduplicated, upon read back of the file, wherever a link is found, the system simply replaces that link with the referenced data chunk. The de-duplication process is intended to be transparent to end users and applications.

Between commercial deduplication implementations, technology varies primarily in chunking method and in architecture. In some systems, chunks are defined by physical layer constraints (e.g. 4KB block size in WAFL). In some systems only complete files are compared, which is called Single Instance Storage or SIS. The most intelligent (but CPU intensive) method to chunking is generally considered to be sliding-block. In sliding block, a window is passed along the file stream to seek out more naturally occurring internal file boundaries.

This is the process where the deduplication hash calculations are initially created on the source (client) machines. Files that have identical hashes to files already in the target device are not sent, the target device just creates appropriate internal links to reference the duplicated data. The benefit of this is that it avoids data being unnecessarily sent across the network thereby reducing traffic load.

By definition, primary storage systems are designed for optimal performance, rather than lowest possible cost. The design criteria for these systems is to increase performance, at the expense of other considerations. Moreover, primary storage systems are much less tolerant of any operation that can negatively impact performance. Also by definition, secondary storage systems contain primarily duplicate, or secondary copies of data. These copies of data are typically not used for actual production operations and as a result are more tolerant of some performance degradation, in exchange for increased efficiency.

To date, data deduplication has predominantly been used with secondary storage systems. The reasons for this are two-fold. First, data deduplication requires overhead to discover and remove the duplicate data. In primary storage systems, this overhead may impact performance. The second reason why deduplication is applied to secondary data, is that secondary data tends to have more duplicate data. Backup application in particular commonly generate significant portions of duplicate data over time. Data deduplication has been deployed successfully with primary storage in some cases where the system design does not require significant overhead, or impact performance.

DRAWBACKS AND CONCERNS

Whenever data is transformed, concerns arise about potential loss of data. By definition, data deduplication systems store data differently from how it was written. As a result, users are concerned with the integrity of their data. The various methods of deduplicating data all employ slightly different techniques. However, the integrity of the data will ultimately depend upon the design of the deduplicating system, and the quality used to implement the algorithms. As the technology has matured over the past decade, the integrity of most of the major products has been well proven.

One method for deduplicating data relies on the use of cryptographic hash functions to identify duplicate segments of data. If two different pieces of information generate the same hash value, this is known as a collision. The probability of a collision depends upon the hash function used, and although the probabilities are small, they are always non zero. Thus, the concern arises that data corruption can occur if a hash collision occurs, and additional means of verification are not used to verify whether there is a difference in data, or not. Both in-line and post-process architectures may offer bit-for-bit validation of original data for guaranteed data integrity. The hash functions used include standards such as SHA-1, SHA-256 and others.

These provide a far lower probability of data loss than the risk of an undetected/uncorrected hardware error in most cases and can be in the order of 0.0000000000000000000000000000000000000000000000000000000000013% ($1.3 \cdot 10^{-49}$) per petabyte (1,000 terabyte) of data. Some cite the computational resource intensity of the process as a drawback of data deduplication. However, this is rarely an issue for stand-alone devices or appliances, as the computation is completely offloaded from other systems. This can be an issue when the deduplication is embedded within devices providing other services.

To improve performance, many systems utilize weak and strong hashes. Weak hashes are much faster to calculate but there is a greater risk of a hash collision. Systems that utilize weak hashes will subsequently calculate a strong hash and will use it as the determining factor to whether it is actually the same data or not. Note that the system overhead associated with calculating and looking up hash values is primarily a function of the deduplication workflow. The reconstitution of files does not require this processing and any incremental performance penalty associated with re-assembly of data chunks is unlikely to impact application performance.

Another area of concern with deduplication is the related effect on snapshots, backup, and archival, especially where deduplication is applied against primary storage (for example inside a NAS filer). Reading files out of a storage device causes full rehydration of the files, so any secondary copy of the data set is likely to be larger than the primary copy. In terms of snapshots, if a file is snapshotted prior to deduplication, the post-deduplication snapshot will preserve the entire original file. This means that although storage capacity for primary file copies will shrink, capacity required for snapshots may expand dramatically.

Another concern is the effect of compression and encryption. Although deduplication is a version of compression, it works in tension with traditional compression. Deduplication achieves better efficiency against smaller data chunks, whereas compression achieves better efficiency against larger chunks. The goal of encryption is to eliminate any discernible patterns in the data. Thus encrypted data will have 0% gain from deduplication, even though the underlying data may be redundant.

Scaling has also been a challenge for dedupe systems because the hash table or dedupe namespace needs to be shared across storage devices. If there are multiple disk backup devices in an infrastructure with discrete dedupe namespaces, then space efficiency is adversely affected. A namespace shared across devices - called Global Dedupe - preserves space efficiency, but is technically challenging from a reliability and performance perspective.

Deduplication ultimately reduces redundancy. If this was not expected and planned for, this may ruin the underlying reliability of the system. (Compare this, for example, to the LOCKSS storage architecture that achieves reliability through multiple copies of data.)

3

Data Storage in Hard Disk Drive

A hard disk drive (HDD) is a non-volatile, random access device for digital data. It features rotating rigid platters on a motor-driven spindle within a protective enclosure. Data is magnetically read from and written to the platter by read/write heads that float on a film of air above the platters.

Introduced by IBM in 1956, hard disk drives have fallen in cost and physical size over the years while dramatically increasing in capacity. Hard disk drives have been the dominant device for secondary storage of data in general purpose computers since the early 1960s. They have maintained this position because advances in their areal recording density have kept pace with the requirements for secondary storage. Today's HDDs operate on high-speed serial interfaces; i.e., serial ATA (SATA) or serial attached SCSI (SAS).

HISTORY

Hard disk drives were introduced in 1956 as data storage for an IBM accounting computer and were developed for use with general purpose mainframe and mini computers.

Driven by areal density doubling every two to four years since their invention, HDDs have changed in many ways, a few highlights include:

- Capacity per HDD increasing from 3.75 megabytes to greater than 1 terabyte, a greater than 270 thousand to 1 improvement.
- Size of HDD decreasing from 87.9 cubic feet (a double wide refrigerator) to 0.002 cubic feet (2½-inch form factor, a pack of cards), a greater than 44 thousand to 1 improvement.
- Price decreasing from about $15,000 per megabyte to less than $0.0001 per megabyte ($100/1 terabyte), a greater than 150 million to 1 improvement.
- Average access time decreasing from greater than 0.1 second to a few thousandths of a second, a greater than 40 to 1 improvement.
- Market application expanding from general purpose computers to most computing applications including consumer applications.

TECHNOLOGY

MAGNETIC RECORDING

HDDs record data by magnetizing ferromagnetic material directionally. Sequential changes in the direction of magnetization represent patterns of binary data bits. The data are read from the disk by detecting the transitions in magnetization and decoding the originally written data. Different encoding schemes, such as Modified Frequency Modulation, group code recording, run-length limited encoding, and others are used.

A typical HDD design consists of a spindle that holds flat circular disks called platters, onto which the data are recorded. The platters are made from a non-magnetic material, usually aluminum alloy or glass, and are coated with a shallow layer of magnetic material typically 10–20 nm in depth, with an outer layer of carbon for protection. For reference, standard copy paper is 0.07–0.18 millimetre (70,000–180,000 nm).

The platters are spun at speeds varying from 3,000 RPM in energy-efficient portable devices, to 15,000 RPM for high performance servers. Information is written to, and read from a platter as it rotates past devices called read-and-write heads that operate very close (tens of nanometers in new drives) over the magnetic surface. The read-and-write head is used to detect and modify the magnetization of the material immediately under it. In modern drives there is one head for each magnetic platter surface on the spindle, mounted on a common arm. An actuator arm (or access arm) moves the heads on an arc (roughly radially) across the platters as they spin, allowing each head to access almost the entire surface of the platter as it spins. The arm is moved using a voice coil actuator or in some older designs a stepper motor. The magnetic surface of each platter is conceptually divided into many small sub-micrometer-sized magnetic regions referred to as magnetic domains. In older disk designs the regions were oriented horizontally and parallel to the disk surface, but beginning about 2005, the orientation was changed to perpendicular to allow for closer magnetic domain spacing. Due to the polycrystalline nature of the magnetic material each of these magnetic regions is composed of a few hundred magnetic grains. Magnetic grains are typically 10 nm in size and each form a single magnetic domain. Each magnetic region in total forms a magnetic dipole which generates a magnetic field.

For reliable storage of data, the recording material needs to resist self-demagnetization, which occurs when the magnetic domains repel each other. Magnetic domains written too densely together to a weakly magnetizable material will degrade over time due to physical rotation of one or more domains to cancel out these forces. The domains rotate sideways to a halfway position that weakens the readability of the domain and relieves the magnetic stresses. Older hard disks used iron(III) oxide as the magnetic material, but current disks use a cobalt-based alloy.

A write head magnetizes a region by generating a strong local magnetic field. Early HDDs used an electromagnet both to magnetize the region and to then read its magnetic field by using electromagnetic induction. Later versions of inductive heads included metal in Gap (MIG) heads and thin film heads. As data density increased, read heads using magnetoresistance (MR) came into use; the electrical resistance of the head changed according to the strength of the magnetism from the platter. Later development made use of spintronics; in these heads, the magnetoresistive effect was much greater than in earlier types, and was dubbed "giant" magnetoresistance (GMR). In today's heads, the read and write elements are separate, but in close proximity, on the head portion of an actuator arm. The read element is typically magneto-resistive while the write element is typically thin-film inductive.

The heads are kept from contacting the platter surface by the air that is extremely close to the platter; that air moves at or near the platter speed. The record and playback head are mounted on a block called a slider, and the surface next to the platter is shaped to keep it just barely out of contact. This forms a type of air bearing.

In modern drives, the small size of the magnetic regions creates the danger that their magnetic state might be lost because of thermal effects. To counter this, the platters are coated with two parallel magnetic layers, separated by a 3-atom layer of the non-magnetic element ruthenium, and the two layers are magnetized in opposite orientation, thus reinforcing each other. Another technology used to overcome thermal effects to allow greater recording densities is perpendicular recording, first shipped in 2005, and as of 2007 the technology was used in many HDDs.

COMPONENTS

A typical hard disk drive has two electric motors; a disk motor to spin the disks and an actuator (motor) to position the read/write head assembly across the spinning disks. The disk motor has an external rotor attached to the disks; the stator windings are fixed in place. Opposite the actuator at the end of the head support arm is the read-write head (near center in photo); thin printed-circuit cables connect the read-write heads to amplifier electronics mounted at the pivot of the actuator. A flexible, somewhat U-shaped, ribbon cable, seen edge-on below and to the left of the actuator arm continues the connection to the controller board on the opposite side.

The head support arm is very light, but also stiff; in modern drives, acceleration at the head reaches 550 *G*s. The silver-colored structure at the upper left of the first image is the top plate of the actuator, a permanent-magnet and moving coil motor that swings the heads to the desired position (it is shown removed in the second image). The plate supports a squat neodymium-iron-boron (NIB) high-flux magnet. Beneath this plate is the moving coil, often referred to as the *voice coil* by analogy to the coil in loudspeakers, which is attached to the actuator hub, and beneath that is a second NIB magnet, mounted on the bottom plate of the motor (some drives only have one magnet).

The voice coil itself is shaped rather like an arrowhead, and made of doubly coated copper magnet wire. The inner layer is insulation, and the outer is thermoplastic, which bonds the coil together after it is wound on a form, making it self-supporting. The portions of the coil along the two sides of the arrowhead (which point to the actuator bearing center) interact with the magnetic field, developing a tangential force that rotates the actuator. Current flowing radially outward along one side of the arrowhead and radially inward on the other produces the tangential force. If the magnetic field were uniform, each side would generate opposing forces that would cancel each other out. Therefore the surface of the magnet is half N pole, half S pole, with the radial dividing line in the middle, causing the two sides of the coil to see opposite magnetic fields and produce forces that add instead of canceling. Currents along the top and bottom of the coil produce radial forces that do not rotate the head.

ERROR HANDLING

Modern drives also make extensive use of Error Correcting Codes (ECCs), particularly Reed–Solomon error correction. These techniques store extra bits for each block of data that are determined by mathematical formulas. The extra bits allow many errors to be fixed. While these extra bits take up space on the hard drive, they allow higher recording densities to be employed, resulting in much larger storage capacity for user data. In 2009, in the newest drives, low-density parity-check codes (LDPC) are supplanting Reed-Solomon. LDPC codes enable performance close to the Shannon Limit and thus allow for the highest storage density available.

Typical hard drives attempt to "remap" the data in a physical sector that is going bad to a spare physical sector—hopefully while the errors in that bad sector are still few enough that the ECC can recover the data without loss. The S.M.A.R.T. system counts the total number of errors in the entire hard drive fixed by ECC, and the total number of remappings, in an attempt to predict hard drive failure.

FUTURE DEVELOPMENT

Because of bit-flipping errors and other issues, perpendicular recording densities may be supplanted by other magnetic recording technologies. Toshiba is promoting bit-patterned recording (BPR), whiles Xyratex are developing heat-assisted magnetic recording (HAMR).

EXTERNAL REMOVABLE DRIVES

External removable hard disk drives connect to the computer using a USB cable or other means. External drives are used for:

- Backup of files and information
- Data recovery
- Disk cloning
- Running virtual machines
- Scratch disk for video editing applications and video recording.

Larger models often include full-sized 3.5" PATA or SATA desktop hard drives. Features such as biometric security or multiple interfaces generally increase cost.

MARKET SEGMENTS

- As of July 2010, the highest capacity consumer HDDs are 3 TB.
- “Desktop HDDs” typically store between 120 GB and 2 TB and rotate at 5,400 to 10,000 rpm, and have a media transfer rate of 0.5 Gbit/s or higher. (1 GB = 10^9 bytes; 1 Gbit/s = 10^9 bit/s)
- Enterprise HDDs are typically used with multiple-user computers running enterprise software. Examples are

- o transaction processing databases;
- o internet infrastructure (email, webserver, e-commerce);
- o scientific computing software;
- o nearline storage management software.

The fastest enterprise HDDs spin at 10,000 or 15,000 rpm, and can achieve sequential media transfer speeds above 1.6 Gbit/s. and a sustained transfer rate up to 1 Gbit/s. Drives running at 10,000 or 15,000 rpm use smaller platters to mitigate increased power requirements (as they have less air drag) and therefore generally have lower capacity than the highest capacity desktop drives.

Enterprise drives commonly operate continuously ("24/7") in demanding environments while delivering the highest possible performance without sacrificing reliability. Maximum capacity is not the primary goal, and as a result the drives are often offered in capacities that are relatively low in relation to their cost.

- Mobile HDDs or laptop HDDs, smaller than their desktop and enterprise counterparts, tend to be slower and have lower capacity. A typical mobile HDD spins at either 4200 rpm, 5200 rpm, 5400 rpm, or 7200 rpm, with 5400 rpm being the most prominent. 7200 rpm drives tend to be more expensive and have smaller capacities, while 4200 rpm models usually have very high storage capacities. Because of smaller platter(s), mobile HDDs generally have lower capacity than their greater desktop counterparts.

The exponential increases in disk space and data access speeds of HDDs have enabled the commercial viability of consumer products that require large storage capacities, such as digital video recorders and digital audio players. In addition, the availability of vast amounts of cheap storage has made viable a variety of web-based services with extraordinary capacity requirements, such as free-of-charge web search, web archiving, and video sharing (Google, Internet Archive, YouTube, etc.).

SALES

Worldwide revenue from shipments of HDDs is expected to reach $27.7 billion in 2010, up 18.4% from $23.4 billion in 2009 corresponding to a 2010 unit shipment forecast of 674.6 million compared to 549.5 million units in 2009.

ICONS

Hard drives are traditionally symbolized as either a stylized stack of platters (in orthographic projection) or, more abstractly, as a cylinder. This is particularly found in schematic diagrams or on indicator lights, as on laptops, to indicate hard drive access. In most modern operating systems, hard drives are instead represented by an illustration or photograph of a hard drive enclosure.

MANUFACTURERS

More than 200 companies have manufactured hard disk drives. Today most drives are made by Seagate, Western Digital, Hitachi, Samsung, and Toshiba (though Toshiba does not manufacture 3.5 inch drives).

COMPACT DISC

A Compact Disc (also known as a CD) is an optical disc used to store digital data. It was originally developed to store and playback sound recordings exclusively, but later expanded to encompass data storage CD-ROM, write-once audio and data storage CD-R, rewritable media CD-RW, Video Compact Discs (VCD), Super Video Compact Discs (SVCD), PhotoCD, PictureCD, CD-i, and Enhanced CD. Audio CDs and audio CD players have been commercially available since October 1982. Standard CDs have a diameter of 120 millimetres (4.7 in) and can hold up to 80 minutes of uncompressed audio (700 MB of data). The Mini CD has various diameters ranging from 60 to 80 millimetres (2.4 to 3.1 in); they are sometimes used for CD singles, storing up to 24 minutes of audio or delivering device drivers.

CD-ROMs and CD-Rs remain widely used technologies in the computer industry. The CD and its extensions are successful: in 2004, worldwide sales of CD audio, CD-ROM, and CD-R reached about 30 billion discs. By 2007, 200 billion CDs had been sold worldwide. Compact Discs are increasingly being replaced or supplemented by other forms of digital distribution and storage, such as downloading and flash drives, with audio CD sales dropping nearly 50% from their peak in 2000.

HISTORY

The Compact Disc is a spin-off of Laserdisc technology. Sony first publicly demonstrated an optical digital audio disc in September 1976. In September 1978 they demonstrated an optical digital audio disc with a 150 minute playing time, and with specifications of 44,056 Hz sampling rate, 16-bit linear resolution, cross-interleaved error correction code, that were similar to those of the Compact

Disc introduced in 1982. Technical details of Sony's digital audio disc were presented during the 62nd AES Convention, held on March 13–16, 1979, in Brussels. On March 8, 1979 Philips publicly demonstrated a prototype of an optical digital audio disc at a press conference called "Philips Introduce Compact Disc" in Eindhoven, Netherlands. On March 6, 2009, Philips received an IEEE Milestone with the following citation: "On 8 March 1979, N.V. Philips' Gloeilampenfabrieken demonstrated for the international press a Compact Disc Audio Player. The demonstration showed that it is possible by using digital optical recording and playback to reproduce audio signals with superb stereo quality. This research at Philips established the technical standard for digital optical recording systems."

Later in 1979, Sony and Philips Consumer Electronics (Philips) set up a joint task force of engineers to design a new digital audio disc. Led by Kees Schouhamer Immink and Toshitada Doi, the research pushed forward laser and optical disc technology that began independently by Philips and Sony in 1977 and 1975, respectively. After a year of experimentation and discussion, the taskforce produced the *Red Book*, the Compact Disc standard. Philips contributed the general manufacturing process, based on video Laserdisc technology. Philips also contributed eight-to-fourteen modulation (EFM), which offers both a long playing time and a high resilience against disc defects such as scratches and fingerprints, while Sony contributed the error-correction method, CIRC. The *Compact Disc Story*, told by a former member of the taskforce, gives background information on the many technical decisions made, including the choice of the sampling frequency, playing time, and disc diameter. The taskforce consisted of around four to eight persons, though according to Philips, the Compact Disc was thus "invented collectively by a large group of people working as a team."

The first test CD was pressed in Langenhagen near Hannover, Germany, by the Polydor Pressing Operations plant. The disc contained a recording of Richard Strauss's *Eine Alpensinfonie* (in English, *An Alpine Symphony*), played by the Berlin Philharmonic and conducted by Herbert von Karajan. The first public demonstration was on the BBC television programme *Tomorrow's World* when The Bee Gees' album *Living Eyes* (1981) was played. In August 1982 the real pressing was ready to begin in the new factory, not far from the place where Emil Berliner had produced his first gramophone record 93 years earlier. By now, Deutsche Grammophon, Berliner's company and the publisher of the Strauss recording, had become a part of PolyGram. The first CD to be manufactured at the new factory was *The Visitors* (1981) by ABBA. The first album to be *released* on CD was Billy Joel's *52nd Street*, that reached the market alongside Sony's CD player CDP-101 on October 1, 1982 in Japan. Early the following year on March 2, 1983 CD players and discs (16 titles from CBS Records) were released in the United States and other markets. This event is often seen as the "Big Bang" of the digital audio revolution. The new audio disc was enthusiastically received, especially in the early-adopting classical music and audiophile communities and its handling quality received particular praise. As the price of players gradually came down, the CD began to gain popularity in the larger popular and rock music markets. The first artist to sell a million copies on CD was Dire Straits, with its 1985 album *Brothers in Arms*. The first major artist to have his entire catalogue converted to CD was David Bowie, whose 15 studio albums were made available by RCA Records in February 1985, along with four Greatest Hits albums. In 1988, 400 million CDs were manufactured by 50 pressing plants around the world.

The CD was planned to be the successor of the gramophone record for playing music, rather than primarily as a data storage medium. From its origins as a musical format, CDs have grown to encompass other applications. In June 1985, the computer readable CD-ROM (read-only memory) and, in 1990, CD-Recordable were introduced, also developed by both Sony and Philips. The CD's compact format has largely replaced the audio cassette player in new automobile applications, and recordable CDs are an alternative to tape for recording music and copying music albums without defects introduced in compression used in other digital recording methods. Other newer video formats such as DVD and Blu-ray have used the same form factor as CDs, and video players can usually play audio CDs as well. With the advent of the MP3 in the 2000s, the sales of CDs has dropped in seven out of the last eight years. In 2008, large label CD sales dropped 20%, although independent and DIY music sales may be tracking better according to figures released March 30, 2009.

PHYSICAL DETAILS

CD is made from 1.2 millimetres (0.047 in) thick, polycarbonate plastic and weighs 15–20 grams. From the center outward, components are: the center spindle hole

(15 mm), the first-transition area (clamping ring), the clamping area (stacking ring), the second-transition area (mirror band), the programme (data) area, and the rim. The inner programme area occupies a radius from 25 to 58 mm

A thin layer of aluminium or, more rarely, gold is applied to the surface making it reflective. The metal is protected by a film of lacquer normally spin coated directly on the reflective layer. The label is printed on the lacquer layer, usually by screen printing or offset printing.

CD data are stored as a series of tiny indentations known as "pits", encoded in a spiral track moulded into the top of the polycarbonate layer. The areas between pits are known as "lands". Each pit is approximately 100 nm deep by 500 nm wide, and varies from 850 nm to 3.5 μm in length. The distance between the tracks, the pitch, is 1.6 μm.

Scanning velocity is 1.2–1.4 m/s (constant linear velocity) – equivalent to approximately 500 rpm at the inside of the disc, and approximately 200 rpm at the outside edge. (A disc played from beginning to end slows down during playback.)

The programme area is 86.05 cm^2 and the length of the recordable spiral is (86.05 cm^2 / 1.6 μm) = 5.38 km. With a scanning speed of 1.2 m/s, the playing time is 74 minutes, or 650 MB of data on a CD-ROM. A disc with data packed slightly more densely is tolerated by most players (though some old ones fail). Using a linear velocity of 1.2 m/s and a track pitch of 1.5 μm yields a playing time of 80 minutes, or a data capacity of 700 MB. Even higher capacities on non-standard discs (up to 99 minutes) are available at least as recordables, but generally the tighter the tracks are squeezed, the worse the compatibility.

A CD is read by focusing a 780 nm wavelength (near infrared) semiconductor laser through the bottom of the polycarbonate layer. The change in height between pits and lands results in a difference in the way the light is reflected. By measuring the intensity change with a photodiode, the data can be read from the disc.

The pits and lands themselves do not directly represent the zeros and ones of binary data. Instead, non-return-to-zero, inverted encoding is used: a change from pit to land or land to pit indicates a one, while no change indicates a series of zeros. There must be at least two and no more than ten zeros between each one, which is defined by the length of the pit. This in turn is decoded by reversing the eight-to-fourteen modulation used in mastering the disc, and then reversing the Cross-Interleaved Reed-Solomon Coding, finally revealing the raw data stored on the disc.

CDs are susceptible to damage from both normal use and environmental exposure. Pits are much closer to the label side of a disc, enabling defects and contaminants on the clear side to be out of focus during playback. Consequently, CDs are more likely to suffer damage on the label side of the disk. Scratches on the clear side can be repaired by refilling them with similar refractive plastic, or by careful polishing.

DISC SHAPES AND DIAMETERS

The digital data on a CD begins at the center of the disc and proceeds toward the edge, which allows adaptation to the different size formats available. Standard CDs are available in two sizes. By far, the most common is 120 millimetres (4.7 in) in diameter, with a 74- or 80-minute audio capacity and a 650 or 700 MB data capacity. This diameter has been adopted by subsequent formats, including Super Audio CD, DVD, HD DVD, and Blu-ray Disc. 80 mm discs ("Mini CDs") were originally designed for CD singles and can hold up to 24 minutes of music or 210 MB of data but never became popular. Today, nearly every single is released on a 120 mm CD, called a Maxi single.

Novelty CDs are also available in numerous shapes and sizes, and are used chiefly for marketing. A common variant is the "business card" CD, a single with portions removed at the top and bottom making the disk resemble a business card.

LOGICAL FORMATS

AUDIO CD

The logical format of an audio CD (officially Compact Disc Digital Audio or CD-DA) is described in a document produced by the format's joint creators, Sony and Philips in 1980. The document is known colloquially as the "Red Book" after the color of its cover. The format is a two-channel 16-bit PCM encoding at a 44.1 kHz sampling rate per channel. Four-channel sound is an allowable option within the Red Book format, but has never been implemented. Monaural audio has no existing standard on a Red Book CD; mono-source material is usually presented as two identical channels on a 'stereo' track.

44.1 KHZ SAMPLE RATE

The selection of the sample rate was based primarily on the need to reproduce the audible frequency range of 20 Hz - 20 kHz. The Nyquist–Shannon sampling theorem states that a sampling rate of more than twice the maximum frequency of the signal to be recorded is needed, resulting in a required rate of at least 40 kHz. The exact sampling rate of 44.1 kHz was inherited from a method of converting digital audio into an analog video signal for storage on U-matic video tape, which was the most affordable way to transfer data from the recording studio to the CD manufacturer at the time the CD specification was being developed. The device that converts an analog audio signal into PCM audio, which in turn is changed into an analog video signal is called a PCM adaptor. This technology could store six samples (three samples per stereo channel) in a single horizontal line. A standard NTSC video signal has 245 usable lines per field, and 59.94 fields/s, which works out to be 44,056 samples/s/stereo channel. Similarly, PAL has 294 lines and 50 fields, which gives 44,100 samples/s/stereo channel. This system could store 14-bit samples with some error correction, or 16-bit samples with almost no error correction.

There was a long debate over the use of 14-bit (Philips) or 16-bit (Sony) quantization, and 44,056 or 44,100 samples/s (Sony) or approximately 44,000 samples/s (Philips). When the Sony/Philips task force designed the Compact Disc, Philips had already developed a 14-bit D/A converter, but Sony insisted on 16-bit. In the end, 16 bits and 44.1 kilosamples per second prevailed. Philips found a way to produce 16-bit quality using its 14-bit DAC by using four times oversampling.

STORAGE CAPACITY AND PLAYING TIME

The partners aimed at a playing time of 60 minutes with a disc diameter of 100 mm (Sony) or 115 mm (Philips). Sony vice-president Norio Ôga suggested extending the capacity to 74 minutes to accommodate Wilhelm Furtwängler's recording of Ludwig van Beethoven's *Symphony Number Nine* from the 1951 Bayreuth Festival.

The additional 14-minute playing time subsequently required changing to a 120 mm disc. Kees Immink, Philips' chief engineer, however, denies this, claiming that the increase was motivated by technical considerations, and that even after the increase in size, the Furtwängler recording would not have fit on one of the earliest CDs. According to a *Sunday Tribune* interview, the story is slightly more involved. In 1979, Philips owned PolyGram, one of the world's largest distributors of music. PolyGram had set up a large experimental CD plant in Hannover, Germany, which could produce huge numbers of CDs having, of course, a diameter of 115 mm. Sony did not yet have such a facility. If Sony had agreed on the 115-mm disc, Philips would have had a significant competitive edge in the market. Sony decided that something had to be done. The long playing time of Beethoven's Ninth Symphony imposed by Ôga was used to push Philips to accept 120 mm, so that Philips' PolyGram lost its edge on disc fabrication.

The 74-minute playing time of a CD, which was longer than the 20 minutes per side typical of long-playing (LP) vinyl albums, was often used to the CD's advantage during the early years when CDs and LPs vied for commercial sales. CDs would often be released with one or more bonus tracks, enticing consumers to buy the CD for the extra material. However, attempts to combine double LPs onto one CD occasionally resulted in the opposite situation in which the CD would actually offer fewer tracks than the equivalent LP.

Playing times beyond 74 minutes are achieved by decreasing track pitch beyond the original red book standard. Most players can accommodate the more closely spaced data. Christian Thielemann's live Deutsche Grammophon recording of Bruckner's Fifth with the Munich Philharmonic in 2004 clocks at 82:34. The Kirov Orchestra recording of Pyotr Ilyich Tchaikovsky's *The Nutcracker* conducted by Valery Gergiev and released by Philips/PolyGram Records (catalogue number 462 114) on October 20, 1998, clocks at 81:14. Disc two of *Gold* (Deutche Grammophon/Universal Classics 477 743) by Herbert von Karajan clocks in at 81:21. The Mission of Burma compilation album *Mission of Burma*, released in 1988 by Rykodisc, previously held the record at 80:08.

DATA STRUCTURE

The smallest entity in a CD is called a *frame*, which consists of 33 bytes and contains six complete 16-bit stereo samples (two bytes × two channels × six samples = 24 bytes). The other nine bytes consist of eight CIRC error-correction bytes and one subcode byte, used for control and display. Each byte is translated into a 14-bit word using eight-to-fourteen modulation, which alternates with three-bit merging words. In total there are 33 × (14 + 3) = 561 bits. A 27-bit unique synchronization word is added, so that the number of bits in a frame totals 588 (of which only 192 bits are music).

These 588-bit frames are in turn grouped into sectors. Each sector contains 98 frames, totaling 98 × 24 = 2352 bytes of music. The CD is played at a speed of 75 sectors per second, which results in 176,400 bytes per second. Divided by two channels and two bytes per sample, this results in a sample rate of 44,100 samples per second.

For CD-ROM data discs, the physical frame and sector sizes are the same. Since error concealment cannot be applied to non-audio data in case the CIRC error correction fails to recover the user data, a third layer of error correction is defined, reducing the payload to 2048 bytes per sector for the Mode-1 CD-ROM format. To increase the data-rate for Video CD, Mode-2 CD-ROM, the third layer has been omitted, increasing the payload to 2336 user-available bytes per sector, only 16 bytes (for synchronization and header data) less than available in Red-Book audio.

Frame: For the Red Book stereo audio CD, the time format is commonly measured in minutes, seconds and frames (mm:ss:ff), where one frame corresponds to one sector, or 1/75th of a second of stereo sound. In this context, the term *frame* is erroneously applied in editing applications and does not denote the physical frame described above. In editing and extracting, the frame is the smallest addressable time interval for an audio CD, meaning that track start and end positions can only be defined in 1/75 second steps.

Logical structure: The largest entity on a CD is called a track. A CD can contain up to 99 tracks (including a data track for mixed mode discs). Each track can in turn have up to 100 indexes, though players which handle this feature are rarely found outside of pro audio, particularly radio broadcasting. The vast majority of songs are recorded under index 1, with the pre-gap being index 0.

Sometimes hidden tracks are placed at the end of the last track of the disc, often using index 2 or 3. This is also the case with some discs offering "101 sound effects", with 100 and 101 being indexed as two and three on track 99. The index, if used, is occasionally put on the track listing as a decimal part of the track number, such as 99.2 or 99.3. (Information Society's *Hack* was one of very few CD releases to do this, following a release with an equally obscure CD+G feature.) The track and index structure of the CD carried forward to the DVD as title and chapter, respectively.

Manufacturing tolerances: Current manufacturing processes allow an audio CD to contain up to 80 minutes (variable from one replication plant to another) without requiring the content creator to sign a waiver releasing the plant owner from responsibility if the CD produced is marginally or entirely unreadable by some playback equipment. Thus, in current practice, maximum CD playing time has crept higher by reducing minimum engineering tolerances; by and large, this has not unacceptably reduced reliability.

CD-TEXT

CD-Text is an extension of the Red Book specification for audio CD that allows for storage of additional text information (e.g., album name, song name, artist) on a standards-compliant audio CD. The information is stored either in the lead-in area of the CD, where there is roughly five kilobytes of space available, or in the subcode channels R to W on the disc, which can store about 31 megabytes.

CD + GRAPHICS

Compact Disc + Graphics (CD+G) is a special audio Compact Disc that contains graphics data in addition to the audio data on the disc. The disc can be played on a regular audio CD player, but when played on a special CD+G player, can output a graphics signal (typically, the CD+G player is hooked up to a television set or a computer monitor); these graphics are almost exclusively used to display lyrics on a television set for karaoke performers to sing along with. The CD+G format takes advantage of the channels R through W. These six bits store the graphics information.

CD + Extended Graphics: Compact Disc + Extended Graphics (CD+EG, also known as CD+XG) is an improved variant of the Compact Disc + Graphics (CD+G) format. Like CD+G, CD+EG utilizes basic CD-ROM features to display text and video information in addition to the music being played. This extra data is stored in subcode channels R-W. Very few, if any, CD+EG discs have been published.

SUPER AUDIO CD

Super Audio CD (SACD) is a high-resolution read-only optical audio disc format that provides much higher fidelity digital audio reproduction than the Red Book. Introduced in 1999, it was developed by Sony and Philips, the same companies that created the Red Book. SACD was in a format war with DVD-Audio, but neither has replaced audio CDs.

In contrast to DVD-Audio, the SACD format has the feature of being able to produce hybrid discs; these discs contain the SACD audio stream as well as a standard audio CD layer which is playable in standard CD players, thus making them backward compatible.

CD-MIDI

CD-MIDI is a format used to store music-performance data which upon playback is performed by electronic instruments that synthesize the audio. Hence, unlike Red Book, these recordings are not audio.

CD-ROM

For the first few years of its existence, the Compact Disc was a medium used purely for audio. However, in 1985 the Yellow Book CD-ROM standard was established by Sony and Philips, which defined a non-volatile optical data computer data storage medium using the same physical format as audio Compact Discs, readable by a computer with a CD-ROM drive.

VIDEO CD (VCD)

Video CD (VCD, View CD, and Compact Disc digital video) is a standard digital format for storing video media on a CD. VCDs are playable in dedicated VCD players, most modern DVD-Video players, personal computers, and some video game consoles. The VCD standard was created in 1993 by Sony, Philips, Matsushita, and JVC and is referred to as the White Book standard.

Overall picture quality is intended to be comparable to VHS video. Poorly compressed VCD video can sometimes be lower quality than VHS video, but VCD exhibits block artifacts rather than analog noise, and does not deteriorate further with each use, which may be preferable.

352x240 (or SIF) resolution was chosen because it is half the vertical, and half the horizontal resolution of NTSC video. 352x288 is similarly one quarter PAL/SECAM resolution. This approximates the (overall) resolution of an analog VHS tape, which, although it has double the number of (vertical) scan lines, has a much lower horizontal resolution.

SUPER VIDEO CD

Super Video CD (Super Video Compact Disc or SVCD) is a format used for storing video media on standard Compact Discs. SVCD was intended as a successor to VCD and an alternative to DVD-Video, and falls somewhere between both in terms of technical capability and picture quality.

SVCD has two-thirds the resolution of DVD, and over 2.7 times the resolution of VCD. One CD-R disc can hold up to 60 minutes of standard quality SVCD-format video. While no specific limit on SVCD video length is mandated by the specification, one must lower the video bit rate, and therefore quality, in order to accommodate very long videos. It is usually difficult to fit much more than 100 minutes of video onto one SVCD without incurring significant quality loss, and many hardware players are unable to play video with an instantaneous bit rate lower than 300 to 600 kilobits per second.

PHOTO CD

Photo CD is a system designed by Kodak for digitizing and storing photos on a CD. Launched in 1992, the discs were designed to hold nearly 100 high quality images, scanned prints and slides using special proprietary encoding. Photo CD discs are defined in the Beige Book and conform to the CD-ROM XA and CD-i Bridge specifications as well. They are intended to play on CD-i players, Photo CD players and any computer with the suitable software irrespective of the operating system. The images can also be printed out on photographic paper with a special Kodak machine. This format is not to be confused with Kodak Picture CD, which is a consumer product in CD-ROM format.

CD-I

The Philips "Green Book" specifies the standard for interactive multimedia Compact Discs designed for CD-i players. This format is unusual because it hides the initial tracks which contains the software and data files used by CD-i players by omitting the tracks from the disc's TOC (table of contents). This causes audio CD players to skip the CD-i data tracks. This is different from the CD-i Ready format, which puts CD-i software and data into the pregap of track 1. CDi was the leading format of its time but was supplanted by the politics of competition. Philips Interactive

Media lead the way in producing breakthrough titles, including the first interactive coloring book, Sesame Street Disc and children's programmes, Groliers and Comptoms encyclopedias and many more pathbreaking programmes.

ENHANCED CD

Enhanced CD, also known as CD Extra and CD Plus, is a certification mark of the Recording Industry Association of America for various technologies that combine audio and computer data for use in both Compact Disc and CD-ROM players.

The primary data formats for Enhanced Compact Disc's are mixed mode (Yellow Book/Red Book), CD-i, hidden track, and multisession (Blue Book).

VINYLDISC

VinylDisc is the hybrid of a standard Audio CD and the vinyl record. The vinyl layer on the disc's label side can hold approximately three minutes of music.

MANUFACTURE

Replicated CDs are mass-produced initially using a hydraulic press. Small granules of heated raw polycarbonate plastic are fed into the press. A screw forces the liquefied plastic into the mold cavity. The mold closes with a metal stamper in contact with the disc surface. The plastic is allowed to cool and harden. Once opened, the disc substrate is removed from the mold by a robotic arm, and a 15 mm diameter center hole (called a stacking ring) is created. The time it takes to "stamp" one CD, is usually 2 to 3 seconds.

This method produces the clear plastic blank part of the disc. After a metallic reflecting layer (usually aluminum, but sometimes gold or other metal) is applied to the clear blank substrate, the disc goes under a UV light for curing and it is ready to go to press. To prepare to press a CD, a glass master is made, using a high-powered laser on a device similar to a CD writer. The glass master is a positive image of the desired CD surface (with the desired microscopic pits and lands). After testing, it is used to make a die by pressing it against a metal disc.

The die is a negative image of the glass master: typically, several are made, depending on the number of pressing mills that are to make the CD. The die then goes into a press and the physical image is transferred to the blank CD, leaving a final positive image on the disc. A small amount of lacquer is applied as a ring around the center of the disc, and rapid spinning spreads it evenly over the surface. Edge protection lacquer is applied before the disc is finished. The disc can then be printed and packed.

Manufactured CDs that are sold in stores are sealed via a process called "polywrapping" or shrink wrapping.

RECORDABLE CD

Recordable Compact Discs, CD-Rs, are injection molded with a "blank" data spiral. A photosensitive dye is then applied, after which the discs are metalized and lacquer-coated. The write laser of the CD recorder changes the color of the dye to allow the read laser of a standard CD player to see the data, just as it would with a standard stamped disc. The resulting discs can be read by most CD-ROM drives and played in most audio CD players. CD-R recordings are designed to be permanent. Over time the dye's physical characteristics may change, however, causing read errors and data loss until the reading device cannot recover with error correction methods. The design life is from 20 to 100 years, depending on the quality of the discs, the quality of the writing drive, and storage conditions. However, testing has demonstrated such degradation of some discs in as little as 18 months under normal storage conditions. This failure is known as CD rot. CD-Rs follow the Orange Book standard.

RECORDABLE AUDIO CD

The recordable audio CD is designed to be used in a consumer audio CD recorder. These consumer audio CD recorders use SCMS (Serial Copy Management System), an early form of digital rights management (DRM), to conform to the AHRA (Audio Home Recording Act). The Recordable Audio CD is typically somewhat more expensive than CD-R due to (a) lower volume and (b) a 3% AHRA royalty used to compensate the music industry for the making of a copy.

HIGH-CAPACITY RECORDABLE CD

A higher density recording format that can hold:

- 98.5 minutes of audio on a 12 cm disc (compared to about 80 minutes for Red Book audio).
- 30 minutes of audio on an 8 cm disc (compared to about 24 minutes for Red Book audio).

REWRITABLE CD

CD-RW is a re-recordable medium that uses a metallic alloy instead of a dye. The write laser in this case is used to heat and alter the properties (amorphous vs. crystalline) of the alloy, and hence change its reflectivity. A CD-RW does not have as great a difference in reflectivity as a pressed CD or a CD-R, and so many earlier CD audio players *cannot* read CD-RW discs, although *most* later CD audio players and stand-alone DVD players can. CD-RWs follow the Orange Book standard.

HIGH-SPEED REWRITABLE CD

Due to technical limitations, the original ReWritable CD could be written no faster than 4x speed. High Speed ReWritable CD has a different design that permits writing at speeds ranging from 4x to 12x.

Original CD-RW drives can only write to original ReWritable CDs. High Speed CD-RW drives can typically write to both original ReWritable CD discs and High Speed ReWritable CD discs. Both types of CD-RW discs can be read in most CD drives.

Higher speed CD-RW discs, Ultra Speed (16x to 24x write speed) and Ultra Speed+ (32x write speed), are now available.

REWRITABLE AUDIO CD

The ReWritable Audio CD is designed to be used in a consumer audio CD recorder, which won't (without modification) accept standard CD-RW discs. These consumer audio CD recorders use the Serial Copy Management System (SCMS), an early form of digital rights management (DRM), to conform to the United States' Audio Home Recording Act (AHRA). The ReWritable Audio CD is typically somewhat more expensive than CD-RW due to (a) lower volume and (b) a 3% AHRA royalty used to compensate the music industry for the making of a copy.

COPY PROTECTION

The Red Book audio specification, except for a simple 'anti-copy' bit in the subcode, does not include any copy protection mechanism. Starting in early 2002, attempts were made by record companies to market "copy-protected" non-standard Compact Discs, which cannot be ripped, or copied, to hard drives or easily converted to MP3s. One major drawback to these copy-protected discs is that most will not play on either computer CD-ROM drives, or some standalone CD players that use CD-ROM mechanisms. Philips has stated that such discs are not permitted to bear the trademarked *Compact Disc Digital Audio* logo because they violate the Red Book specifications. Numerous copy-protection systems have been countered by readily available, often free, software.

CAPACITY MEASUREMENTS IN HARD DISK

Hard disk manufacturers quote disk capacity in multiples of SI-standard powers of 1000, where a *terabyte* is 1000 gigabytes and a *gigabyte* is 1000 megabytes. With file systems that report capacity in powers of 1024, available space appears somewhat less than advertised capacity. The discrepancy between the two methods of reporting sizes had serious financial consequences for at least one hard drive manufacturer when a class action suit argued the different methods effectively misled consumers.

Semiconductor memory chips are organized so that memory sizes are expressed in multiples of powers of two. Hard disks by contrast have no inherent binary size. Capacity is the product of the number of heads, number of tracks, number of sectors per track, and the size of each sector. Sector sizes are standardized for convenience at 256 or 512 and more recently 4096 bytes, which are powers of two. This can cause some confusion because operating systems may report the formatted capacity of a hard drive using binary prefix units which increment by powers of 1024. For example, Microsoft Windows reports disk capacity both in a decimal integer to 12 or more digits and in binary prefix units to three significant digits.

A one terabyte (1 TB) disk drive would be expected to hold around 1 trillion bytes (1,000,000,000,000) or 1000 GB; and indeed most 1 TB hard drives will contain slightly more than this number. However some operating system utilities would report this as around 931 GB or 953,674 MB. (The actual number for a formatted capacity will be somewhat smaller still, depending on the file system.) Following are the several ways of reporting one Terabyte.

ADDRESSING DATA ON LARGE DRIVES

The capacity of an HDD can be calculated by multiplying the number of cylinders by the number of heads by the number of sectors by the number of bytes/sector (most commonly 512). Drives with the ATA interface and a capacity of eight gigabytes or more behave as if they were structured into 16383 cylinders, 16 heads, and 63 sectors, for compatibility with older operating systems. Unlike in the 1980s, the cylinder, head, sector (C/H/S) counts reported to the CPU by a modern ATA drive are no longer actual physical parameters since the reported numbers are constrained by historic operating-system interfaces and with zone bit recording the actual number of sectors varies by zone. Disks with SCSI interface address each sector with a unique integer number; the operating system remains ignorant of their head or cylinder count.

The old C/H/S scheme has been replaced by logical block addressing. In some cases, to try to "force-fit" the C/H/S scheme to large-capacity drives, the number of heads was given as 64, although no modern drive has anywhere near 32 platters.

Not all the space on a hard drive is available for user files. The operating system file system uses some of the disk space to organize files on the disk, recording their file names and the sequence of disk areas that represent the file. Examples of data structures stored on disk to retrieve files include the MS DOS file allocation table (FAT), and UNIX inodes, as well as other operating system data structures. This file system overhead is usually less than 1% on drives larger than 100 MB.

For RAID drives, data integrity and fault-tolerance requirements also reduce the realized capacity. For example, a RAID1 drive will be about half the total capacity as a result of data mirroring. For RAID5 drives with x drives you would lose 1/x of your space to parity. RAID drives are multiple drives that appear to be one drive to the user, but provides some fault-tolerance. A general rule of thumb to quickly convert the manufacturer's hard disk capacity to the standard Microsoft Windows formatted capacity is 0.93*capacity of HDD from manufacturer for HDDs less than a terabyte and 0.91*capacity of HDD from manufacturer for HDDs equal to or greater than 1 terabyte.

HDD FORMATTING

The presentation of an HDD to its host is determined by its controller. This may differ substantially from the drive's native interface particularly in mainframes or servers.

Modern HDDs, such as SAS and SATA drives, appear at their interfaces as a contiguous set of logical blocks; typically 512 bytes long but the industry is in the process of changing to 4,096 byte logical blocks; see Advanced Format.

The process of initializing these logical blocks on the physical disk platters is called *low level formatting* which is usually performed at the factory and is not normally changed in the field. *High level formatting* then writes the file system structures into selected logical blocks to make the remaining logical blocks available to the host OS and its applications.

FORM FACTORS

Mainframe and minicomputer hard disks were of widely varying dimensions, typically in free standing cabinets the size of washing machines (e.g. HP 7935 and DEC RP06 Disk Drives) or designed so that dimensions enabled placement in a 19" rack (e.g. Diablo Model 31). In 1962, IBM introduced its model 1311 disk, which used 14 inch (nominal size) platters. This became a standard size for mainframe and minicomputer drives for many years, but such large platters were never used with microprocessor-based systems.

With increasing sales of microcomputers having built in floppy-disk drives (FDDs), HDDs that would fit to the FDD mountings became desirable, and this led to the evolution of the market towards drives with certain Form factors, initially derived from the sizes of 8-inch, 5.25-inch, and 3.5-inch floppy disk drives. Smaller sizes than 3.5 inches have emerged as popular in the marketplace and/or been decided by various industry groups.

- 8 inch: 9.5 in × 4.624 in × 14.25 in (241.3 mm × 117.5 mm × 362 mm)

  In 1979, Shugart Associates' SA1000 was the first form factor compatible HDD, having the same dimensions and a compatible interface to the 83 FDD.

- 5.25 inch: 5.75 in × 3.25 in × 8 in (146.1 mm × 82.55 mm × 203 mm)

  This smaller form factor, first used in an HDD by Seagate in 1980, was the same size as full-height 5+¹D ₄-inch-diameter (130 mm) FDD, 3.25-inches high. This is twice as high as "half height"; i.e., 1.63 in (41.4 mm). Most desktop models of drives for optical 120 mm disks (DVD, CD) use the half height 5¼3 dimension, but it fell out of fashion for HDDs. The Quantum Bigfoot HDD was the last to use it in the late 1990s, with "low-profile" (H"25 mm) and "ultra-low-profile" (H"20 mm) high versions.

- 3.5 inch: 4 in × 1 in × 5.75 in (101.6 mm × 25.4 mm × 146 mm) = 376.77344 cm^3

  This smaller form factor, first used in an HDD by Rodime in 1983, was the same size as the "half height" 3½3 FDD, i.e., 1.63 inches high. Today it has been largely superseded by 1-inch high "slimline" or "low-profile" versions of this form factor which is used by most desktop HDDs.
- 2.5 inch: 2.75 in × 0.275–0.59 in × 3.945 in (69.85 mm × 7–15 mm × 100 mm) = 48.895–104.775 cm^3

  This smaller form factor was introduced by PrairieTek in 1988; there is no corresponding FDD. It is widely used today for hard-disk drives in mobile devices (laptops, music players, etc.) and as of 2008 replacing 3.5 inch enterprise-class drives. It is also used in the Playstation 3 and Xbox 360 video game consoles. Today, the dominant height of this form factor is 9.5 mm for laptop drives (usually having two platters inside), but higher capacity drives have a height of 12.5 mm (usually having three platters). Enterprise-class drives can have a height up to 15 mm. Seagate has released a wafer-thin 7mm drive aimed at entry level laptops and high end netbooks in December 2009.
- 1.8 inch: 54 mm × 8 mm × 71 mm = 30.672 cm^3

  This form factor, originally introduced by Integral Peripherals in 1993, has evolved into the ATA-7 LIF with dimensions as stated. It was increasingly used in digital audio players and subnotebooks, but is rarely used today. An original variant exists for 2–5GB sized HDDs that fit directly into a PC card expansion slot. These became popular for their use in iPods and other HDD based MP3 players.

- 1 inch: 42.8 mm × 5 mm × 36.4 mm

  This form factor was introduced in 1999 as IBM's Microdrive to fit inside a CF Type II slot. Samsung calls the same form factor "1.3 inch" drive in its product literature.

- 0.85 inch: 24 mm × 5 mm × 32 mm

  Toshiba announced this form factor in January 2004 for use in mobile phones and similar applications, including SD/MMC slot compatible HDDs optimized for video storage on 4G handsets. Toshiba currently sells a 4 GB (MK4001MTD) and 8 GB (MK8003MTD) version and holds the Guinness World Record for the smallest hard disk drive.

- 3.5-inch and 2.5-inch hard disks currently dominate the market.

  By 2009 all manufacturers had discontinued the development of new products for the 1.3-inch, 1-inch and 0.85-inch form factors due to falling prices of flash memory, which is slightly more stable and resistant to damage from impact and/or dropping.

The inch-based nickname of all these form factors usually do not indicate any actual product dimension (which are specified in millimeters for more recent form factors), but just roughly indicate a size relative to disk diameters, in the interest of historic continuity.

PERFORMANCE CHARACTERISTICS

ACCESS TIME

The factors that limit the time to access the data on a hard disk drive (Access time) are mostly related to the mechanical nature of the rotating disks and moving heads. Seek time is a measure of how long it takes the head assembly to travel to the track of the disk that contains data. Latency is rotational delay incurred because the desired disk sector may not be directly under the head when data transfer is requested. These two delays are on the order of milliseconds each. The bit rate or data transfer rate once the head is in the right position creates delay which is a function of the number of blocks transferred; typically relatively small, but can be quite long with the transfer of large contiguous files. Delay may also occur if the drive disks are stopped to save energy, see Power management.

An HDD's Average Access Time is its average Seek time which technically is the time to do all possible seeks divided by the number of all possible seeks, but in practice is determined by statistical methods or simply approximated as the time of a seek over one-third of the number of tracks

Defragmentation is a procedure used to minimize delay in retrieving data by moving related items to physically proximate areas on the disk. Some computer operating systems perform defragmentation automatically. Although automatic defragmentation is intended to reduce access delays, the procedure can slow response when performed while the computer is in use.

Access time can be improved by increasing rotational speed, thus reducing latency and/or by decreasing seek time. Increasing areal density increases throughput by increasing data rate and by increasing the amount of data under a set of heads, thereby potentially reducing seek activity for a given amount of data. Based on historic trends, analysts predict a future growth in HDD areal density (and therefore capacity) of about 40% per year. Access times have not kept up with throughput increases, which themselves have not kept up with growth in storage capacity.

SEEK TIME

Average Seek time ranges from 3 ms for high-end server drives, to 15 ms for mobile drives, with the most common mobile drives at about 12 ms and the most common desktop type typically being around 9 ms. The first HDD had an average seek time of about 600 ms and by the middle 1970s HDDs were available with seek times of about 25 ms. Some early PC drives used a stepper motor to move the heads, and as a result had seek times as slow as 80–120 ms, but this was quickly improved by voice coil type actuation in the 1980s, reducing seek times to around 20 ms. Seek time has continued to improve slowly over time.

LATENCY

Latency is the delay for the rotation of the disk to bring the required disk sector under the read-write mechanism. It depends on rotational speed of a disk, measured in revolutions per minute (RPM). Average rotational delay is shown in the table below, based on the empirical relation that the average latency in milliseconds for such a drive is one-half the rotational period:

Spindle [rpm]	*Average latency [ms]*
4200	7.14
5400	5.56
7200	4.17
10000	3
15000	2

DATA TRANSFER RATE

As of 2010, a typical 7200 rpm desktop hard drive has a sustained “disk-to-buffer” data transfer rate up to 1030 Mbits/sec. This rate depends on the track location, so it will be higher for data on the outer tracks (where there are more data sectors) and lower toward the inner tracks (where there are fewer data sectors); and is generally somewhat higher for 10,000 rpm drives. A current widely used standard for the “buffer-to-computer” interface is 3.0 Gbit/s SATA, which can send about 300 megabyte/s from the buffer to the computer, and thus is still comfortably ahead of today’s disk-to-buffer transfer rates. Data transfer rate (read/write) can be measured by writing a large file to disk using special file generator tools, then reading back the file. Transfer rate can be influenced by file system fragmentation and the layout of the files.

HDD data transfer rate depends upon the rotational speed of the platters and the data recording density. Because heat and vibration limit rotational speed, advancing density becomes the main method to improve sequential transfer rates. Areal density advances by increasing both the number of tracks across the disk and the number of sectors per track, the later will increase the data transfer rate (for a given RPM). Since data transfer rate performance only tracks one of the two components of areal density, its performance improves at lower rate,

POWER CONSUMPTION

Power consumption has become increasingly important, not only in mobile devices such as laptops but also in server and desktop markets. Increasing data center machine density has led to problems delivering sufficient power to devices (especially for spin up), and getting rid of the waste heat subsequently produced, as well as environmental and electrical cost concerns. Heat dissipation directly tied to power consumption, and as drive age, disk failure rates increase at higher drive temperatures. Similar issues exist for large companies with thousands of desktop PCs. Smaller form factor drives often use less power than larger drives. One interesting development in this area is actively controlling the seek speed so that the head arrives at its destination only just in time to read the sector, rather than arriving as quickly as possible and then having to wait for the sector to come around (i.e. the rotational latency). Many of the hard drive companies are now producing Green Drives that require much less power and cooling. Many of these Green Drives spin slower (<5,400 rpm compared to 7,200, 10,000 or 15,000 rpm) and also generate less waste heat. Power consumption can also be reduced by parking the drive heads when the disk is not in use reducing friction, adjusting spin speeds according to transfer rates, and disabling internal components when not in use. Also in systems where there might be multiple hard disk drives, there are various ways of controlling when the hard drives spin up since the highest current is drawn at that time.

- On SCSI hard disk drives, the SCSI controller can directly control spin up and spin down of the drives.
- On Parallel ATA (aka PATA) and Serial ATA (SATA) hard disk drives, some support power-up in standby or PUIS. The hard disk drive will not spin up until the controller or system BIOS issues a specific command to do so. This limits the power draw or consumption upon power on.
- Some SATA II hard disk drives support staggered spin-up, allowing the computer to spin up the drives in sequence to reduce load on the power supply when booting.

POWER MANAGEMENT

Most hard disk drives today support some form of power management which uses a number of specific power modes that save energy by reducing performance. When implemented an HDD will change between a full power mode to one or more power saving modes as a function of drive usage. Recovery from the deepest mode, typically called Sleep, may take as long as several seconds.

AUDIBLE NOISE

Measured in dBA, audible noise is significant for certain applications, such as DVRs, digital audio recording and quiet computers. Low noise disks typically use fluid bearings, slower rotational speeds (usually 5,400 rpm) and reduce the seek speed under load (AAM) to reduce audible clicks and crunching sounds. Drives in smaller form factors (e.g. 2.5 inch) are often quieter than larger drives.

SHOCK RESISTANCE

Shock resistance is especially important for mobile devices. Some laptops now include active hard drive protection that parks the disk heads if the machine is dropped, hopefully before impact, to offer the greatest possible chance of survival in such an event. Maximum shock tolerance to date is 350 g for operating and 1000 g for non-operating.

ACCESS AND INTERFACES

Hard disk drives are accessed over one of a number of bus types, including parallel ATA (P-ATA, also called IDE or EIDE), Serial ATA (SATA), SCSI, Serial Attached SCSI (SAS), and Fibre Channel. Bridge circuitry is sometimes used to connect hard disk drives to buses that they cannot communicate with natively, such as IEEE 1394, USB and SCSI.

For the ST-506 interface, the data encoding scheme as written to the disk surface was also important. The first ST-506 disks used Modified Frequency Modulation (MFM) encoding, and transferred data at a rate of 5 megabits per second. Later controllers using 2,7 RLL (or just "RLL") encoding caused 50% more data to appear under the heads compared to one rotation of an MFM drive, increasing data storage and data transfer rate by 50%, to 7.5 megabits per second.

Many ST-506 interface disk drives were only specified by the manufacturer to run at the 1/3 lower MFM data transfer rate compared to RLL, while other drive models (usually more expensive versions of the same drive) were specified to run at the higher RLL data transfer rate. In some cases, a drive had sufficient margin to allow the MFM specified model to run at the denser/faster RLL data transfer rate (not recommended nor guaranteed by manufacturers). Also, any RLL-certified drive could run on any MFM controller, but with 1/3 less data capacity and as much as 1/3 less data transfer rate compared to its RLL specifications.

Enhanced Small Disk Interface (ESDI) also supported multiple data rates (ESDI disks always used 2,7 RLL, but at 10, 15 or 20 megabits per second), but this was usually negotiated automatically by the disk drive and controller; most of the time, however, 15 or 20 megabit ESDI disk drives were not downward compatible (i.e. a 15 or 20 megabit disk drive would not run on a 10 megabit controller). ESDI disk drives typically also had jumpers to set the number of sectors per track and (in some cases) sector size.

Modern hard drives present a consistent interface to the rest of the computer, no matter what data encoding scheme is used internally. Typically a DSP in the electronics inside the hard drive takes the raw analog voltages from the read head and uses PRML and Reed–Solomon error correction to decode the sector boundaries and sector data, then sends that data out the standard interface. That DSP also watches the error rate detected by error detection and correction, and performs bad sector remapping, data collection for Self-Monitoring, Analysis, and Reporting Technology, and other internal tasks. SCSI originally had just one signaling frequency of 5 MHz for a maximum data rate of 5 megabytes/ second over 8 parallel conductors, but later this was increased dramatically. The SCSI bus speed had no bearing on the disk's internal speed because of buffering between the SCSI bus and the disk drive's internal data bus; however, many early disk drives had very small buffers, and thus had to be reformatted to a different interleave (just like ST-506 disks) when used on slow computers, such as early Commodore Amiga, IBM PC compatibles and Apple Macintoshes.

ATA disks have typically had no problems with interleave or data rate, due to their controller design, but many early models were incompatible with each other and could not run with two devices on the same physical cable in a master/slave setup. This was mostly remedied by the mid-1990s, when ATA's specification was standardized and the details began to be cleaned up, but still causes problems occasionally (especially with CD-ROM and DVD-ROM disks, and when mixing Ultra DMA and non-UDMA devices).

Serial ATA does away with master/slave setups entirely, placing each disk on its own channel (with its own set of I/O ports) instead. FireWire/IEEE 1394 and USB(1.0/2.0) HDDs are external units containing generally ATA or SCSI disks with ports on the back allowing very simple and effective expansion and mobility.

Most FireWire/IEEE 1394 models are able to daisy-chain in order to continue adding peripherals without requiring additional ports on the computer itself. USB however, is a point to point network and does not allow for daisy-chaining. USB hubs are used to increase the number of available ports and are used for devices that do not require charging since the current supplied by hubs is typically lower than what's available from the built-in USB ports.

DISK INTERFACE FAMILIES USED IN PERSONAL COMPUTERS

Notable families of disk interfaces include:

- Historical bit serial interfaces connect a hard disk drive (HDD) to a hard disk controller (HDC) with two cables, one for control and one for data. (Each drive also has an additional cable for power, usually connecting it directly to the power supply unit). The HDC provided significant functions such as serial/ parallel conversion, data separation, and track formatting, and required matching to the drive (after formatting) in order to assure reliability. Each control cable could serve two or more drives, while a dedicated (and smaller) data cable served each drive.
  - ST506 used MFM (Modified Frequency Modulation) for the data encoding method.
  - ST412 was available in either MFM or RLL (Run Length Limited) encoding variants.
  - Enhanced Small Disk Interface (ESDI) was an industry standard interface similar to ST412 supporting higher data rates between the processor and the disk drive.
- Modern bit serial interfaces connect a hard disk drive to a host bus interface adapter (today typically

integrated into the "south bridge") with one data/ control cable. (As for historical *bit serial interfaces* above, each drive also has an additional power cable, usually direct to the power supply unit.)

- o Fibre Channel (FC), is a successor to parallel SCSI interface on enterprise market. It is a serial protocol. In disk drives usually the Fibre Channel Arbitrated Loop (FC-AL) connection topology is used. FC has much broader usage than mere disk interfaces, and it is the cornerstone of storage area networks (SANs). Recently other protocols for this field, like iSCSI and ATA over Ethernet have been developed as well. Confusingly, drives usually use *copper* twisted-pair cables for Fibre Channel, not fibre optics. The latter are traditionally reserved for larger devices, such as servers or disk array controllers.
- o Serial ATA (SATA). The SATA data cable has one data pair for differential transmission of data to the device, and one pair for differential receiving from the device, just like EIA-422. That requires that data be transmitted serially. Similar differential signaling system is used in RS485, LocalTalk, USB, Firewire, and differential SCSI.
- o Serial Attached SCSI (SAS). The SAS is a new generation serial communication protocol for devices designed to allow for much higher speed data transfers and is compatible with SATA. SAS uses a mechanically identical data and power connector to standard 3.5-inch SATA1/SATA2 HDDs, and many server-oriented SAS RAID

controllers are also capable of addressing SATA hard drives. SAS uses serial communication instead of the parallel method found in traditional SCSI devices but still uses SCSI commands.

- Word serial interfaces connect a hard disk drive to a host bus adapter (today typically integrated into the "south bridge") with one cable for combined data/ control. (As for all *bit serial interfaces* above, each drive also has an additional power cable, usually direct to the power supply unit.) The earliest versions of these interfaces typically had a 8 bit parallel data transfer to/from the drive, but 16-bit versions became much more common, and there are 32 bit versions. Modern variants have serial data transfer. The word nature of data transfer makes the design of a host bus adapter significantly simpler than that of the precursor HDD controller.
    - o Integrated Drive Electronics (IDE), later renamed to ATA, with the alias P-ATA ("parallel ATA") retroactively added upon introduction of the new variant Serial ATA. The original name reflected the innovative integration of HDD controller with HDD itself, which was not found in earlier disks. Moving the HDD controller from the interface card to the disk drive helped to standardize interfaces, and to reduce the cost and complexity. The 40-pin IDE/ ATA connection transfers 16 bits of data at a time on the data cable. The data cable was originally 40-conductor, but later higher speed requirements for data transfer to and from the hard drive led to an

"ultra DMA" mode, known as UDMA. Progressively swifter versions of this standard ultimately added the requirement for a 80-conductor variant of the same cable, where half of the conductors provides grounding necessary for enhanced high-speed signal quality by reducing cross talk. The interface for 80-conductor only has 39 pins, the missing pin acting as a key to prevent incorrect insertion of the connector to an incompatible socket, a common cause of disk and controller damage.

- EIDE was an unofficial update (by Western Digital) to the original IDE standard, with the key improvement being the use of direct memory access (DMA) to transfer data between the disk and the computer without the involvement of the CPU, an improvement later adopted by the official ATA standards. By directly transferring data between memory and disk, DMA eliminates the need for the CPU to copy byte per byte, therefore allowing it to process other tasks while the data transfer occurs.
- Small Computer System Interface (SCSI), originally named SASI for Shugart Associates System Interface, was an early competitor of ESDI. SCSI disks were standard on servers, workstations, Commodore Amiga, and Apple Macintosh computers through the mid-1990s, by which time most models had been transitioned to IDE (and later, SATA) family disks. Only in 2005 did the capacity of SCSI disks fall behind IDE disk technology, though the highest-performance disks

are still available in SCSI and Fibre Channel only. The range limitations of the data cable allows for external SCSI devices. Originally SCSI data cables used single ended (common mode) data transmission, but server class SCSI could use differential transmission, either low voltage differential (LVD) or high voltage differential (HVD). ("Low" and "High" voltages for differential SCSI are relative to SCSI standards and do not meet the meaning of low voltage and high voltage as used in general electrical engineering contexts, as apply e.g. to statutory electrical codes; both LVD and HVD use low voltage signals (3.3 V and 5 V respectively) in general terminology.)

INTEGRITY

Due to the extremely close spacing between the heads and the disk surface, hard disk drives are vulnerable to being damaged by a head crash—a failure of the disk in which the head scrapes across the platter surface, often grinding away the thin magnetic film and causing data loss. Head crashes can be caused by electronic failure, a sudden power failure, physical shock, contamination of the drive's internal enclosure, wear and tear, corrosion, or poorly manufactured platters and heads.

The HDD's spindle system relies on air pressure inside the disk enclosure to support the heads at their proper *flying height* while the disk rotates. Hard disk drives require a certain range of air pressures in order to operate properly. The connection to the external environment and pressure occurs through a small hole in the enclosure (about 0.5 mm in breadth), usually with a filter on the inside (the *breather filter*). If the air pressure is too low, then there is not enough lift for the flying head, so the head gets too close to the disk, and there is a risk of head crashes and data loss. Specially manufactured sealed and pressurized disks are needed for reliable high-altitude operation, above about 3,000 m (10,000 feet). Modern disks include temperature sensors and adjust their operation to the operating environment. Breather holes can be seen on all disk drives—they usually have a sticker next to them, warning the user not to cover the holes.

The air inside the operating drive is constantly moving too, being swept in motion by friction with the spinning platters. This air passes through an internal recirculation (or "recirc") filter to remove any leftover contaminants from manufacture, any particles or chemicals that may have somehow entered the enclosure, and any particles or outgassing generated internally in normal operation. Very high humidity for extended periods can corrode the heads and platters.

For giant magnetoresistive (GMR) heads in particular, a minor head crash from contamination (that does not remove the magnetic surface of the disk) still results in the head temporarily overheating, due to friction with the disk surface, and can render the data unreadable for a short period until the head temperature stabilizes (so called "thermal asperity", a problem which can partially be dealt with by proper electronic filtering of the read signal).

ACTUATION OF MOVING ARM

The hard drive's electronics control the movement of the actuator and the rotation of the disk, and perform reads and writes on demand from the disk controller. Feedback of the drive electronics is accomplished by means of special segments of the disk dedicated to servo feedback. These are either complete concentric circles (in the case of dedicated servo technology), or segments interspersed with real data (in the case of embedded servo technology). The servo feedback optimizes the signal to noise ratio of the GMR sensors by adjusting the voice-coil of the actuated arm. The spinning of the disk also uses a servo motor. Modern disk firmware is capable of scheduling reads and writes efficiently on the platter surfaces and remapping sectors of the media which have failed.

LANDING ZONES AND LOAD/UNLOAD TECHNOLOGY

Modern HDDs prevent power interruptions or other malfunctions from landing its heads in the data zone by parking the heads either in a landing zone or by unloading (i.e., load/unload) the heads. Some early PC HDDs did not park the heads automatically and they would land on data. In some other early units the user manually parked the heads by running a programme to park the HDD's heads.

A landing zone is an area of the platter usually near its inner diameter (ID), where no data are stored. This area is called the Contact Start/Stop (CSS) zone. Disks are designed such that either a spring or, more recently, rotational inertia in the platters is used to park the heads in the case of unexpected power loss. In this case, the spindle motor temporarily acts as a generator, providing power to the actuator.

Spring tension from the head mounting constantly pushes the heads towards the platter. While the disk is spinning, the heads are supported by an air bearing and experience no physical contact or wear. In CSS drives the sliders carrying the head sensors (often also just called *heads*) are designed to survive a number of landings and takeoffs from the media surface, though wear and tear on these microscopic components eventually takes its toll. Most manufacturers design the sliders to survive 50,000 contact cycles before the chance of damage on startup rises above 50%. However, the decay rate is not linear: when a disk is younger and has had fewer start-stop cycles, it has a better chance of surviving the next startup than an older, higher-mileage disk (as the head literally drags along the disk's surface until the air bearing is established). For example, the Seagate Barracuda 7200.10 series of desktop hard disks are rated to 50,000 start-stop cycles, in other words no failures attributed to the head-platter interface were seen before at least 50,000 start-stop cycles during testing.

Around 1995 IBM pioneered a technology where a landing zone on the disk is made by a precision laser process (*Laser Zone Texture* = LZT) producing an array of smooth nanometer-scale "bumps" in a landing zone, thus vastly improving stiction and wear performance. This technology is still largely in use today (2008), predominantly in desktop and enterprise (3.5 inch) drives. In general, CSS technology can be prone to increased stiction (the tendency for the heads to stick to the platter surface), e.g. as a consequence of increased humidity. Excessive stiction can cause physical damage to the platter and slider or spindle motor.

Load/Unload technology relies on the heads being lifted off the platters into a safe location, thus eliminating the risks of wear and stiction altogether. The first HDD RAMAC and most early disk drives used complex mechanisms to load and unload the heads. Modern HDDs use ramp loading, first introduced by Memorex in 1967, to load/unload onto plastic "ramps" near the outer disk edge.

All HDDs today still use one of these two technologies listed above. Each has a list of advantages and drawbacks in terms of loss of storage area on the disk, relative difficulty of mechanical tolerance control, non-operating shock robustness, cost of implementation, etc.

Addressing shock robustness, IBM also created a technology for their ThinkPad line of laptop computers called the Active Protection System. When a sudden, sharp movement is detected by the built-in accelerometer in the Thinkpad, internal hard disk heads automatically unload themselves to reduce the risk of any potential data loss or scratch defects. Apple later also utilized this technology in their PowerBook, iBook, MacBook Pro, and MacBook line, known as the Sudden Motion Sensor. Sony, HP with their HP 3D DriveGuard and Toshiba have released similar technology in their notebook computers.

This accelerometer-based shock sensor has also been used for building cheap earthquake sensor networks.

DISK FAILURES AND THEIR METRICS

Most major hard disk and motherboard vendors now support S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology), which measures drive characteristics such as operating temperature, spin-up time, data error rates, etc. Certain trends and sudden changes in these parameters are thought to be associated with increased likelihood of drive failure and data loss.

However, not all failures are predictable. Normal use eventually can lead to a breakdown in the inherently fragile device, which makes it essential for the user to periodically back up the data onto a separate storage device. Failure to do so can lead to the loss of data. While it may sometimes be possible to recover lost information, it is normally an extremely costly procedure, and it is not possible to guarantee success. A 2007 study published by Google suggested very little correlation between failure rates and either high temperature or activity level; however, the correlation between manufacturer/model and failure rate was relatively strong. Statistics in this matter is kept highly secret by most entities. Google did not publish the manufacturer's names along with their respective failure rates, though they have since revealed that they use Hitachi Deskstar drives in some of their servers. While several S.M.A.R.T. parameters have an impact on failure probability, a large fraction of failed drives do not produce predictive S.M.A.R.T. parameters. S.M.A.R.T. parameters alone may not be useful for predicting individual drive failures.

A common misconception is that a colder hard drive will last longer than a hotter hard drive. The Google study seems to imply the reverse—"lower temperatures are associated with higher failure rates". Hard drives with S.M.A.R.T.-reported average temperatures below 27 °C (80.6 °F) had higher failure rates than hard drives with the highest reported average temperature of 50 °C (122 °F), failure rates at least twice as high as the optimum S.M.A.R.T.-reported temperature range of 36 °C (96.8 °F) to 47 °C (116.6 °F).

SCSI, SAS, and FC drives are typically more expensive and are traditionally used in servers and disk arrays, whereas inexpensive ATA and SATA drives evolved in the home computer market and were perceived to be less reliable. This distinction is now becoming blurred.

The mean time between failures (MTBF) of SATA drives is usually about 600,000 hours (some drives such as Western Digital Raptor have rated 1.4 million hours MTBF), while SCSI drives are rated for upwards of 1.5 million hours. However, independent research indicates that MTBF is not a reliable estimate of a drive's longevity. MTBF is conducted in laboratory environments in test chambers and is an important metric to determine the quality of a disk drive before it enters high volume production. Once the drive product is in production, the more valid metric is annualized failure rate (AFR). AFR is the percentage of real-world drive failures after shipping.

SAS drives are comparable to SCSI drives, with high MTBF and high reliability.

Enterprise S-ATA drives designed and produced for enterprise markets, unlike standard S-ATA drives, have reliability comparable to other enterprise class drives. Typically enterprise drives (all enterprise drives, including SCSI, SAS, enterprise SATA, and FC) experience between 0.70%–0.78% annual failure rates from the total installed drives.

Eventually all mechanical hard disk drives fail, so to mitigate loss of data, some form of redundancy is needed, such as RAID or a regular backup system.

4

Data Store and Design

A data store is a data repository of a set of integrated objects. These objects are modeled using classes defined in database schemas. Data store includes not only data repositories like databases, it is a more general concept that includes also flat files that can store data. Some data stores do represent data in only one schema, while other data stores use several schemas for this task. An example are RDBMS-based data stores like MySQL or ORACLE.

TYPES

Data stores can be of different types.

- The widely used data store type are relational databases. (Examples: Microsoft SQL Server, MySQL, ORACLE, Postgresql)
- Object-oriented databases, like Caché or ConceptBase. They can save objects of an object-oriented design.

- Operational data stores
- Schemaless data stores, like Distributed data stores (e.g. Apache Cassandra or Dynamo).

DATABASE DESIGN

Database design is the process of producing a detailed data model of a database. This logical data model contains all the needed logical and physical design choices and physical storage parameters needed to generate a design in a Data Definition Language, which can then be used to create a database. A fully attributed data model contains detailed attributes for each entity. The term database design can be used to describe many different parts of the design of an overall database system. Principally, and most correctly, it can be thought of as the logical design of the base data structures used to store the data. In the relational model these are the tables and views. In an object database the entities and relationships map directly to object classes and named relationships. However, the term database design could also be used to apply to the overall process of designing, not just the base data structures, but also the forms and queries used as part of the overall database application within the database management system (DBMS). The process of doing database design generally consists of a number of steps which will be carried out by the database designer. Usually, the designer must:

- Determine the relationships between the different data elements.
- Superimpose a logical structure upon the data on the basis of these relationships.

ER DIAGRAM (ENTITY-RELATIONSHIP MODEL)

Database designs also include ER(Entity-relationship model) diagrams. An ER diagram is a diagram that helps to design databases in an efficient way. Attributes in ER diagrams are usually modeled as an oval with the name of the attribute, linked to the entity or relationship that contains the attribute. Within the relational model the final step can generally be broken down into two further steps, that of determining the grouping of information within the system, generally determining what are the basic objects about which information is being stored, and then determining the relationships between these groups of information, or objects. This step is not necessary with an Object database.

THE DESIGN PROCESS

The design process consists of the following steps:

1. Determine the purpose of your database - This helps prepare you for the remaining steps.
2. Find and organize the information required - Gather all of the types of information you might want to record in the database, such as product name and order number.
3. Divide the information into tables - Divide your information items into major entities or subjects, such as Products or Orders. Each subject then becomes a table.
4. Turn information items into columns - Decide what information you want to store in each table. Each item becomes a field, and is displayed as a column in the table. For example, an Employees table might include fields such as Last Name and Hire Date.

5. Specify primary keys - Choose each table's primary key. The primary key is a column that is used to uniquely identify each row. An example might be Product ID or Order ID.
6. Set up the table relationships - Look at each table and decide how the data in one table is related to the data in other tables. Add fields to tables or create new tables to clarify the relationships, as necessary.
7. Refine your design - Analyze your design for errors. Create the tables and add a few records of sample data. See if you can get the results you want from your tables. Make adjustments to the design, as needed.
8. Apply the normalization rules - Apply the data normalization rules to see if your tables are structured correctly. Make adjustments to the tables

DETERMINING DATA TO BE STORED

In a majority of cases, a person who is doing the design of a database is a person with expertise in the area of database design, rather than expertise in the domain from which the data to be stored is drawn e.g. financial information, biological information etc. Therefore the data to be stored in the database must be determined in cooperation with a person who does have expertise in that domain, and who is aware of what data must be stored within the system. This process is one which is generally considered part of requirements analysis, and requires skill on the part of the database designer to elicit the needed information from those with the domain knowledge. This is because those with the necessary domain knowledge frequently cannot express clearly what their system requirements for the database are as they are unaccustomed to thinking in terms of the discrete data elements which must be stored. Data to be stored can be determined by Requirement Specification.

NORMALIZATION

In the field of relational database design, normalization is a systematic way of ensuring that a database structure is suitable for general-purpose querying and free of certain undesirable characteristics—insertion, update, and deletion anomalies—that could lead to a loss of data integrity. A standard piece of database design guidance is that the designer should create a fully normalized design; selective denormalization can subsequently be performed, but only for performance reasons. However, some modeling disciplines, such as the dimensional modeling approach to data warehouse design, explicitly recommend non-normalized designs, i.e. designs that in large part do not adhere to 3NF.

TYPES OF DATABASE DESIGN

CONCEPTUAL SCHEMA

Once a database designer is aware of the data which is to be stored within the database, they must then determine where dependancy is within the data. Sometimes when data is changed you can be changing other data that is not visible. For example, in a list of names and addresses, assuming a situation where multiple people can have the same address, but one person cannot have more than one address, the name is dependent upon the address, because if the address is different, then the associated name is different too.

However, the other way around is different. One attribute can change and not another. (*NOTE:* A common misconception is that the relational model is so called because of the stating of relationships between data elements therein. This is not true. The relational model is so named because it is based upon the mathematical structures known as relations.)

LOGICALLY STRUCTURING DATA

Once the relationships and dependencies amongst the various pieces of information have been determined, it is possible to arrange the data into a logical structure which can then be mapped into the storage objects supported by the database management system. In the case of relational databases the storage objects are tables which store data in rows and columns.

Each table may represent an implementation of either a logical object or a relationship joining one or more instances of one or more logical objects. Relationships between tables may then be stored as links connecting child tables with parents. Since complex logical relationships are themselves tables they will probably have links to more than one parent. In an Object database the storage objects correspond directly to the objects used by the Object-oriented programming language used to write the applications that will manage and access the data. The relationships may be defined as attributes of the object classes involved or as methods that operate on the object classes.

PHYSICAL DATABASE DESIGN

The physical design of the database specifies the physical configuration of the database on the storage media. This includes detailed specification of data elements, data types, indexing options and other parameters residing in the DBMS data dictionary. It is the detailed design of a system that includes modules & the database's hardware & software specifications of the system.

DATABASE THEORY

Database theory encapsulates a broad range of topics related to the study and research of the theoretical realm of databases and database management systems. Theoretical aspects of data management include, among other areas, the foundations of query languages, computational complexity and expressive power of queries, finite model theory, database design theory, dependency theory, foundations of concurrency control and database recovery, deductive databases, temporal and spatial databases, real time databases, managing uncertain data and probabilistic databases, and Web data. Most research work has traditionally been based on the relational model, since this model is usually considered the simplest and most foundational model of interest. Corresponding results for other data models, such as object-oriented or semi-structured models, or, more recently, graph data models and XML, are often derivable from those for the relational model. A central focus of database theory is on understanding the complexity and power of query languages and their connection to logic. Starting from relational algebra and first-order logic (which are equivalent by Codd's theorem) and the insight that important queries such as graph reachability are not expressible in this language, more powerful language based on logic programming and fixpoint logic such as datalog were studied.

Another focus was on the foundations of query optimization and data integration. Here most work studied conjunctive queries, which admit query optimization even under constraints using the chase algorithm. The main research conferences in the area are the ACM Symposium on Principles of Database Systems (PODS) and the International Conference on Database Theory (ICDT).

5

Computer DVD Storage Device

DVD (Digital Versatile Disk) is an optical disc storage media format, invented and developed by Philips, Sony, Toshiba, and Time Warner in 1995. DVDs offer higher storage capacity than compact discs while having the same dimensions. Pre-recorded DVDs are mass-produced using molding machines that physically stamp data onto the DVD. Such discs are known as DVD-ROM, because data can only be read and not written nor erased. Blank recordable DVDs (DVD-R and DVD+R) can be recorded once using a DVD recorder and then function as a DVD-ROM. Rewritable DVDs (DVD-RW, DVD+RW, and DVD-RAM) can be recorded and erased multiple times. DVDs are used in DVD-Video consumer digital video format and in DVD-Audio consumer digital audio format, as well as for authoring AVCHD discs. DVDs containing other types of information may be referred to as DVD data discs.

HISTORY

In 1993, two optical disc storage formats were being developed. One was the MultiMedia Compact Disc (MMCD) also called CDi, backed by Philips and Sony, and the other was the Super Density (SD) disc, supported by Toshiba, Time Warner, Matsushita Electric, Hitachi, Mitsubishi Electric, Pioneer, Thomson, and JVC.

Representatives of the SD camp approached IBM, asking for advice on the file system to use for their disc as well as seeking support for their format for storing computer data. Alan E. Bell, a researcher from IBM's Almaden Research Center got that request and also learned of the MMCD development project. Wary of being caught in a repeat of the costly videotape format war between VHS and Betamax in the 1980s, he convened a group of computer industry experts, including representatives from Apple, Microsoft, Sun, Dell, and many others. This group was referred to as the Technical Working Group, or TWG.

The TWG voted to boycott both formats unless the two camps agreed on a single, converged standard. Lou Gerstner, president of IBM, was recruited to apply pressure on the executives of the warring factions. Eventually, the computer companies won the day, and a single format, now called DVD, was agreed upon. The TWG also collaborated with the Optical Storage Technology Association (OSTA) on the use of their implementation of the ISO-13346 file system (known as Universal Disc Format) for use on the new DVDs.

Philips and Sony decided it was in their best interest to avoid another format war over their MultiMedia Compact Disc, and agreed to unify with companies backing the Super

Density Disc to release a single format with technologies from both. The specification was mostly similar to Toshiba and Matsushita's Super Density Disc, except for the dual-layer option (MMCD was single-sided and optionally dual-layer, whereas SD was single-layer but optionally double-sided) and EFMPlus modulation.

EFMPlus was chosen because of its great resilience to disc damage, such as scratches and fingerprints. EFMPlus, created by Kees Immink (who also designed EFM), is 6% less efficient than the modulation technique originally used by Toshiba, which resulted in a capacity of 4.7 GB, as opposed to the original 5 GB. The result was the DVD specification, finalized for the DVD movie player and DVD-ROM computer applications in December 1995.

The DVD Video format was first introduced by Toshiba in Japan in November 1996, in the United States in March 1997 (test marketed), in Europe in October 1998, and in Australia in February 1999.

In May 1997, the DVD Consortium was replaced by the DVD Forum, which is open to all other companies.

SPECIFICATIONS

DVD specifications created and updated by the DVD Forum are published as so-called *DVD Books* (e.g. DVD-ROM Book, DVD-Audio Book, DVD-Video Book, DVD-R Book, DVD-RW Book, DVD-RAM Book, DVD-AR Book, DVD-VR Book, etc.).

Some specifications for mechanical, physical and optical characteristics of DVD optical discs can be downloaded as *freely available standards* from the ISO website. Also, the DVD+RW Alliance publishes competing DVD specifications such as DVD+R, DVD+R DL, DVD+RW or DVD+RW DL. These DVD formats are also ISO standards.

Some of DVD specifications (e.g. for DVD-Video) are not publicly available and can be obtained only from the DVD Format/Logo Licensing Corporation for a fee of US $5000. Every subscriber must sign a non-disclosure agreement as certain information in the DVD Book is proprietary and confidential.

ETYMOLOGY

The official DVD charter documents specify that the basis of the DVD name stems from the term "digital versatile disc". Usage in the present day varies, with *Digital Versatile Disc*, *Digital Video Disc*, and *DVD* being the most common.

DVD was originally used as an initialism for the unofficial term *digital videodisk*.

A newsgroup FAQ written by Jim Taylor (a prominent figure in the industry) claims that four years later, in 1999, the DVD Forum stated that the format name was simply the three letters "DVD" and did not stand for anything.

The DVD Forum website has a section called "DVD Primer" in which the answer to the question, "What does DVD mean?" reads, "The keyword is 'versatile.' Digital Versatile Discs provide superb video, audio and data storage and access—all on one disc."

IDENTIFICATION (MID)

The DVD is made of a spiral groove read or written starting at the center. The form of the groove encodes unalterable identification data known as Media Identification Code (MID). The MID contains data such as the manufacturer and model, byte capacity, allowed data rates (also known as speed), etc.

DESIGN

AS A MOVIE DELIVERY MEDIUM

DVD was adopted by movie and home entertainment distributors to replace the ubiquitous VHS tape as the primary means of distributing films to consumers in the home entertainment marketplace. DVD was chosen for its superior ability to reproduce moving pictures and sound, for its superior durability, and for its interactivity. Interactivity had proven to be a feature which consumers, especially collectors, favoured when the movie studios had released their films on laser disk. When the price point for a laser disk at approximately $100 per disk moved to $20 per disk at retail, this luxury feature became available for mass consumption. Simultaneously the movie studios decided to change their home entertainment release model from a rental model to a for purchase model, and large numbers of dvds were sold.

At the same time, a demand for interactive design talent and services was created. Movies in the past had uniquely designed title sequences. Suddenly every movie being released required information architecture and interactive design components that matched the film's tone and were at the quality level that Hollywood demanded for its product. Whole mini-studios to perform this type of work, such as

1K Studios and Canned Interactive, either formed or created service departments to fulfill this need in Hollywood.

New DVD releases are released weekly by all major studios. DVDs are typically released on Tuesdays of every week. With the advent of Blu-ray releases, studios now rely on both Blu-ray and DVDs to supplement their revenue for a particular movie.

AS AN INTERACTIVE MEDIUM

DVD as a format had two qualities at the time that were not available in any other interactive medium: 1. Enough capacity and speed to provide high quality, full motion video and sound, and 2. low cost delivery mechanism provided by consumer products retailers who quickly moved to sell their players for under $200 and eventually for under $50 at retail. In addition, the medium itself was small enough and light enough to mail using general first class postage. Almost over night, this created a new business opportunity and model for business innovators like Netflix to re-invent the home entertainment distribution model. It also opened up the opportunity for business and product information to be inexpensively provided on full motion video through direct mail.

Again, a demand for unique interactive design talent and services was created. Companies as large as Nike and Mattel went in search of new start-ups such as 1K Studios and Canned Interactive to fulfill their information architecture and design needs.

The basic types of DVD (12 cm diameter, single-sided or homogeneous double-sided) are referred to by a rough approximation of their capacity in gigabytes. In draft versions of the specification, DVD-5 indeed held five gigabytes, but some parameters were changed later on as explained above, so the capacity decreased. Other formats, those with 8 cm diameter and hybrid variants, acquired similar numeric names with even larger deviation.

The 12 cm type is a standard DVD, and the 8 cm variety is known as a MiniDVD. These are the same sizes as a standard CD and a mini-CD, respectively. The capacity by surface (MiB/cm) varies from 6.92 MiB/cm^2 in the DVD-1 to 18.0 MiB/cm^2 in the DVD-18.

As with hard disk drives, in the DVD realm, gigabyte and the symbol GB are usually used in the SI sense (i.e., 10^9, or 1,000,000,000 bytes). For distinction, gibibyte (with symbol GiB) is used (i.e., 1024^3 (2^{30}), or 1,073,741,824 bytes). Each DVD sector contains 2,418 bytes of data, 2,048 bytes of which are user data. There is a small difference in storage space between + and - (hyphen) formats:

TECHNOLOGY

DVD uses 650 nm wavelength laser diode light as opposed to 780 nm for CD. This permits a smaller pit to be etched on the media surface compared to CDs (0.74 μm for DVD versus 1.6 μm for CD), allowing in part for a DVD's increased storage capacity.

In comparison, Blu-ray Disc, the successor to the DVD format, uses a wavelength of 405 nm, and one dual-layer disc has a 50 GB storage capacity. Writing speeds for DVD were 1×, that is, 1350 kB/s (1,318 KiB/s), in the first drives and media models. More recent models, at 18× or 20×, have 18 or 20 times that speed. Note that for CD drives, 1× means 153.6 kB/s (150 KiB/s), about one-ninth as swift.

INTERNAL MECHANISM OF A DRIVE

This mechanism is shown right side up; the disc would sit on top of it. The laser and optical system scans the underside of the disc.

With reference to the photo, just to the right of image center is the disc spin motor, a gray cylinder, with its gray centering hub and black resilient drive ring on top. A clamp (not in the photo, retained in the drive's cover), pulled down by a magnet, clamps the disc when this mechanism rises, after the disc tray stops moving inward. This motor has an external rotor – every part of it that you can see spins.

The gray metal chassis is shock-mounted at its four corners to reduce sensitivity to external shocks, and to reduce drive noise when running fast. The soft shock mount grommets are just below the brass-colored washers at the four corners (the left one is obscured). Running through those grommets are screws to fasten them to the black plastic frame that's underneath.

Two parallel precision guide rods that run between upper left and lower right in the photo carry the "sled", the moving optical read-write head. As shown, this "sled" is close to, or at the position where it reads or writes at the edge of the disc.

A dark gray disc with two holes on opposite sides has a blue lens surrounded by silver-colored metal. This is the lens that's closest to the disc; it serves to both read and write by focusing the laser light to a very small spot. It's likely that this disc rotates half a turn to position a different set of optics (the other "hole") for CDs vs. DVDs.

Under the disc is an ingenious actuator comprising permanent magnets and coils that move the lens up and down to maintain focus on the data layer. As well, the actuator moves the lens slightly toward and away from the spin-motor spindle to keep the spot on track. Both focus and tracking are relatively quite fast and very precise. The same actuator rotates the lens mount half a turn as described.

To select tracks (or files) as well as advancing the "sled" during continuous read or write operations, a stepping motor rotates a coarse-pitch leadscrew to move the "sled" throughout its total travel range. The motor, itself, is the gray cylinder just to the left of the most-distant shock mount; its shaft is parallel to the support rods. The leadscrew, itself, is the rod with evenly-spaced darker details; these are the helical groove that engages a pin on the "sled".

The irregular orange material is flexible etched copper foil supported by thin sheet plastic; these are "flexible printed circuits" that connect everything to the electronics (which is not shown).

DVD RECORDABLE AND REWRITABLE

HP initially developed recordable DVD media from the need to store data for backup and transport. DVD recordables are now also used for consumer audio and video recording. Three formats were developed: DVD-R/RW, DVD+R/RW (plus), and DVD-RAM. DVD-R is available in two formats, General (650 nm) and Authoring (635 nm), where Authoring discs may be recorded with encrypted content but General discs may not.

Although most DVD writers can nowadays write the DVD+R/RW and DVD-R/RW formats (usually denoted by "DVD±RW" and/or the existence of both the DVD Forum logo and the DVD+RW Alliance logo), the "plus" and the "dash" formats use different writing specifications. Most DVD readers and players will play both kinds of discs, although older models can have trouble with the "plus" variants. Some first generation DVD players would cause damage to DVD±R/RW/DL when attempting to read them.

DUAL-LAYER RECORDING

Dual-layer recording (sometimes also known as double-layer recording) allows DVD-R and DVD+R discs to store significantly more data—up to 8.54 gigabytes per disc, compared with 4.7 gigabytes for single-layer discs. Along with this, DVD-DLs have slower write speeds as compared to ordinary DVDs and when played on a DVD player, a slight transition can be seen between the layers. DVD-R DL was developed for the DVD Forum by Pioneer Corporation; DVD+R DL was developed for the DVD+RW Alliance by Philips and Mitsubishi Kagaku Media (MKM).

A dual-layer disc differs from its usual DVD counterpart by employing a second physical layer within the disc itself. The drive with dual-layer capability accesses the second layer by shining the laser through the first semitransparent layer. In some DVD players, the layer change can exhibit a noticeable pause, up to several seconds. This caused some viewers to worry that their dual-layer discs were damaged or defective, with the end result that studios began listing a standard message explaining the dual-layer pausing effect on all dual-layer disc packaging.

DVD recordable discs supporting this technology are backward-compatible with some existing DVD players and DVD-ROM drives. Many current DVD recorders support dual-layer technology, and the price is now comparable to that of single-layer drives, although the blank media remain more expensive. The recording speeds reached by dual-layer media are still well below those of single-layer media.

There are two modes for dual-layer orientation. With *Parallel Track Path* (PTP), used on DVD-ROM, both layers start at the inside diameter (ID) and end at the outside diameter (OD) with the lead-out. With *Opposite Track Path* (OTP), used on many Digital Video Discs, the lower layer starts at the ID and the upper layer starts at the OD, where the other layer ends; they share one lead-in and one lead-out.

DVD-VIDEO

DVD-Video is a standard for storing and distributing video/audio content on DVD media. The format went on sale in Japan on November 1, 1996, in the United States on March 1, 1997, in Europe on October 1, 1998 and in Australia on February 1, 1999.

DVD-Video became the dominant form of home video distribution in Japan when it first went on sale in 1996, but did not become the dominant form of home video distribution in the United States until June 15, 2003, when weekly DVD-Video in the United States rentals began outnumbering weekly VHS cassette rentals, reflecting the rapid adoption rate of the technology in the U.S. marketplace. Currently, DVD-Video is the dominant form of home video distribution worldwide, although in Japan it was surpassed by Blu-ray Disc when Blu-ray first went on sale in Japan on March 31, 2006.

CONSUMER RIGHTS

The rise of filesharing and “piracy” has prompted many copyright holders to display notices on DVD packaging or displayed on screen when the content is played that warn consumers of the illegality of certain uses of the DVD. It is commonplace to include a 90 second advert warning that most forms of copying the contents are illegal. Many DVDs prevent skipping past or fast-forwarding through this warning, forcing the consumer to watch. Arrangements for renting and lending differ by geography. In the U.S., the right to re-sell, rent, or lend out bought DVDs is protected by the first-sale doctrine under the Copyright Act of 1976. In Europe, rental and lending rights are more limited, under a 1992 European Directive that gives copyright holders broader powers to restrict the commercial renting and public lending of DVD copies of their work.

DVD-AUDIO

DVD-Audio is a format for delivering high fidelity audio content on a DVD. It offers many channel configuration options (from mono to 5.1 surround sound) at various sampling frequencies (up to 24-bits/192 kHz versus CDDA’s

16-bits/44.1 kHz). Compared with the CD format, the much higher-capacity DVD format enables the inclusion of considerably more music (with respect to total running time and quantity of songs) and/or far higher audio quality (reflected by higher sampling rates and greater sample resolution, and/or additional channels for spatial sound reproduction). Despite DVD-Audio's superior technical specifications, there is debate as to whether the resulting audio enhancements are distinguishable in typical listening environments. DVD-Audio currently forms a niche market, probably due to the very sort of format war with rival standard SACD that DVD-Video avoided.

SECURITY

DVD-Audio discs employ a DRM mechanism, called Content Protection for Prerecorded Media (CPPM), developed by the 4C group (IBM, Intel, Matsushita, and Toshiba).

Although CPPM was supposed to be much harder to crack than DVD-Video's CSS, it too was eventually cracked in 2007 with the release of the *dvdcpxm* tool. The subsequent release of the libdvdcpxm library (which is based on dvdcpxm) allowed for the development of open source DVD-Audio players and ripping software, such as DVD-Audio Explorer. As a result, making 1:1 copies of DVD-Audio discs is now possible with relative ease, much like DVD-Video discs.

IMPROVEMENTS AND SUCCESSION

HD DVD AND BLU-RAY DISC

In 2006, two new formats called HD DVD and Blu-ray Disc were released as the successor to DVD. HD DVD competed successfully with Blu-ray Disc in the format war of 2006–2008. A dual layer HD DVD can store up to 30GB and a dual layer Blu-ray disc can hold up to 50GB.

However, unlike previous format changes, e.g., audio tape to compact disc or VHS videotape to DVD, there is no immediate indication that production of the standard DVD will gradually wind down, as they still dominate, with around 87% of video sales and approximately one billion DVD player sales worldwide. In fact experts claim that the DVD will remain the dominant medium for at least another five years as Blu-ray technology is still in its introductory phase, write and read speeds being poor as well as the fact of necessary hardware being expensive and not readily available.

Consumers initially were also slow to adopt Blu-ray due to the cost. By 2009, 85% of stores were selling Blu-ray Discs. A high-definition television and appropriate connection cables are also required to take advantage of Blu-ray disc. Some analysts suggest that the biggest obstacle to replacing DVD is due to its installed base; a large majority of consumers are satisfied with DVDs. The DVD succeeded because it offered a compelling alternative to VHS. In addition, Blu-ray players are designed to be backward-compatible, allowing older DVDs to be played since the media are physically identical; this differed from the change from vinyl to CD and from tape to DVD, which involved a complete change in physical medium. As of 2011 it is still commonplace for major releases to be issued in “combo pack” format, including both a DVD and a Blu-ray disc (as well as, in many cases, a third disc with an authorized digital copy). Also, some multi-disc sets use Blu-ray for the main feature, but DVDs for supplementary features (examples of this include the *Harry Potter* "Ultimate Edition" collections, the 2009 re-release of the 1967 *The Prisoner* TV series, and a 2007 collection related to *Blade Runner*).

This situation can be best compared to the changeover from 78 rpm shellac recordings to 45 rpm and 33S! rpm vinyl recordings; because the medium used for the earlier format was virtually the same as the latter version (a disc on a turntable, played using a needle), phonographs continued to be built to play obsolete 78s for decades after the format was discontinued. Manufacturers continue to release standard DVD titles as of 2011, and the format remains the preferred one for the release of older television programmes and films, with some programmes such as *Star Trek: The Original Series* needing to be re-scanned to produce a high definition version from the original film recordings (certain special effects were also updated in order to be better received in high-definition viewing). In the case of *Doctor Who*, a series primarily produced on standard definition videotape between 1963 and 1989, BBC Video reportedly intends to continue issuing DVD-format releases of that series until at least November 2013 (since there would be very little increase in visual quality from upconverting the standard definition videotape masters to high definition).

HOLOGRAPHIC VERSATILE DISC

The Holographic Versatile Disc (HVD) is an optical disc technology that may one day hold up to 6 terabytes (TB) of information, although the current maximum is 500GB. It employs a technique known as collinear holography.

5D DVD

The 5D DVD, being developed in the Swinburne University of Technology in Melbourne, Australia, uses a multilaser system to encode and read data on multiple layers. Disc capacities are estimated at up to 10 terabytes, and the technology could be commercially ready within ten years.

USE AS BACKUP MEDIUM

Durability of DVDs is measured by how long the data may be read from the disc, assuming compatible devices exist that can read it: that is, how long the disc can be stored until data is lost. Five factors affect durability: sealing method, reflective layer, organic dye makeup, where it was manufactured, and storage practices.

The longevity of the ability to read from a DVD+R or DVD-R is largely dependent on manufacturing quality, ranging from 2 to 15 years, and is believed to be an unreliable medium for backup unless great care is taken for storage conditions and handling.

According to the Optical Storage Technology Association (OSTA), "manufacturers claim life spans ranging from 30 to 100 years for DVD, DVD-R and DVD+R discs and up to 30 years for DVD-RW, DVD+RW and DVD-RAM".