Introduction

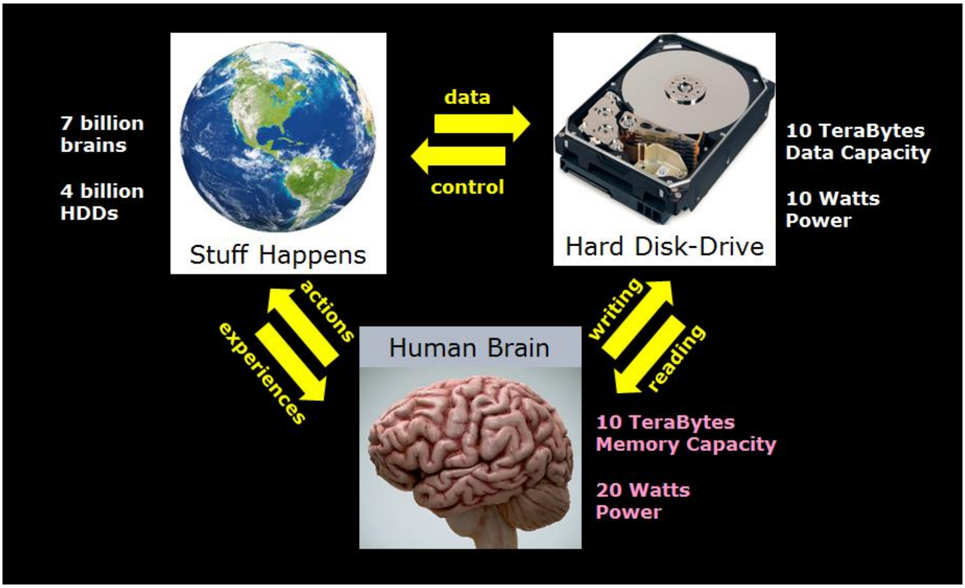

Memory is useful. An organism that can recall earlier situations and their outcomes or earlier actions and their consequences will fare better than one that cannot. Memory may be internal to an organism or it may be external. They say elephants never forget, though it may actually be humans that have the best internal memory of any organism. What cannot be denied is that humans have developed *external memory or data storage* to an absolutely unprecedented degree. Modern society is underpinned by and totally dependent on the institutional memory of all the various organizations and businesses and governments the world over. Some of that institutional memory lies with the individuals involved and some is written down in books and other documents, but, more and more, it is memory in the form of machine-readable digital data manipulated and stored by computers. This data may be stored locally, but, increasingly, data is being stored in the world-wide-web on the 'cloud' in huge concentrated data-centers. Whether saved locally or on the cloud, the bulk of this enormous quantity of data is stored on a ubiquitous but sadly underappreciated device – that 'perfect' invention - the remarkable Hard Disk Drive (almost invariably called an 'HDD').

Memory involves first the creation or formatting or writing of information in a form suitable for storage; second, the maintenance and safe storage of that information, possibly over considerable time-spans; and third, the retrieval or reading of the information. It can be well viewed as the transmission of information through time rather than space and, as we'll see later, has much in common with data communications and information theory. The simplest forms of memory occur very naturally. Any mechanism or pathway that becomes easier to use the more often it is used, can form a basis for memory or data storage. At a basic level, this is roughly how the neurons work that store information internally in the nervous system and brain. But it can also be true for external memory - it is how trails are created automatically by animals going to and from their favorite drinking hole or food source. Trails leading to poisoned waterholes or frequented by voracious predators rapidly disappear from the landscape.

Many animals go one step further and more deliberately modify the environment by laying down pheromone trails (like ants) to remind them of the correct path or by using scent to mark territory (like dogs). Some actually build physical pathways like tunnels (moles) or canals (beavers) to get from A to B. In this manner, memory and information gets stored by being literally carved into the landscape. An advantage of this kind of physical external data storage is that the information is readily shared. Anyone can use a deer trail, not just the deer.

Biological memory – internal data storage

The advantage of external physical storage of information like a deer trail is, of course, also its disadvantage. Competitors and predators can see the very same signs and pathways. In contrast, information stored on neurons inside a brain can only be accessed by the owner of that brain. The grey squirrel common to North America and Europe has a prodigious memory. It hides food (acorns, for example) in hundreds of small caches for later use in the winter. The sites are each identified and remembered based on visual landmarks and a 'world' map and ultimately stored on neurons (nerve cells) within the squirrel's brain. This approach has the advantage that competitors like jays don't know where the food is hidden and cannot steal it. An alternative would be to mark the caches visually, but that has the likelihood that another less thrifty species might pilfer the stores. Jays are notoriously intelligent birds. If the squirrels were to use some conspicuous way of marking each cache, there would quickly be some very fat jays and some very thin squirrels. The downside of internal memory is that the knowledge is available only to the individual. When a squirrel meets an untimely end, those food caches are lost to the entire squirrel community (though perhaps to the benefit of the oak tree community). Obviously the best approach for squirrels would be to note the coordinates of all the caches and record them at a separate secure location readable only by an authorized group of squirrels. Obviously squirrels never got this far, but humans certainly did!

Having a brain and a good memory, however, is not without its drawbacks. Aside from the additional mass and weight and the need for a protective enclosure, a brain needs a constant large supply of energy and nutrients that is very disproportionate to its size. But the most profound and least understood fact is the need for sleep[1]. Sleep-like behavior has been observed in invertebrates as simple as jellyfish. The need for sleep has been documented in many reptiles and fish. Sleep is an absolute necessity for all mammals and all birds. The brain quickly starts to operate incorrectly if it is deprived of sleep. In particular, sleep appears to be essential in creating long-term memories. A prevailing theory is that short-term memory consists of electro-chemical impulses circulating in the brain and the purpose of sleep is to 'fix' these memories into relatively permanent physical changes in the brain cells. The essential nature of sleep is underscored by the fact that the central nervous system is almost completely switched off for extended periods during sleep thus leaving the organism unable to forage for food and highly vulnerable to predation and accidental damage.

The human brain is a remarkable device. It is a compact dense mass of heavily-interconnected nerve-cells (neurons) strategically located as close as possible to the major sense organs: eyes, ears, nose, and tongue. The entire assembly forms a roughly-spherical hemi-hirsute protuberance mounted at the top of the body on a short flexible stem. The human brain is estimated to contain some 86 billion neurons. There are many different kinds of neuron, but they all have the property of being electro-chemically active and capable of processing and transmitting small electrical signals.

---

[1] https://en.wikipedia.org/wiki/Sleep

Their general form is of a heavily-branched filamentary structure. The slender filaments span a few tenths of a millimeter within the brain but can stretch to around one meter for major motor neurons that transmit signals to the limbs. Each neuron is capable of interacting with a very large number of nearby neurons, perhaps as many as ten thousand. The interfaces between neurons are called synapses. These synapses are where long-term memories are believed to be stored. Pairs of neurons that are frequently simultaneously excited gradually develop a much stronger connection across their common synapse. A strong connection means that when one neuron is stimulated, the other is likely to follow.

The storage capacity of the human brain is the subject of much debate. There are at least 100 trillion synapses and the strength of the connection across the synapse is an analog (not digital) quantity (so possibly several bits[2] worth of information per synapse). However there is probably much duplication and redundancy and inefficiency. Conservatively, the storage capacity has been estimated to be of the order of ten TeraBytes[3]. The access time to human memories is similarly difficult to estimate. The speed of propagation of electrical impulses (action potentials) covers a very wide range (0.1 to 100 meters/second) depending on the diameter and environment and type of the neuron. Taking 10 m/s as typical and 10 cm (~4 inches) as a characteristic dimension of the brain, gives a transit time of 10 milliseconds. The brain consumes a disproportionately large amount of energy in comparison to the rest of the human body and dissipates a total of about 20 Watts of power. Curiously, as we will see later in this book, these three numbers, 10 TB storage-capacity, 10 milliseconds access-time, and 20 Watts power-requirement are very similar to the numbers for the modern hard disk drive.

Analog vs. Digital

In the world of electronics and information storage, there is a distinction between systems that operate on a continuum (analog) and those that

---

[2] A bit is a unit of information. It refers to the amount of information contained in defining or choosing one of two possible states. If someone tells you the result of a coin toss is, say, heads, they have given you one bit of information. If there are more than two possibilities or states involved then there may be several bits of information.

[3] Ten TeraBytes (10 TB) = 10 trillion Bytes or 80 trillion bits (8×10^{13} bits). Data Capacity is typically quoted in Bytes (with upper-case B) rather than bits (lower-case b). One Byte = 8 bits.

operate with discrete values (digital). The brain is essentially an analog machine with the synapses storing arbitrary values corresponding to the strength of the connection across each synapse. The advantage in working in the analog domain is that the values and processing can have a direct correspondence to real-world quantities. The disadvantage is that there are limitations on how accurately these values can be initially set then finally read back. Particularly during the storage phase the values may drift or change for myriad reasons.

Operating in the digital domain means working with just a small set of distinct known values rather than arbitrary values on a continuum.. (Digital does not necessarily suggest ten values and often just two values are used – binary ‘bits’). Since the exact values are prescribed and known in advance, small changes in the stored quantities can be recognized and tolerated and the quantities frequently reset to the nearest acceptable value as required[4]. The big disadvantage of digital is the need to translate the initial analog values into a digital form with sufficient accuracy and resolution (analog to digital conversion or ADC) and finally the converse (digital to analog conversion or DAC). A typical situation might take a single analog value and translate it into an 8-bit or 16-bit digital representation (corresponding to 256 or 65536 levels). Simple operations like the linear addition of two signals now become highly nonlinear functions of a large number of binary quantities. Nevertheless the ability to withstand noise and distortion and use to really tiny devices for each bit has resulted in digital techniques supplanting analog universally in computers and more recently in almost all consumer devices.

Some devices like semiconductor Static-RAM (Random Access Memory) which is built of an array of bi-stable transistor latches are inherently binary devices storing ‘on’ or ‘off’. Similarly Magnetic-RAM has only an up and down state. It is physically impossible to store any intermediate value. Other devices like Flash-Memory are intrinsically analog devices but are deliberately arranged to store two, four, eight, or sixteen discrete

[4] Noise and distortion rapidly build up in making successive copies of an analog audio tape recording. However digital recordings can generally be copied through many generations without any degradation.

levels per cell and are thus used in a digital mode to store one, two, three, or four bits[5] per cell.

However, it should be noted that digital storage is not confined solely to the world of computers and consumer devices, it is actually a universal and essential feature of all biological systems. There is one unique digital storage mechanism used by all living creatures. Some people would say that it defines life itself. Twenty-three chromosomes contain pretty much a complete set of instructions on how to build a human being, starting almost from scratch. The instructions provide not only the details of how to make each organ and where to position it and what to connect it to, but, more importantly, the exact sequence of steps and the exact timing required to grow and assemble an entire new person from very basic raw materials. The information is strictly digital and is recorded on strands of the famous DNA[6] double helix. The information is recorded in a quaternary format (effectively two bits at a time) by choosing one of four possible chemical bases: cytosine, guanine, adenine, or thymine. The bases are about 1/3 nm apart, so the linear recording density is 6 bits/nm or 6 Gbits/m or 150 Mbits/inch (about 75 times higher than data on a hard-disk drive). However, the diameter of the double-helix DNA strand is just a few nm, so DNA wins hands-down not only in linear-density but also in volumetric efficiency by many many orders of magnitude over a hard-disk drive. It is speculated that DNA might one day provide the basis for a very high capacity very high-density machine-readable data store, but currently, despite huge advances, there are severe limitations on the rate at which data can written and retrieved that preclude any real application.

Nevertheless the successful storage and transmission of this information from generation to generation without too many mistakes and without getting mixed up with other DNA is exactly what defines a living species. There are occasional transcription errors, but these errors together with the forces of natural selection are exactly the reason that such complex organisms ultimately evolved. That this remarkable DNA storage

---

[5] For N equally-likely possibilities, the amount of information stored is $\log_2(N)$ bits. For DNA, N=4, so the information per base is $\log_2(4) = 2$ bits. Claude Shannon developed these concepts and the field of information theory in 1948 while working at Bell Labs.

[6] DNA or DeoxyriboNucleic Acid is the long chain molecule that carries the genetic information on all living organisms on Earth. There are about 2 meters of DNA in the chromosomes in a human cell, assuming it could all be stretched out into a straight line.

mechanism exists at all and that we who are a result of this mechanism can observe and comment on it is perhaps the ultimate miracle. But now we should move from the philosophical and get back to more practical topics and talk about data storage and, ultimately, the remarkable hard disk drive.

Solid-state

Nowadays a key desirable attribute of memory or data-storage is that it be ‘solid-state’. Solid-state originally advertised a device that used semiconductors as opposed to one that used vacuum tubes. However, in this context, it simply means that it’s not a disk-drive or tape-recorder or an optical disk (CD, DVD, Blu-ray, etc.). Solid-state now basically means just that there are no moving parts in the storage-device (beyond the movement of individual ions or electrons or spins[7]). An old-fashioned electric light switch is about the simplest example of a binary memory element. It can be toggled between two mechanical states. It is obviously solid, but it does not count as a solid-state memory because of the moving parts. Latching relays are similar devices that actually allow electrical control. However devices like this (including HDDs) that involve macroscopic mechanical motions generally do not lend themselves to rapid access or high data rates. Another concern is that sliding surfaces or flexing structures can exhibit wear and fatigue.

The key advantage of solid-state is that access is achieved by electronic selection via lithographically pre-defined paths. This means that the access can be very fast, often in mere nanoseconds[8]. In contrast, in hard-drives and tape-drives, access is primarily mechanical with the medium moving at a steady velocity along tracks in one dimension and with the read/write head moving cross-track to address different tracks. (Often there are multiple read/write heads, but switching between heads can be done electrically and thus is relatively fast). Access-times for hard drives are necessarily measured in milliseconds (slow, but still much faster than the proverbial blink-of-an-eye). Tape-recorders are even slower with access times measured in seconds. But the down-side of solid-state is that

---

[7] Electron spin is responsible for magnetization in ferromagnetic materials like cobalt, iron, and nickel and thus provides the storage mechanism in disk drives and tape.

[8] The prefix ‘nano’ means ‘one billionth of’ (see table at end of chapter). To try to get this on human scale: light can travel about one foot (30 cm) in one nanosecond - illustrating both how very short a nanosecond is and how very fast the speed of light is.

every single tiny storage cell must be physically defined and uniquely addressable. This places a huge burden on the necessary lithographic tools and processes. In contrast, disk and tape media have uniform unfeatured surfaces and are inherently low-cost. Media production can literally be measured in acres or hectares. These big differences between solid-state versus HDD or tape technology get reflected in cost. Solid-state cost per bit is roughly an order-of-magnitude higher than for HDD and tape. This trade-off between performance and cost leads to a hierarchy in storage types. Expensive rapid-access memory is used close to the computer-processor unit (CPU) and less-expensive slower-access devices are used for data that is accessed less-frequently or less-urgently. The terms 'memory' and 'data storage' have similar meanings, though 'memory' suggests rapid-access devices close to the CPU, while "data storage" tends to suggest disk drives and tape drives. Data can be shuffled up or down the storage hierarchy by the system as the usage requirements change. There will be more about this trade-off and about the memory or storage hierarchy in later chapters.

As mentioned earlier, mechanical motion can lead to wear and fatigue. Non-solid-state devices with moving parts like HDDs are very carefully designed taking such issues into account in order to achieve the required high levels of reliability. Even solid state devices, however, are not immune to such wear and aging effects. In particular, if the storage mechanism involves the movement of atoms or ions[9] then the process of switching between states may not be perfectly reversible and there will be a limit on how many times the memory can be cycled or used. In phase-change memory, for example, the atoms switch between an amorphous state and a crystalline state. Such devices come with relatively strict limits on the number of times a cell can be re-written. Some types of memory are strictly Read-Only-Memories (ROM) with the information fixed in the factory (CDs and phonograph records, for example). Other types of memory are deliberately designed to be written by the customer but only once though again can be read many times (examples include CD-R and

[9] An ion is an atom that has lost or gained an electron. An ion has a net electric charge because the number of electrons no longer balances the number of protons in the nucleus. Ions and electrons can be readily manipulated by electric and magnetic fields.

DVD-R[10]). Others such as flash memory are more forgiving and can be re-written many thousands of times. Flash memory theoretically involves only the movement of electrons but nevertheless still exhibits significant wear and this is a major limiting factor for the technology. Later we will be talking a good deal more about "NAND-flash" and the corresponding Solid-State Drives (SSD). These devices currently account for only a small fraction of data stored but they are becoming increasingly important because of their attributes of fast access, low idle power, and physical robustness.

There is one storage mechanism that does not involve the movement even of electrons. The reversal of electron spins or magnetization is about as gentle and innocuous as one can get for a storage mechanism. It is viewed as indefinitely reversible. This means there is no wear and no limit on the number of times the memory can be written and rewritten. Magnetic recording exemplifies this fact. The spins in the magnetic recording medium can be switched back and forth any number of times without any degradation in their characteristics.

Volatile vs. Non-volatile

This is another important distinction in memory or data-storage terminology. Human memory is volatile. If the flow of energy to a person's brain is interrupted, within a few minutes their memories quickly disappear – sadly forever. On the other hand, their epitaph carved into a gravestone would be definitely classified as non-volatile. That message could easily last for centuries with no power supplied at all! Similarly, the contents of high-speed static-RAM semiconductor memory in a CPU will disappear almost immediately if the computer power is lost. However, the data on the magnetic hard-drives and tapes is non-volatile and can last for decades - standing ready to 'reboot' the computer when power is restored. Magnetic-RAM (MRAM) is similarly non-volatile and is a serious contender for a place in the storage hierarchy to replace volatile memory. In this case, the data is stored in a tiny permanent magnet that can point either up or down (inherently binary) and the magnetic state and the data persist even if there is no power supplied. Resistive-RAM (ReRAM) operates along lines somewhat analogous to a synapse on a neuron. The

[10] The "-R" stands for "recordable". The recording mechanism is a high-contrast color change that occurs in a thin-layer of organic polymer dye when exposed to laser light

ease with which current flows through of an individual storage element depends on how much current has previously passed through that element. In contrast to biological neurons, ReRAM is a non-volatile memory.

The distinction between volatile and non-volatile can be a little fuzzy. In dynamic-RAM (DRAM) which plays an important role in the computer memory hierarchy, the data persists on time-scales of seconds. It thus needs to be refreshed every few milliseconds to ensure reliable storage. DRAM is normally viewed as a volatile memory but can be used as non-volatile memory with suitable battery backup.

(Don't despair at trying to track this alphabet soup of "-RAMs". There will be more explanation and discussion in later chapters).

Evolution of Data Storage

The storage of information by humans has had a long history that has seen continued growth over several millennia starting with simple tallies on rocks or bone, through more sophisticated writing and accounting systems on parchment or paper, through to today's mass reproduction and distribution of books and magazines and newspapers. What must also be included is all the information (music and movies) distributed on CD's and DVD's. Until fairly recently, this type of content was distributed in analog form, on video VHS tapes and audio cassette tapes and earlier on phonograph records and photographic movies. But this scenario is shifting again rapidly as optical fibers and wireless transmission and cellular networks replace copper wires. The result is that inexpensive high-speed data-transmission and connectivity are becoming almost universal. Information can thus be stored more centrally in a computer-based system and distributed electronically on-demand directly to the customer's 'device'. A lot of the information or data is still based on the written word such as in news stories or emails or texts, but nowadays the bulk of the data comprises compressed[11] audio or pictures or video (think social media like Facebook and WhatsApp - both with over a billion active users).

---

[11] Data-compression can be applied very effectively to audio and still-images and video because the waveforms involved are often very repetitive. Compression algorithms take advantage of this redundancy to dramatically reduce file-size. Examples include the .mp3 and .jpg and .mpg file-suffixes for audio and images and movies, respectively.

The use of machine-readable data has grown exponentially over the last 50 years or so with the advent of sophisticated computer control and data transmission and storage systems. The recent growth has been so explosive that every few years one is forced to learn a new Greek prefix to describe the total amount of data stored in the world or even in a single data-center (mega- ➔giga-➔ tera- ➔ peta- ➔ exa- ➔ zetta- where each step is a factor of 1000). For example, it is estimated that the total HDD storage capacity shipped in 2015 alone exceeded half a zettaByte. The prefix zetta- means 10^{21} or 1 followed by 21 zeros[12]. Each Byte is defined as 8 bits, so half a zettabyte means about four thousand billion billion bits. As a reference point, it has been estimated that the human brain can hold of the order of 10 TeraBytes of data. How much of this capacity is actually used is a matter for debate and even jokes. But the world is certainly changing. The rote memorization that characterized the world's formal education systems throughout the 20th century is much less emphasized nowadays. The ease with which information can be accessed on hand-held devices in real-time during daily activities or conversations means there is much less need or incentive to commit the information to memory. A very good example is navigation in automobiles. It is no longer necessary to remember the route to drive from A to B one simply punches or dictates it into one's smart-phone or car navigator. "The Knowledge", as the London taxi-drivers' infamously difficult examination is called, requires the cab-driver know all the details of the network of 25,000 streets in central London including all the hotels, restaurants, clubs, theatres, hospitals, police stations, embassies, bus-stops, stations, places of worship, cemeteries, parks, sports-centers, schools, colleges, etc. There is evidence that mastering the "Knowledge" produces measurable changes in the brain, but this may all be moot – "The Knowledge" is viewed by many as an anachronism.

It can be argued that the ease and speed with which information on almost any topic can be accessed from a simple hand-held device will cause the storage capabilities of the human brain to atrophy from lack of exercise. This is all very interesting to speculate about, but this book takes the very positive and hopefully correct perspective that this almost universally-accessible massive store of information written primarily on HDDs is a

[12] This scientific notation and the corresponding Greek prefixes are shown in the table at the end of the chapter.

great boon to civilization allowing better and faster and more universal decision-making to the enormous benefit of mankind.

Background and Structure of the Book

This book has an interesting background and a perhaps curious structure. The author is veteran of many years in the hard-disk drive industry and has witnessed first-hand the amazing progress in the technology and can speak directly and enthusiastically to all the superlatives that the technology embraces. However, the book also reflects the author's interests in other fields ranging from astronomy and spaceflight to nanotechnology and micromagnetics. In particular, the book attempts to include some historical context and some of the key personalities involved. Above all the book tries to create an easily readable narrative but at the same time includes numerous interesting diversions and adds many simple intuitive explanations of important terms and physical phenomena. This allows the reader to treat the book either a popular science text or, with a little more effort, as an introductory textbook on the basic principles behind magnetic recording and the hard disk drive in particular. Although the authors has not been directly involved in solid state memory, the technology and attributes of NAND-Flash are introduced in some detail so that the reader can better appreciate the changing data storage landscape.

The unusual structure of the book follows from a slide presentation that one of the authors was invited to give at a banquet at an ASME[13] conference in 2016. The topic was information storage processing systems. However, this was not to be a formal technical talk, but something that would be perhaps more entertaining and exciting and would hold the audience's attention. It is unclear exactly what prompted the approach used here where the information is presented in descending factors of ten in scale. But, often, as engineers designing or testing components, we do fail to truly appreciate the tiny scales at which the recording processes actually take place. Conversely, the internet and the cloud operate on a truly global scale and, beyond that, magnetic storage has been an essential feature of space exploration. So an objective of the original slide presentation was to bring a realization of the actual physical scale on which devices operate and furthermore a better appreciation of

[13] American Society of Mechanical Engineers

the huge *range* of scales involved. One page in the original presentation was devoted to each factor of ten in scale. The starting point was 10^{+13} meters and the ending point was 10^{-10} meters.

So the book is structured in the same way, in decreasing factors of ten in size. In some ways it is a very artificial way of proceeding with the story and at some scales one struggles a little to find interesting relevant discussion points. But nevertheless that rigor is enforced here, so there is exactly one chapter dedicated to each length scale. The chapters are of very variable length and contain varying amounts of information that is pertinent to data storage, but altogether it makes for an interesting journey with plenty of diversions, while still maintaining that common thread of data storage.

We had to decide the starting and ending points of our logarithmic[14] journey. The starting point for the journey is 10^{+13} meters which takes us way outside the solar system and allows is to talk about the Voyager spacecraft and space exploration and all the technologies involved – including the tape recorder, of course. If we were to go to even larger scales, there would not be much to talk about until we get to our nearest star, Proxima Centauri, and the exoplanet Proxima-b discovered in 2016. This system is about $4 \times 10^{+16}$ meters distant (4.2 light-years). Proxima-b is not only the closest known exoplanet but it also happens to lie in the star's so-called, Goldilocks zone, where liquid water and hence life could possibly exist. However, unless we suddenly get a message from the inhabitants of Proxima-b detailing their HDD technology, there's not likely to be anything much to say about data storage at this scale.

Similarly at the other end of the scale, 10^{-10} meters is one Ångstrom[15] which is a useful measure for describing atomic structures and the diameter of atoms and interatomic spacings. The properties of magnetic materials are determined at this length-scale. Below the Ångstrom length-scale, the next thing to talk about would be the nucleus of an atom,

[14] Logarithm refers to the number of zeros in the number. Examples: $\log_{10}(10000) = 4$ and $\log_{10}(1) = 0$. The concept generalizes so that $\log_{10}(2) = 0.3$ and $\log_{10}(0.1) = -1$.

[15] The Ångstrom is not an officially recognized SI unit, but it is conveniently about the size of an atom so it has become widely used in solid-state physics and chemistry. Anders Jonas Ångström was a Swedish physicist known for his pioneering work in spectroscopy in the 1800s.

however, atomic nuclei are measured in femtometers and a typical atomic nucleus is roughly 10^{-14} meters in diameter. Perhaps one could talk about Magnetic Resonance Imaging (MRI) which operates at the level of a nucleus but again it is difficult to link this to data storage other than saying that MRI images can consume quite a lot of space on a disk drive. So again there is a natural gap between 10^{-10} and 10^{-14} meters. So we stop the journey at 10^{-10} meters or one Ångstrom.

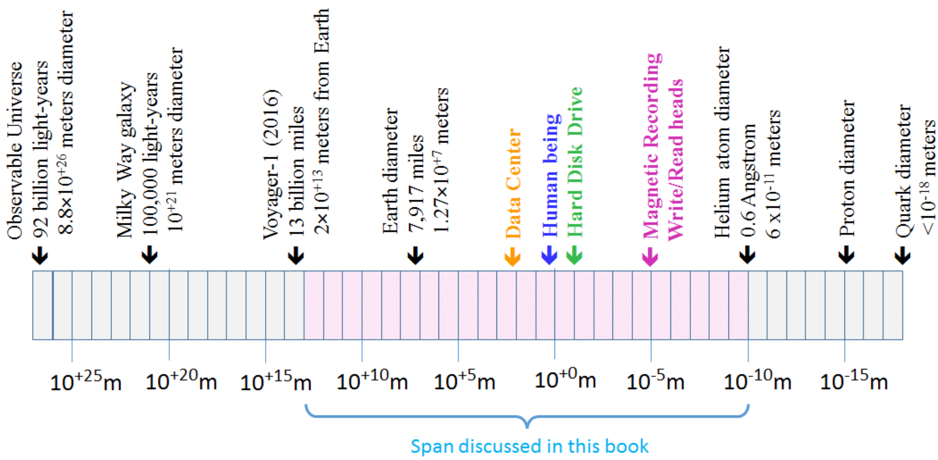

Undoubtedly, an inspiration for the structure of this book was the 1977 short documentary film "Powers of Ten" by Charles and Ray Eames. Forty years later it is still well worth viewing[16]. This short film was based on the earlier 1957 graphic essay by Kees Boeke. Charles and Ray Eames were already well known for their work in architecture and furniture. The couple turned their hand to film as a medium and created some 125 short films often popularizing science topics. The 1977 version of "Powers of Ten" film "Powers of Ten" covers length-scales in some forty steps from a billion light-years, 10^{+25} meters, down to the scale of quarks, 10^{-16} meters. The film was narrated by Philip Morrison. He was a Professor of Physics at MIT and a participant in the Manhattan project though later became an activist for nuclear non-proliferation. He was also a popular science writer who with wife, Phyllis, created the book[17] version of the film.

---

[16] Eames Office, "Powers of Ten": https://www.youtube.com/watch?v=0fKBhvDjuy0

[17] P. Morrison et al., "Powers of Ten", Scientific American Library; Revised. (Aug. 1994).

The most recent version of the powers of ten approach comes from Caleb Scharf claiming to span from the size of the observable universe, 10^{+27} meters, down to the Planck length of 10^{-35} meters[18]. Even more ambitiously, but covering a journey in *time* rather than space, is a book from Nobel Laureate Gerard 't Hooft and physicist Stefan Vandoren that covers seventy-five orders of magnitude from 10^{+32} seconds down to 10^{-43} seconds[19].

From Interstellar space to interatomic distances

Our journey in this book starts with the Voyager-1 spacecraft, the most distant human-made object, and then works its way inwards through the solar system to near earth orbit and then to systems on the Earth. In the early chapters, we spend much more time on discussing the spacecraft, the various missions, the instrumentation, the space environment, and the new discoveries, rather than on the on-board data-storage. But, at the same time, some basic concepts of magnetic fields and flux and the essence of magnetic recording are introduced. Once we are down on Earth, the discussion centers on the internet and cloud storage and data-centers, and also the back-ground and history and personalities involved in the development of magnetic recording and hard disk drives. Many of the important innovations in both tape and hard drives occurred in the San Jose area in California so this fills in the kilometer-scale chapters. These developments were certainly not exclusive to San Jose, but the United States and Japan together have always dominated the industry. The current (2016) leader in Flash or Solid-State Memory is Samsung, a Korean company although the other major players are again based in the US and Japan.

When we finally get to the 10 cm scale, the size of a modern HDD, it is easy to start talking about all the superb engineering that culminates in the ability to store in excess of ten TeraBytes in a standard form-factor 3.5" (95 mm) hard disk drive. Zooming in to much smaller scales, we are able to focus on the magnetic components within the drive and on the details of the writing and reading processes. One of the miraculous aspects in HDD technology is the achievement of the non-contact head-disk interface which reliably maintains the head at sub-nanometer clearance over a disk

[18] C. Scharf, "The Zoomable Universe", Scientific American, New York, 2017
[19] G. 't Hooft et al., "Time in Powers of Ten", World Scientific Publishing Co. (May, 2014)

surface rotating at about 70 miles per hour (30 meters/second). The fanciful analogy is often made with a 747 airliner flying less than a millimeter above the ground. At the end of the journey, we are down at atomic scales (Ångstrom scales) and have the opportunity to discuss the atomic structure of important materials and the effects that occur at these scales – including the effect called magnetism! Since magnetism plays such a central role in this story we introduce some of the history and concepts in the inset boxes at the end of this chapter.

A final chapter summarizes the journey through space and time and also talks to the uncertain future of the remarkable Hard Disk Drive and to current efforts to extend the technology and satisfy the world's apparently insatiable appetite for data storage. In some ways the replacement of Hard Disk Drives by Solid State Drives seems inevitable, but it is very clear that this will not happen quickly.

Table: Metric prefixes in everyday use			
Text	Symbol	Factor	Power
zetta	Z	1000000000000000000000	10^{+21}
exa	E	1000000000000000000	10^{+18}
peta	P	1000000000000000	10^{+15}
tera	T	1000000000000	10^{+12}
giga	G	1000000000	10^{+9}
mega	M	1000000	10^{+6}
kilo	k	1000	10^{+3}
hecto	h	100	10^{+2}
deca	da	10	10^{+1}
(none)	(none)	1	10^{0}
deci	d	0.1	10^{-1}
centi	c	0.01	10^{-2}
milli	m	0.001	10^{-3}
micro	µ	0.000001	10^{-6}
nano	n	0.000000001	10^{-9}
pico	p	0.000000000001	10^{-12}
femto	f	0.000000000000001	10^{-15}
atto	a	0.000000000000000001	10^{-18}