

Memory and Storage

Russell Hitchcock

In today's typical networks there are many different forms of data storage. Understanding these methods of data storage is essential to understanding your network. In this series of articles I will explain some of the more common forms of data storage, some exciting, up-and-coming technologies, and I will outline some of the more common memory addressing protocols.

Magnetic Storage

Hard drives, the most common form of magnetic storage, store data on disk shaped platters. These platters are typically made of aluminum, glass, or ceramics, and are coated with a ferromagnetic material which is often a cobalt alloy. This ferromagnetic coating allows read/write heads to magnetize small regions of the platter which represent a digital bit.



Figure 1: Hard Drive

Platters, yes there is more than one on a hard drive, are separated by spacers on a single spindle. This spindle is controlled by a motor which can spin the platters. This motor's speed is constant, and is the speed advertised as the speed of the hard drive.

Read/write heads, one per platter side, are attached to a single actuator arm. The actuator arm is controlled with a servo motor which can move the heads closer to, or further away from the spindle in unison.

There are two ways to write the data onto the platter: longitudinal and perpendicular. Longitudinal is the traditional way to write the data onto the platter. You can think of the bits like a bar magnet laying flat on the platter surface, end to end. You can easily imagine that these regions can take up quite a lot of space. This is why hard drive manufacturers have been working to reduce the size of these regions. Currently we are reaching the limits of how small engineers can make these regions. This is because of the superparamagnetic effect. Basically, this effect describes how random thermal effects can flip the polarity of a very small magnet. On a hard drive, if the polarity of one of the magnetized regions is flipped then this would mean that the data is changed from a '1' to a '0'. This is bad.

Memory and Storage

Russell Hitchcock

Perpendicular recording technology can allow engineers to pack more data on the same area of the platter without having to worry about the superparamagnetic effect for similarly sized regions. If you picture the bar magnets again, it should be obvious why you can fit more of them in the same area when they are stood on their ends, i.e. perpendicular. It is a little less obvious why they do not have to worry about the superparamagnetic effect. Essentially this is because the direction of the magnetic fields has changed, and therefore they react with their neighbours differently. This reaction is important in determining whether the superparamagnetic effect is in play.

Similarly to hard drives, tape drives store bits by polarizing a small magnetic region. There are basically two types of tape drives: linear and helical.

Linear tape drives have linear tracks. On the tape there are several dozen tracks which extend from one end of the tape to the other. Each track consists of many magnetic small regions which can be used to represent a '1' or a '0'.

Helical tapes have tracks that run diagonally up and down the tape. This means that the tracks will overlap each other. Normally this would be bad; however this type of tape drive uses two write heads, each using an opposite polarization which allows the read heads to distinguish between the tracks. This allows for a higher capacity on the tape.

Semiconductor Memory

One of the most common types of semiconductor memory is RAM, shown below in figure 2. There are two general types of RAM, static and dynamic. Static RAM or SRAM stores data in a collection of 6 transistors, commonly known as a flip-flop. Dynamic RAM or DRAM stores data in capacitors which require continuous refreshing, and is the reason why DRAM loses the data when power is shut off. The advantage of DRAM is that it only takes one transistor and one capacitor for each bit. This gives a very high memory capacity compared to a similarly sized SRAM chip. The advantage of SRAM is that the transistors do not require refreshing and react faster than capacitors.



Figure 2: Ram

Another form of semiconductor memory growing in popularity is Flash memory. There are two basic types of flash memory, NOR and NAND. NOR (Not OR) refers to the NOR logic gate while NAND (Not AND) refers to the NAND logic gate.

Both NAND and NOR logic gates consist of a collection of transistors, and do not contain capacitors. This means that they do not require refreshing and therefore hold data when power is no longer applied.

Memory and Storage

Russell Hitchcock



Figure 3: A USB Flash Drive

Although both NAND and NOR Flash are similar in some respects they are also quite different. NAND flash being a sequential access technology is better suited to be used to store data. NOR flash is a random access technology which makes it better suited to store programs which use little memory. NOR flash is usually used in applications such as running a cell phone's operating system. NAND flash is typically used in applications such as USB memory sticks. According to www.appleinsider.com, Apple's new iPhone uses both NOR and NAND flash.

Optical Memory

The most common type of optical storage is the CD. CDs are made from injection molded polycarbonate plastic which has microscopic bumps, sometimes called pits, arranged in a continuous spiral around the disk. It is these bumps that represent the data. Over this polycarbonate is a thin layer of reflective material, usually aluminum or gold and over that is an acrylic layer to protect the disk.

When a CD is being read, a laser is shown through the polycarbonate layer and reflected off of the reflective material. The reflected laser light is in turn detected by an optical sensor which converts the received laser signal into electricity. Depending on whether the laser was focused on a bump or not, the electrical signal will be different because the reflected laser light will be different. The difference in the electrical signals is how a computer can read data off of the CD. That is the case for regular CDs, but what about recording data onto CD-Rs, and CD-RWs?

A CD-R is similar to a CD in their construction except for two key aspects. First, there are no bumps. Second, between the polycarbonate and the reflective aluminum there is a layer of transparent dye. To save data onto a CD-R the writing laser is focused onto the desired part of the spiral (which doesn't really exist until you create it by writing data) and heats up the dye. The chemical properties of the dye are such that when it is heated to sufficiently high temperatures its level of opacity changes. So the writing laser can move along the spiral and change the opacity of small regions, this difference in opacity is how you can create a '1' or a '0'. Data is then read from the CD-R in the same manner as a CD. CD-Rs of course can only be written to once. This is because once you make the dye opaque you cannot then make it transparent. So then what about CD-RWs? CD-RWs use a different dye which starts out opaque and when heated turns transparent. This dye also has the amazing property of turning back opaque if heated to an even higher temperature. This allows you to erase the data previously written to the disk.

Memory and Storage

Russell Hitchcock




Figure 4: An Image of the Bumps on a DVD

DVDs work exactly the same as CDs. DVDs can store more data because there are essentially many thin CDs stacked on top of each other. That is, they are made of several layers of polycarbonate and reflective material. The lasers and optical sensors are also more advanced, in that the laser has the ability to pass through the different layers and the optical sensor can detect all of these different layers. Those are some of the more common methods of data storage available. Keep an eye out for my next article where I will discuss some newer, more advanced technologies such as pulse-change and holographic memory.

Memory Technologies

In my last article I explained to you how some of the more common memory storage technologies work. All of these technologies have found their place in today's typical network. In this article I will discuss technologies which have not yet found wide acceptance. In some cases these technologies are still only found in laboratories, in other cases the products available today do not fully reflect the potential of that technology.

Molecular Memory

So what is wrong with the memory storage technologies explained in my last article? Well, nothing. The motivation for developing new memory storage technologies is that we are quickly reaching the limits of how small and how fast we can make these things while users are demanding more capacity and better performance. New technologies will soon be needed. Is that technology molecular memory? Maybe. What makes molecular memory attractive is that even large molecules are very small and could provide a memory density many times greater than current technologies.

To hold a bit in a molecule is, in theory, quite simple. You simply add or subtract electrons to the molecule. The difficult part is reading and writing the bits. You cannot exactly have a copper wire leading to a molecule; the size difference would be inhibiting.

Memory and Storage

Russell Hitchcock

In order to access the molecules for reading and writing, some researchers have been arranging arrays of molecules around tiny nanotubes capable of carrying electric charge. This method is shown below in figure 1. Other researchers are trying to manipulate bits via radio waves. Did you say radio waves? Absolutely. Pretty cool eh! They do this by creating an electromagnetic pulse of a specific frequency which would alter the charge of the molecule. To read the bits another pulse of a different frequency is then created. The effect the molecule has on this second pulse can tell you what the first pulse did to the molecule, hence allowing you to store and then read a bit.




Figure 1: A Diagram of a Molecular Memory Device

As you can see, molecular memory holds great promise to provide the user with a large memory density. Currently, molecular memory is still in the laboratory stage so we will all have to wait a few years to find out what this technology can really do for us.

Phase Change Memory

Unlike molecular memory, phase change memory is available today. In fact, the technology behind phase change memory has been around for decades. In the 1960s Stanford Ovshinsky invented a way to crystallize amorphous materials, that is, materials without a crystalline structure.

As mentioned in my last article, CD-Rs, and CD-RWs work by a laser changing the opacity of a small region on a disc. What changes the opacity is the fact that the material changes from amorphous to crystalline, or vice versa. This is the same technology invented by Ovshinsky. Ovshinsky actually made a prototype CD-RW in 1970!

The difference between CDR technology and phase change technology is that with phase change memory the crystalline state of a small region is altered with an electrical current not lasers. Since we are not using lasers to read and write the data we do not refer to the opacity of the region but rather the resistivity of the region. Once the region is changed to either crystalline or amorphous, the resistivity of the region can be measured and based on the amount of resistivity the region is considered to be either a '1' or a '0'.

As a side note, you should now be able to see that electrical resistivity is really quite similar to opacity. Where a resistive material does not allow much electricity to flow through it and an opaque material does not allow much light to pass through it. You should also know that opaque materials actually reflect light. You may not realize that resistive materials also reflect. More correctly, it is the impedance of a material which reflects the electricity. Resistance is one aspect of what makes up

Memory and Storage

Russell Hitchcock

impedance; the others are capacitance and inductance. In many applications, limiting this reflection by impedance matching is a major design consideration.

Phase-change memory has the potential to replace flash memory in only a few years. How does it compare to flash? Like flash, phase change memory is a non-volatile random access memory making it well suited for both code execution and data storage. In 2006 IBM, along with Macronix and Qimonda, announced research results which stated that they had designed, built, and demonstrated a prototype phase-change memory device. This device was 500 times faster than flash while using less than half of the power. The prototype device was also far smaller than flash memory.

Another sign the phase change memory will soon be upon us in great numbers comes from Intel. In April of this year Intel announced that they would begin shipping samples of phase change modules in 128mb sizes. Look for these modules to appear in your electronics soon.

Holographic Memory

Many people think holographic technology is a futuristic, far off technology, but it is available today from InPhase technologies. Of course, it is not widely available and is quite expensive. This will soon change because there are a lot of advantages to storing your data on holographic memory.



Figure 2: A Holographic Memory Device

Holographic memory works by shining two coherent beams of light onto a light sensitive medium, a data beam and a reference beam. The three dimensional interference pattern created by these two beams of light is stored as a hologram. This interference pattern can be read by shining only the reference beam of light onto the interference pattern; the resulting beam resembles the original data beam.

This type of three dimensional memory means that we can store and access memory pages at a time. It also means that holographic memory devices will have an incredibly dense memory capacity.

Memory and Storage

Russell Hitchcock

Because of these advantages, I feel that it is inevitable that holographic memory will become a powerhouse in the tertiary storage arena. I am not convinced however, that holographic devices will ever reach the popularity of CDs and DVDs.

Magneto-resistive RAM

Magneto-resistive Random Access Memory, like holographic memory, is available today. In July 2006 Freescale semiconductor announced the world's first mass produced MRAM product. Even today, one year later, there are few MRAM products available. Those that are available are expensive, are of low memory density, and are of interest to only a small niche market. Within a few years I believe that MRAM will become much more popular.

Like hard drives, MRAM stores data in magnetic storage media, making MRAM a non-volatile storage medium. This is an important feature of MRAM which will allow it to compete with DRAM and SRAM which lose their data when the power is no longer applied. Although not as fast as SRAM, MRAM chips provide faster read and write times than DRAM. MRAM does however, have a much higher memory density than SRAM. This will give engineers designing future CPUs another option as they will have to choose between a large, slightly slower cache, and a smaller, slightly faster cache. Personally, I think MRAM will win in the end, as it takes much longer to put data into the cache from your hard drive than to read data from the cache. So the less often you have to load the cache the better performance you will see.

Watch for MRAM to take on flash memory within a few short years. This is because MRAM memory, while much faster than flash, is also cheaper to make. I feel that the cost advantage of MRAM will be what propels the MRAM chips to take on the flash market, while the increased speed will be a nice side benefit to the customer.

In my next article I will set away from talking about the memory itself and discuss memory addressing techniques. These techniques are extremely important for both speed and error correction. I know you will enjoy it.

Bus Specifications

In my last two articles I wrote about various memory technologies including; hard disk drives, RAM, MRAM, and holographic data storage. Of course, the technology used for the physical data storage is only part of what will determine the performance you see. To fully understand the performance of your memory, you must understand how the memory connects to the CPU and how it can be manipulated. Generally, memory connects to the CPU the same way that any peripheral connects to the CPU; via a system bus. The next time you have your computer apart, take a look at the underside of your motherboard. Here you will see a lot of wires; this is the system bus.

Different system buses are defined with system specifications. A system bus specification will include things like; the numbers of wires, the size of the wires, what each wire should be used for, and what connectors should be used. Modern bus specifications have advanced features and are quite complex, with multiple wires for data transmission in both directions, and multiple wires for control signals. The complexities of these system buses are really blurring the lines between buses and networks.

One of the oldest and continuously used bus specifications is the inter integrated circuit or I2C, specification. With only two bidirectional wires, one for data and one for a clock signal, this is a relatively simple specification.

Memory and Storage

Russell Hitchcock

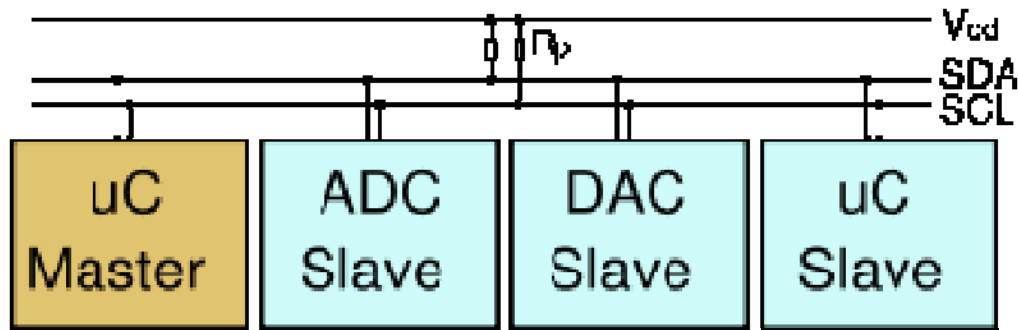


Figure 1: A Representation of an I2C Bus
V_{cd} is the power, SDA is for data, and SCL is for the clock signal

Inter Integrated Circuit

Developed in the early 1980s by Philips Semiconductor, I2C was developed as a low speed standard for onboard serial communications. Originally developed for use in television sets, it only took a few short years to become widely adopted. I2C is still used today in many applications, like computer monitors.

Before I2C every device used by a microcontroller had to be connected directly to the microcontroller. The problem with this was that as you increased the complexity of the system, you had to increase the amount of pins on the microcontroller to accommodate the additional devices.

With the addition of a bus on the circuit board, each device could connect to the bus and the microcontroller would only need to connect to the bus. All that was needed was a set of rules for how to use the bus. Even this early bus sounds a lot like a network, doesn't it?

The current version of the I2C specification is version 3.0. This version supports four categories of bus speeds:

Standard	100 kb/s
Fast	400 kb/s
Fast mode plus	1 Mb/s
High speed mode	3.4 Mb/s

One really cool capability of I2C is called clock stretching. What clock stretching means is that, the slave device can hold the clock signal down in a low position. This is useful because when a slave receives a read request, it can hold down the clock while it prepares to send the information to the master and then release the clock when it is ready to send.

Other companies were also busy in the early 1980s developing their own bus specifications. IBM was one of these companies. One of their early busses was the Industry Standard Architecture or ISA, bus. Actually, when it was developed IBM called it the XT bus.

IBM Specifications

Developed in 1981 as an 8bit bus, ISA was later modified to be a 16 bit bus architecture. The ISA bus supports speeds of up to 8Mb/s. During the 1980s, ISA gained widespread popularity, but it wasn't perfect. One major drawback to ISA was that the specification lacked detailed information on bus timing and the rules governing bus use. As a result many companies were forced to implement their own version of ISA with unique specifications. These unique versions of ISA were rarely compatible and caused a lot of headaches.

Memory and Storage

Russell Hitchcock

Another drawback to ISA was the intensive configurations required to connect a device to the bus. In 1993, ISA Plug and Play was developed to alleviate this problem. This plug and play architecture allowed the computer's operating system to do the configuring instead of the user. This was a major step forward in computer architecture and many of today's computers still support ISA plug and play.

In the late 1980s IBM attempted to replace the aging ISA bus with the Micro Channel Architecture or MCA, bus. This specification offered a 32 bit bus and automatic configuration.

While MCA alleviated the constraints of ISA's 16 bit bus and allowed for communications at up to about 40MB/s, there were also disadvantages which greatly hindered industry adoption of this bus specification. One disadvantage was that MCA was a proprietary technology, which of course meant that any company which wanted to use MCA had to pay IBM. Needless to say, not many people liked this idea.

Extended Industry Standard Architecture

Shortly after IBM developed MCA, IBM's competitors, including Compaq and HP, came up with a solution. They organized themselves, often referred to as the "Gang of Nine", and developed the Extended Industry Standard Architecture or EISA. Actually, they also renamed IBM's XT bus the ISA bus so that they wouldn't infringe on IBM's trademark.

Like MCA, EISA was a 32 bit bus architecture. But, like its name suggests, EISA was an extension of the early ISA bus specification which meant that it was backwards compatible with ISA devices. And of course EISA was an open specification. EISA supported parallel communications at a speed of 8.33MB/s.



Figure 2: Picture of Three EISA Slots

Memory and Storage

Russell Hitchcock

Largely due to the fact that EISA was an open architecture, it quickly gained popularity in the industry. This popularity however, was short lived. EISA was soon superseded by Intel's PCI bus.

The PCI bus is now being replaced by PCI Express or PCIe. This bus specification is really cool and blurs the line between network and bus even further. PCIe is a serial bus but pretends to be a parallel bus. It does this by having a hub on the motherboard. This hub can route device to device communications. It also allows more than one pair of devices to communicate at the same time, thus leading to parallelism. The PCIe specification allows for communications at up to 8GB/s, while PCI allows for just 133MB/s.

Direct Memory Access

As I mentioned earlier, modern computer busses can be quite complex. One complexity which greatly improves performance (at least from the user's perspective) is direct memory access, or DMA.

DMA allows certain hardware subsystems, like graphics cards and sound cards, to access memory directly for read and write operations independently of the CPU. Of course, the CPU is still required to initiate the transaction, but DMA eliminates the need for the CPU to be involved in tedious memory transactions and thus allows the CPU to be used for other tasks.

One newer bus specification is the AMD-backed HyperTransport specification, which is developed by the HyperTransport consortium. HyperTransport supports even faster communication speeds and advanced features such as DDR, or Double Data Rate. DDR allows devices to send data on both the rising and falling edges of a clock cycle. Basically this doubles the throughput.

Another interesting feature of HyperTransport is that it supports two kinds of write commands. One write command, used for DMA, is referred to as a posted write command. Unlike the unposted write command, the posted write command does not require a confirmation signal from the target device. This combined with the fact that HyperTransport is packet based, means that HyperTransport could really be considered a network. In fact it can also be used to connect multiple computers together.