Digital Recording Methods for Electromechanical Computer Hard Drives (HDD)

Dr. Andrew Blyth, PhD (andrew.blyth@adisa.global)

ADISA Research Centre, ADISA, Thrales End Business Centre, Thrales End Lane,

Harpenden AL5 3NS

1.0 Introduction

The role of this white paper is to define and document the current methods being utilised for the encoding and recoding of data onto a computer hard drive. The first hard drive was invented at IBM in 1957 and comprised a recording medium and head for reading and writing data. At an abstract level an electromechanical hard disk drive (HDD) functions to transform digital data into a magnetic wave form so that it can be written to, and read from, a recording medium. Over the past sixty years the size of a hard drive as increased from a few megabytes (2^{20} Bytes) to 14 terabytes (2^{40} Bytes). In 1957 a computer hard drive was 68 cubic feet in volume, by comparison a modern hard drive being 2.1 cubic inches on volume. In addition, in 1959 the data density of a hard drive is 10^3 bits per square inch, as compared with a data density today of 10^{12} bits per square inch [1].

This white paper will be structured as follows. Firstly, I will begin by defining what a computer electromechanical hard drive is and its key components, along with key definitions and terminology. Then I will document the method by which digital data is written to, and read from, the device.

2.0 An Introduction to an Electromechanical Hard Drive

At a physical level a HDD consists of the following elements.

2.1 A Metallic Case

The role and function of the metallic case is twofold. Firstly it is to provide a solid and stable base for the other elements of the disk drive to operate from. Secondly it is to shield the drive from electromagnetic interference as much as possible.

2.2 A Disk Platter

The disk platter is typically made from either glass or ceramic and is coated with a cobalt-based metallic alloy. This metallic oxide is the medium that will hold the magnetic wave form.

2.3 A Spindle and Spindle Motor

The disk platter is powered by a motor and rotates on an air bearing at speeds ranging from 7,200 rpm to 15,000 rpm.

2.4 Read/Write Head

The read/write head reads/writes the magnetic wave form to the platter and floats at a distance of between 7 – 10 nanometres (10^{-9} meters) above the platter. The read/write head is designed along the lines of an aircraft wing to generate lift. As the disk platter rotates under the head it moves the air with it. This movement of air over the read/write head functions to generate lift and thus keep the head at a fixed distance over the platter. The distance that the head is above the platter is a function of the number of air molecules that can fit between the head and the platter.

2.5 Actuator and Actuator Arm

The actuator arm is the device on which the read/write head is mounted. The Actuator is the device that positions the read/write head so that data can be read/written in tracks on a platter. A track is a lane on the platter which holds the data.

2.6 PCB/Controller Board

The PCB controller board is responsible for providing an interface between the computer hard drive and the computer. The PCB will also read/write digital data to/from the interface and position the read/write head via moving the actuator-arm so that data stored on the platter can be accessed. On modern computer hard drives the PCB will consist of the following components [2]:

2.6.1 Controller

On all modern computer hard drives the controller is typically an ARM based processor. It functions to execute the firmware (located in the ROM), and via hardware interfaces to the motor and actuator, it controls the location of the disk head so that data can be read/written to/from the platter. It also makes use of SDRAM as a cache for reading and writing data.

2.6.2 ROM/Firmware and SDRAM Memory

The SDRAM functions as a cache allowing for large volumes of data to be a) read from the hard drive and then forwarded to a recipient via the interface, or b) read from the interface and written to the disk platter. The firmware contained in the Read Only Memory (ROM) can be considered as the operating system that the controller executes. The role and function of the firmware will be discussed in greater detail in later sections of this white paper.

2.6.3 Interface

The typical interface for a computer hard drive is either Serial Advanced Technology Attachment (SATA) or Serial Attached SCSI (SAS). Legacy drives make use of interface standards such as IDE/PATA or SCSI [2]. The data transfer capacities for SATA are as follows:

- SATA revision 1.0 – Raw Data Rate: 1.5 Gbit/s & Data Rate: 150 MB/s
- SATA revision 2.0 – Raw Data Rate: 3 Gbit/s & Data Rate: 300 MB/s
- SATA revision 3.0 – Raw Data Rate: 6 Gbit/s & Data Rate: 600 MB/s

The data transfer capacities for SAS are as follows:

- SAS revision 1.0 – Raw Data Rate: 3 Gbit/s & Data Rate: 300 MB/s
- SAS revision 2.0 – Raw Data Rate: 6 Gbit/s & Data Rate: 600 MB/s
- SAS revision 3.0 – Raw Data Rate: 12 Gbit/s & Data Rate: 1.2 GB/s
- SAS revision 4.0 – Raw Data Rate: 22.5 Gbit/s & Data Rate: 2.25 GB/s

2.6.4 Motor Controller

The motor controller is responsible for regulating the speed of the platter and ensuring a fixed number of revolutions per minute are maintained.

3.0 An Electromechanical Hard Drive in Detail

The challenge that a HDD faces is how to take digital data and transform it into an analogue signal in such a manner that the transformation can be viewed as a two-way function, i.e. we can read and write data. This challenge is addressed via the placing of a magnetic wave form onto a platter. Thus, when considering an electromagnetic computer hard drive in detail, there are two components that we need to understand. The first is the disk head and how a Magnetic Wave Form (MWF) is read and written to a platter and the second is the role of the firmware.

3.1 The Magnetic Wave Form

A magnetic wave form is a series of electromagnetic pulses that are used to magnetise elements on a platter in either a north-south, or south-north orientation. These two orientations are used to denote either a one or a zero in the digital domain.

There are two recording methods utilised when read/writing data to a platter as a magnetic wave form, and they are a) longitudinal recording and perpendicular recording. In longitudinal recording is wave form is oriented along the platter and with perpendicular recording the wave form is oriented perpendicular to the platter – See Figure 1.

The magnetic wave form is recorded as a series of north-south pole orientations, with a north/South orientation corresponding to a one and a South/north orientation corresponding to a zero. When a bit of data is written to a disk platter in the form of a magnetic wave form the previous magnetic orientation is over written. The act of writing a bit of data is viewed as an atomic action, and thus it is not possible to recover the previous magnetic orientation for a bit of information.

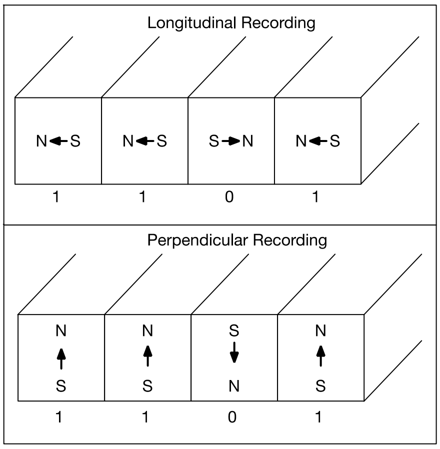

Figure 1 - Longitudinal Recording and Perpendicular Recording

To allow the hard drive to know the physical location where data is stored on a disk platter, a number of keys concepts are utilised. These are: Cylinders, Heads and Sectors, known as CHE. Figure 2 depicts the basic concepts. A hard drive will have a number of disk heads that are used to read/write data. Typically, there are two heads per platter. One disk head reads the top of the platter and the other reads the bottom of the platter. The surface of the platter is divided up into a number of tracks and a track is divided into a number of sectors. These three attributes allow a hard drive to locate an item of data on any platter surface.

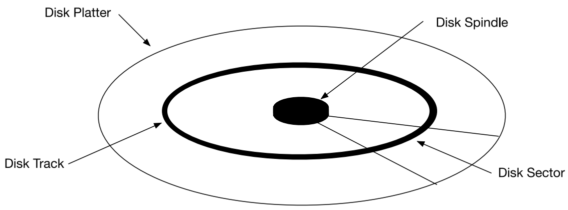

Figure 2 – Cylinder, Head and Sector.

A sector can typically store 512 bytes of user data, but in order for the hard disk to write/read the data it needs to first make use of a packet structure and error correction. On early computer hard drives, between each track there was free space called the inter-spatial gap. Post the year 2000, computer hard drives no longer make use of the inter-spatial gap and tracks are aligned next to each other to maximise data density. The packet structure allows the disk to know where a packet of data is, and how to encode/decode the packet. The basic structure of a sector is depicted in Figure 3.

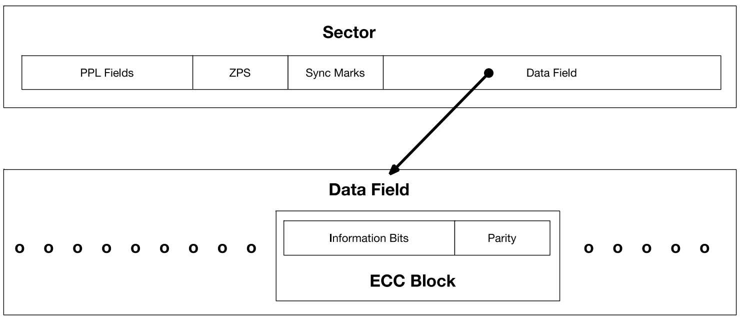

Figure 3 – The Structure of a Disk Sector

The following are the key elements of a Sector and are used to align the read head with the data field.

- PPL stands for Phase Lock Loop and is a pattern of 100 bits that is used for synchronisation.

- ZIPS is the zero-phase start pattern and is of 8-16 bits in length.

- The Sync Mark is used as a sector address marker for increased reliability and is 20 bits in length.

Error Correction Coding (ECC) is used to protect information bits from random noise interference and other unpredictable disturbances. The data field is made up of a set of ECC blocks, where each ECC block is made up of a set of information bits and a parity check. The role of the parity check is to identify errors in the information bits. While various algorithms exist for encoding information and parity checking, each disk vendor tend to use their own propriety algorithms for encoding the information bits and parity checking. In addition, each vendor will decide how many EEC blocks make up a data field. Vendors may also use different algorithms for encoding and parity checking for each EEC Block in a data field. To allow the disk head to position itself between sectors servo wedges are place on the disk between two sectors – See Figure 4

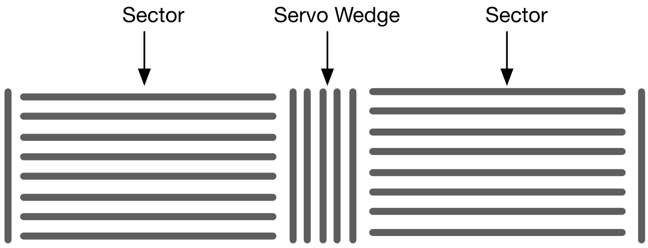

Figure 4 – Servo Wedge Alignment

As we can observe in Figure 4 sectors are written in a concentric horizontal manner in relation to the platter spindle, while server wedges are written in a perpendicular manner to the platter spindle. Specific fields within the servo wedges are used to determine the radial position of the heads relative to the array of concentric tracks and angular position. The servo wedges are read only and are read all the time while the drive is spinning [1]. This allows the firmware to position a head so that a sector can be read and written.

3.2 The Firmware

The firmware that the controller chip on the HDD PCB executes, functions as a mini-operating system. The firmware performs the following functions for the HDD:

- Interfaces with the computer via the ATA [3]/SCSI [4] interface such that functionality contained in the ATA/SCSI standards is implemented.

- Controls the ability to locate a specific cylinder, head and sector within the hard drive so as to locate a specific data field.

- Controls the read function for reading the magnetic wave form from the disk platter and transforming it into a digital signal.

- Controls the write function in taking a digital signal and transforming it into an analogue signal, writing it to the platter.

- Controls the spin speed of the disk platter and ensures that the disk spin speed is maintained at a fixed value.

4.0 Data Erasure Discussion

Before the year 2000, the ability of read/write heads to wobble gave rise to the possibility for data to be written into the interspatial gap on a platter. This meant that as data was written into a sector, the central location of the sector could move in relation to the horizontal axis of the platter in relation to the disk spindle via a very small amount, and that this small amount would not hinder the data being read back. This is called the splatter effect.

Because of the splatter effect data could now be written into the interspatial gap due to disk head wobble, this gave rise to a possible attack vector for data recovery. This attack said that even if the sector had been over written it may be possible to access the data due to it being written into the interspatial gap via magnetic force microscopy. To mitigate this attack vector various data erase standards proposed that data be written to a sector multiple times. These standards include the following:

- Peter Gutmann's Algorithm
- Bruce Schneier's magnetic orientation
- U.S. Navy Staff Office Publication NAVSO P-5239-26
- U.S. Airforce System Security Instruction 5020

After the year 2000, disk head drive technology had matured removing the possibility for data to be written into the interspatial gap and hence sectors could now be written next to each other. As sectors are now being written next to each other and there is no interspatial gap between sectors we can assert that overwriting data to a sector once is sufficient to render the previous data held in the data unrecoverable.

5.0 References

[1] B. Vasic and E. Kutas, **Coding and Signal Processing for Magnetic Recording Systems**, CRC Press, 2005.

[2] R. Thompson and B Thompson, **PC Hardware in a Nutshell**, 3rd Edition, O'Reilly, 2003

[3] **The Path from 3Gb/s to SATA 6Gb/s: How to Migrate Current Designs to the SATA Revision 3.0 Specification**, https://sata-io.org, Last Accessed 16th July 2019.

[4] **24G SAS Data Storage Specification Development Complete; SCSI Trade Association Spotlights Technology at 2017 Flash Memory Summit,** http://www.scsita.org, Last Accessed 16th July 2019.