

DATA STORAGE AND RETRIEVAL

5 FIELD OF INVENTION

The present invention relates broadly to a method of storing data on a track of a storage disk, a method of retrieving data on a track of a storage disk, and to a controller system for a storage disk drive.

10

BACKGROUND

Data storage devices such as magnetic hard disk drives are required by systems to guarantee a certain BER (bit error rate). For example, in a typical magnetic recording system a BER of 10^{-12} is desired. This BER of 10^{-12} translates to an overall SNR (Signal to Noise Ratio) of approximately 19-20dB. The two major components to providing the desired SNR are the front end electronics detecting readback waveform from head medium and the back end signal processing.

20 The front end SNR is largely determined by the recording density. In order to store more bits in a given magnetic disk, hard disk manufacturers have sought to increase recording densities. However, the higher recording density directly translates to a lower SNR. In order to overcome the lower SNR, hard disk manufacturers have sought to improve signal processing techniques at the back end. For example, R. Wood, IEEE
25 TRANSACTIONS ON MAGNETICS, VOL. 36, No. 1, JANUARY 2000, discloses a special coding scheme to tolerate low SNR at the recording media front end. However, even with the improvements in the signal processing techniques at the back end, an SNR of about 10dB at the recording media front end is still required to meet typical system requirements on BER.

30

Therefore, there exists a need to provide a system and method for data storage and retrieval to address one or more of the problems mentioned above.

SUMMARY

In accordance with a first aspect of the present invention, there is provided a method of storing data on a track of a storage disk comprising the steps of: storing
5 the data in a first zone of the track; identifying errors in the stored data; and storing information associated with the identified errors in a second zone of the track.

The information may identify locations of the respective errors in the first
10 zone.

The information may identify a sector number and a bit number of each error.

Identifying the error may comprise: reading the data stored in the first zone;
and comparing the read data with original data intended for storage.

15 The comparing may comprise performing a XOR operation between the read data and the original data.

The information associated with the identified errors may comprise an error
20 matrix based on the XOR operation for consecutive bits in the read data and the original data.

The second zone may be disposed at an end of the track following the first
25 zone.

The steps of identifying errors in the stored data and storing information
associated with the identified errors in a second zone of the track may occur in one
revolution of the storage disk.

30 The first and second zones may each comprise two or more segments
alternatingly disposed along the track.

The step of storing information associated with the identified errors in a
second zone of the track may comprise storing information associated with the

identified errors from a first segment of the first zone in a first segment of the second zone in one revolution of the storage disk; and storing information associated with the identified errors from a second segment of the first zone in a second segment of the second zone in a next revolution of the storage disk.

5

The data in the first zone may be stored at a higher density than the information in the second zone.

10 In accordance with a second aspect of the present invention, there is provided a method of retrieving data stored on a track of a storage disk comprising the steps of: reading data from a first zone of the track; reading information associated with errors in the read data from a second zone of the track; and correcting errors in the read data based on the retrieved information.

15 The information may identify locations of the respective errors in the first zone.

The information may identify a sector number and a bit number of each error.

20 The data in the first zone may be stored at a higher density than the information in the second zone.

Correcting the error may comprise performing a XOR operation between the read data and the retrieved information.

25

The second zone may be disposed at an end of the track following the first zone.

30 The first and second zones may each comprise two or more segments alternatingly disposed along the track.

In accordance with a third aspect of the present invention, there is provided a controller system for a storage disk drive, the system comprising: means for controlling the storage disk drive to store the data in a first zone of a track of a

magnetic disk; means for identifying errors in the stored data; and means for controlling the disk drive to store information associated with the identified errors in a second zone of the track.

5 In accordance with a fourth aspect of the present invention, there is provided a controller system for a storage disk drive, the system comprising: means for controlling the storage disk drive to read data from a first zone of a track of a magnetic disk; means for controlling the disk drive to read information associated with errors in the read data from a second zone of the track; and means for
10 correcting errors in the read data based on the retrieved information.

BRIEF DESCRIPTION OF THE DRAWINGS

 Embodiments of the invention will be better understood and readily apparent
15 to one of ordinary skill in the art from the following written description, by way of example only, and in conjunction with the drawings, in which:

 Figure 1 is a plan view of a hard disk drive implementing an example embodiment of the present invention.
20

 Figure 2a shows a read and write head of the hard disk drive implemented in an example embodiment of the present invention.

 Figure 2b is a plan view of a read and write head of the hard disk drive
25 implemented in an example embodiment of the present invention.

 Figures 3a and 3b are charts illustrating the BER versus SNR curve of magnetic recording systems.

30 Figure 4a shows an example grain structure which provides high media SNR.

 Figure 4b shows an example grain structure which doubles the areal density provided by the example grain structure of Figure 3a.

Figure 5 is a chart illustrating the increase in BER when the areal density is doubled.

Figure 6 illustrates a conventional disk track and a disk track in an embodiment of the present invention.

Figure 7 illustrates the write process in an example embodiment of the present invention.

Figure 8 is a schematic diagram of error identification based on the comparison between the reproduced user data sequence and the original user data in the buffer.

Figure 9 is a schematic diagram of a read process in an embodiment of the present invention.

Figure 10 illustrates a disk track implemented in another embodiment of the present invention.

Figure 11 shows a flow chart illustrating a method of storing data on a track of a magnetic disk implemented in an example embodiment.

Figure 12 shows a flow chart illustrating a method of retrieving data stored on a track of a magnetic disk implemented in an example embodiment.

DETAILED DESCRIPTION

Embodiments of the present invention seek to reserve a portion of disk track space in each disk track of a storage device, and to record location(s) of bit errors of the current disk track data in the reserved zone. In particular, in one embodiment of the present invention, the storage device is a magnetic hard disk wherein a portion of magnetic disk track space in each disk track is reserved, and location(s) of bit errors of the current disk track data are recorded in the reserved zone. This reserved portion may comprise e.g. 1-2% of the total disk track space. During the data reproducing process, error location(s) can be recalled from the reserved zone to allow error(s) to be corrected.

As errors can be corrected, data may be stored in the data zone of the disk track space with a lower SNR (Signal to Noise Ratio) and corresponding higher BER (Bit error rate). Thus, data bits can be recorded with a reduced number of magnetic grains per bit, and the areal density of the disk can be increased significantly.

5

Some portions of the description which follows are explicitly or implicitly presented in terms of algorithms and functional or symbolic representations of operations on data within a computer memory. These algorithmic descriptions and functional or symbolic representations are the means used by those skilled in the data processing arts to convey most effectively the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities, such as electrical, magnetic or optical signals capable of being stored, transferred, combined, compared, and otherwise manipulated.

10

15

Unless specifically stated otherwise, and as apparent from the following, it will be appreciated that throughout the present specification, discussions utilizing terms such as "receiving", "amplifying", "reading", "processing", "writing", or the like, refer to the action and processes of a computer system, or similar electronic device, that manipulates and transforms data represented as physical quantities within the system into other data similarly represented as physical quantities within the system or other information storage, transmission or display devices.

20

The present specification also discloses apparatus for performing the operations of the methods. Such apparatus may be specially constructed for the required purposes, or may comprise a general purpose computer or other device selectively activated or reconfigured by a computer program stored in the computer. The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose machines may be used with programs in accordance with the teachings herein. Alternatively, the construction of more specialized apparatus to perform the required method steps may be appropriate. The structure of a conventional general purpose computer will appear from the description below.

25

30

In addition, the present specification also implicitly discloses a computer program, in that it would be apparent to the person skilled in the art that the individual steps of the method described herein may be put into effect by computer code. The computer program is not intended to be limited to any particular programming language and implementation thereof. It will be appreciated that a variety of programming languages and coding thereof may be used to implement the teachings of the disclosure contained herein. Moreover, the computer program is not intended to be limited to any particular control flow. There are many other variants of the computer program, which can use different control flows without departing from the spirit or scope of the invention.

Furthermore, one or more of the steps of the computer program may be performed in parallel rather than sequentially. Such a computer program may be stored on any computer readable medium. The computer readable medium may include storage devices such as magnetic or optical disks, memory chips, or other storage devices suitable for interfacing with a general purpose computer. The computer readable medium may also include a hard-wired medium such as exemplified in the Internet system, or wireless medium such as exemplified in the GSM mobile telephone system. The computer program when loaded and executed on such a general-purpose computer effectively results in an apparatus that implements the steps of the preferred method.

In an embodiment of the present invention, a writing process uses two revolutions of the magnetic disk to complete recording of one track of data. In the first revolution, the magnetic writing head records data at a high linear density in a data zone. For a second revolution, a read head reads the previously recorded data, transmits the read data through the channel and detects reproduced data from the readback waveform. The reproduced data is compared with the original data temporarily stored in the buffer, to identify error locations within the recorded data stream. As the read and write head accesses the entire data zone followed by the error recording zone, when the read head reaches the end of the data zone, the comparison can be completed and the identified error locations can now be written into the error recording zone. The magnetic writing head arrives at the reserved error recording zone shortly after the reading head, as the magnetic writing head trails the read head. Therefore, the magnetic writing head preferably has enough time to respond to the comparisons and record error location information in the error recording zone. The error location information is written at a

lower linear density, as compared with the data in the data zone, such that the BER of the error recording zone data is preferably below e.g. 10^{-12} .

In the embodiments of the present invention, a read process uses one revolution
5 of the magnetic disk, wherein the read head reads the data in the data zone, as well as the error locations in the reserved zone. After the reading of the data stream in the entire track data, the errors in the data zone after the channel detection (at the BER level of e.g. 10^{-4}) are corrected based on the pre-recorded error location information. Although
10 there is a significant BER degradation as a result of the reduction in SNR by for example 6dB, overall performance of the magnetic recording system can be maintained at a desired threshold, while allowing the linear or areal density of the magnetic recording system to increase.

Figure 1 shows a plan view of a hard disk drive 100 comprising a head media
15 system 101 which further comprises suspension arm 102 with a read and write head 104 in relation to a magnetic disk/media 106. The read and write head comprises of separately operable read head 104a and write head 104b (illustrated in Figure 2). The disk track presently being assessed by the read and write head 104 is shown as numeral 108. It will be appreciated that the magnetic disk 106 comprises a plurality of circular
20 tracks 107 forming progressively larger concentric rings centered about the center 109 of the magnetic disk 106.

Figure 2a shows a read and write head 104 of the hard disk drive 100 implemented in an example embodiment of the present invention. The write head 104b
25 comprises a write pole 202 and a return pole 204. The read head 104a comprises of a reader 208 flanked by a pair of reader shields 210, 212. In the example embodiment, the write head 104b trails the read head 104a as the disk revolves in a manner such that the read head 104a accesses a sector of the disk before the write head 104b. Figure 2a shows the disk moving from right to left. Figure 2b shows a plan view of the read and
30 write head 104 implemented in an example embodiment.

The hard disk drive 100 further comprises a controller 110, which manages data transfers to and from the hard disk drive 100. The controller 110 comprises a channel encoder 114 for encoding user input data 112 into recording data 116. The recording

data 116 is stored in a write data buffer 118 before it is written by the write head 104b into the magnetic disk 106. To obtain data from the hard disk drive 100, the controller further comprises a pre-amplifier 120 for receiving and amplifying the data read from the disk media 106 using the read head 104a into a Readback waveform 122. The
 5 Readback waveform 122 passes through a read channel 124, to generate reproduced data 126. The reproduced data 126 is then decoded by the channel decoder 128 to provide the output user data. In the example embodiments, the read channel 124 performs corrections to the read data using error information that has been stored in an error zone to generate reproduced data 126, thus providing additional gain to the overall
 10 SNR of the hard disk drive 100. It will be appreciated that the channel encoder 114 and channel decoder 128 advantageously provide signal processing to the data signals to further improve the overall SNR performance of the hard disk drive 100.

It will be appreciated by a person skilled in the art that the controller 110 may
 15 be implemented as a programmable integrated circuit. It may be implemented as software such as a program being executed within the integrated circuit and instructing the controller 110 to conduct the method of the example embodiment. The integrated circuit may be connected to a computer via suitable connectors e.g. SATA or other Bus interface. The integrated circuit may comprise a processor, a
 20 Random Access Memory (RAM) and a Read Only Memory (ROM). The components of the integrated circuit typically communicate via an interconnected bus and in a manner known to the person skilled in the relevant art.

From the user perspective, the user data is preferable to be recorded at a high
 25 head media SNR with a BER level of e.g. below 10^{-12} . Figure 3a shows the BER versus SNR curve of a magnetic recording system. As shown in Figure 3a, in order to have a final BER of below 10^{-12} , the overall system SNR is preferably in the region of around 19 dB. Of the total 19 dB SNR, the head media system 101 (Figure 1) may provide an SNR of around 13 to 14 dB (see arrow 202 of Figure 3b), whilst the channel and signal
 30 processing gains from the controller 110 (Figure 1) amount to another 5 to 6 dB SNR (see arrow 304 of Figure 3b) for the overall system to have the BER below 10^{-12} .

It will be appreciated by a person skilled in the art that the recording media e.g. 106 (figure 1) of hard disk drives uses a certain number of magnetic grains to form one

recording bit. As this number (e.g. 22) of magnetic grains for each bit is high, the media SNR is also high. Figure 4a shows an example grain structure 400 which provides high media SNR where e.g. 22 magnetic grains e.g. 402 form one recording bit 404. When the magnetic grain number per bit 406 reduces to e.g. 11, as shown in Figure 4b, the media SNR decreases significantly. The reduced number of magnetic grains per bit will reduce the SNR and this may deteriorate the overall system BER to the level around 10^{-4} , as shown by arrow 502 in Figure 5. However, with the magnetic grain number per bit reducing by half from the bit grain configuration in Figure 4a to that in Figure 4b, the areal density of magnetic recording system is doubled, thereby allowing more bits to be stored in the same area.

Figure 6 shows a comparison of a conventional disk track 600 and a disk track 610 in an embodiment of the present invention. It will be appreciated that the disk tracks 600 and 610 are shown as linear in shape for illustrative purposes only. As shown in the hard disk drive 100 of Figure 1, tracks e.g. 107 are concentric rings formed on the recording media 107. Returning to Figure 6, the conventional disk track 600 comprises of a single data zone 602. In contrast, a portion of space in the disk track 610, e.g. 1-2% of the overall track space, is reserved as the error recording zone 614. Therefore, the disk track 610 has two zones: a data zone 612 and an error recording zone 614. Compared with the data zone 602 of the conventional disk track 600, the data zone 612 of the disk track 610 records the user data at a higher bit density. In other words, more bits are stored in each inch of the data zone 612, than in each inch of the conventional data zone 602. For illustrative purposes, the conventional data zone 602 may store data at a bit density of 1 MBPI (Megabits per inch), while the data zone 612 in one example embodiment of the data track 610 stores data at a bit density of 2 MPBI. The higher bit density of e.g. 2MBPI in the data track 610 translates to a lower media SNR and hence higher bit error rate.

Given the higher bit error rate of the data track 610, the error recording zone 614 records the data locations/addresses within the data zone 612 wherein incorrect or erroneous data is stored. In this way, errors in the data stored in the data zone 612 can be corrected, based on their error locations, during the read process. A more detailed description of the read process is provided below. The linear recording density in the error recording zone is preferably kept low, e.g. at the conventional level of 1MBPI.

The write process in accordance to an embodiment of the present invention uses two revolutions to accomplish one track recording. Figure 7 illustrates the write process in the example embodiment of the present invention. As shown in figure 7 and with
 5 reference to Figure 1, in the first revolution, encoded user data or recording data 116 (Figure 1) stored in the buffer 118 (Figure 1) is written in the data zone 612 of the track 610 using the write head 104b (Figure 1) of the hard disk drive at a higher recording density of e.g. 2MBPI. No data is written into the error recording zone 614. Arrow 702 shows the direction and area where the write head 104b (Figure 1) writes data in the
 10 track 610.

At the second revolution, the read head 104a (Figure 1) reads the previously written data in the data zone 612, in the direction and track area as shown by arrow 704. The read data passes through the pre-amplifier 120 (Figure 1) and read channel 124
 15 (Figure 1) before reproduced data 126 (Figure 1) is obtained. The reproduced data 126 (Figure 1) is compared with the recording data 116 (Figure 1) in the buffer 118 (Figure 1) for errors. The location(s)/address(es) of error data in the data zone 704 of the disk track is determined and the error location information is recorded in the error recording zone 614, at a lower bit density than that in the data zone 612. This recording of error
 20 locations is represented by arrow 706 in Figure 7.

Figure 8 shows the schematic diagram of error identification based on the comparison between the reproduced data sequence 804 and the original data 802 in the buffer. An "XOR" operation is conducted in sequence to both reproduced user data and
 25 original user data. The resultant output sequence or error matrix 806 will show a "1" at bit locations where an error is identified e.g. 808, 810 and a "0" where the bit is correct.

After the comparison of all the data in the present disk track is complete, the bit locations of all the errors on this track can be identified. The location of each error
 30 comprises of the information of the sector number and the bit number of the error within the sector. In an example configuration of a hard disk drive, wherein there are below 4000 sectors per track, with each sector comprising 4096 bits, 12 bits can be used to store the sector number where the error occurs, and a further 12 bits can be used to further record the bit number of the error within the sector. In another example

embodiment where the number of sectors is about 8000, an additional bit or a total of 13 bits can be used to record the sector number where the error occurs.

In the example where 24 bits are used to record one error, the bits can be stored
 5 at the lower bit/linear density of e.g. 1 MBPI. In such an example embodiment where 48 bits (adjusted for the space not available for data writing at 2MBPI in the data zone) are "sacrificed" to record the error information. Assuming the overall system BER is at 10^{-4} , an error occurs for every 10000 bits in the data zone. As such, 48×10^{-4} , ~0.48% of the
 10 total disk track space can be used as the error recording zone to record the errors of the data zone. If it is assumed that the probability of error occurrence is a Gaussian distribution, a reservation of 1.5% of the total disk track space can cover three times of the standard deviation of the error occurrence. This translates to almost full coverage of all errors in a data zone by using just 1.5% of the total disk track space as the error recording zone in one example embodiment.

15 Figure 9 shows a schematic diagram of a read process in accordance to one embodiment of the present invention. The error location information e.g. 902, 904 are converted into a data sequence 906 of the same length as the reproduced data 908 from the data zone. In the data sequence 906, alternatively known as the error matrix, each
 20 error location is represented by a "1" e.g. 916, 918 while a "0" indicates correct data. An "XOR" operation is conducted between each corresponding bit of the data sequence 906 and the reproduced data 908 such that errors e.g. 910, 912 can be corrected with the subsequent output of the operation generating the originally stored data 914.

25 In the earlier example embodiments described above, with reference to e.g. Figure 7, the second revolution of the write process involves the steps of:

- obtaining the reproduced data from the read channel;
- comparing the reproduced data with the data stored in the write data buffer;
- identifying the error locations in the reproduced data; and
- 30 writing error information based on the error locations in the error recording zone.

It will be appreciated that the steps of obtaining the reproduced data from the read channel; comparing the reproduced data with the data stored in the write data buffer; and identifying the error locations in the read data, are preferably completed

before the subsequent writing step can proceed. In general, as magnetic writing heads of hard disk drives tend to trail the reading heads when accessing the track, there is a brief period available between the time where the reading head completes reading, and the time when the magnetic writing head begins writing beyond the previously read area.

- 5 This brief period is preferably sufficient for the obtaining, comparing and identifying steps to be completed, such that the writing step can be performed.

The inventors have recognised that there may be instances where the period is insufficient for the obtaining, comparing and identifying steps to be complete. For
10 example, obtaining the reproduced data from the read channel may be time consuming given the signal processing involved. An alternative embodiment of the track structure may therefore be implemented. Figure 10 shows a disk track 1010 implemented in an alternative embodiment of the present invention. A conventional disk track 1000 is also shown for comparison. It will be appreciated that the disk tracks 1000 and 1010 are
15 shown as linear in shape for illustrative purposes only. As shown in the hard disk drive 100 of Figure 1, tracks e.g. 107 are concentric rings formed on the recording media 107. Returning to Figure 10, the conventional disk track 1000 comprises of a single data zone 1002. In contrast, two continuous portions of space in the disk track 1010 occupying, e.g. 1-2% of the overall track space is reserved as the error recording zone 1014a,
20 1014b. Therefore, the disk track 1010 has four zones: e.g. a first data zone 1012a, a second data zone 1012b, a first error recording zone 1014a, and a second error recording zone 1014b. While the example embodiment described in Figure 10 show two separate data and respective error recording zones, it will be appreciated that any number of such data and respective error recording zones may be implemented. Where
25 data was previously stored as a single data stream in the embodiment illustrated in Figure 6, the same data is stored as e.g. two separate streams. The data zones 1012a, 1012b of the disk track 1010 each record user data at a higher bit density than the data zone 1002 of disk track 1000. In other words, more bits are stored in each inch of the data zones 1012a, 1012b, than each inch of the conventional data zone 1002. For
30 illustrative purposes, the conventional data zone 602 stores data at a bit density of 1 MBPI (Megabits per inch), while the data zones 1012a, 1012b in the example embodiment of the data track 1010 each store data at a bit density of 2 MBPI. The higher bit density of e.g. 2MBPI in the data track 1010 translates to a lower media SNR and hence higher bit error rate.

Each error recording zone 1014a, 1014b caters to the data zones 1012b, 1012a respectively by interleave. In this way, errors in the data stored in the data zone 1012a can be corrected by the error information located in the error recording zone 1014b, while errors in the data stored in the data zone 912b can be corrected by the error information located in the error recording zone 1014a. This provides more time for channel detection. During a write process, data is first written in the data zones 1012a and 1012b. Thereafter, to obtain reproduced data (for comparison with original the data zone 1012a is read and subsequently detected. At the same time as the data detection of data zone 1012a, data zone 1012b can be read. This may result in time savings as data detection of data zone 1012a can be completed while data zone 1012b is read. The determined error locations in data zone 1012a are written in error zone 1014b after the finish of the reading of data zone 1012b. Similarly, the error locations in 1012b are recorded in 1014a with a half revolution delay. In this regard, the example embodiment uses 2.5 revolutions to finish writing of a single track.

It will be appreciated that the improvements to the density of the recording of data in the magnetic drives may be dependent on the shield to shield spacing of the read head. The pitch of each bit is preferably chosen such that interference from adjacent bits is reduced or minimised. The smaller shield to shield spacing of read head is preferred for higher readback resolution. This can enable the read head to pick up stronger signals from high density bits.

It will be appreciated that the example embodiments of the present invention can provide bit error correction so as to reduce the overall system BER. Further, with the reduced overall BER, the example embodiments of the present invention can be implemented to increase the density of recording of data in magnetic drives, while keeping the BER of the magnetic drive at a desired level.

It will also be appreciated that while the example embodiments describe a doubling of the bit density in the data zones, the density in the data zones may be increased by any factor, with the corresponding error recording zones sized accordingly.

Figure 11 shows a flow chart illustrating a method 1100 of storing data on a track of a storage disk implemented in an example embodiment. At step 1102, data is stored in a first zone of the track. At step 1104, errors in the stored data are identified. At step 1106, information associated with the identified errors is stored in a second zone of the track.

Figure 12 shows a flow chart 1200 illustrating a method of retrieving data stored on a track of a storage disk implemented in an example embodiment. At step 1202, data is read from a first zone of the track. At step 1204, information associated with errors in the read data is read from a second zone of the track. At step 1206 errors in the read data are corrected based on the retrieved information.

It will be appreciated by a person skilled in the art that numerous variations and/or modifications may be made to the present invention as shown in the specific embodiments without departing from the spirit or scope of the invention as broadly described. The present embodiments are, therefore, to be considered in all respects to be illustrative and not restrictive.

It will be appreciated that embodiments of the present invention can be extended to other recording mediums such as optical recording mediums e.g. CD, DVD, Blue-Ray etc. While the readback resolution of optical recording is dependant on the laser spot size, embodiments of the present invention may still provide for increased media densities.

CLAIMS

1. A method of storing data on a track of a storage disk comprising the steps of:

5 storing the data in a first zone of the track
identifying errors in the stored data; and
storing information associated with the identified errors in a second zone of the track.

10 2. The method as claimed in claim 1, wherein the information identifies locations of the respective errors in the first zone.

3. The method as claimed in claim 2, wherein the information identifies a sector number and a bit number of each error.

15 4. The method as claimed in any one of the preceding claims, wherein identifying the error comprises:
reading the data stored in the first zone; and
comparing the read data with original data intended for storage.

20 5. The method as claimed in claim 4, wherein the comparing comprises performing a XOR operation between the read data and the original data.

25 6. The method as claimed in claim 5, wherein the information associated with the identified errors comprises an error matrix based on the XOR operation for consecutive bits in the read data and the original data.

7. The method as claimed in any one of the preceding claims, wherein the second zone is disposed at an end of the track following the first zone.

30 8. The method as claimed in claim 7 wherein the steps of identifying errors in the stored data and storing information associated with the identified errors in a second zone of the track occur in one revolution of the storage disk.

9. The method as claimed in any one of claims 1 to 6, wherein the first and second zones each comprises two or more segments alternatingly disposed along the track.

5 10. The method as claimed in claim 9, wherein the step of storing information associated with the identified errors in a second zone of the track comprises

storing information associated with the identified errors from a first segment of the first zone in a first segment of the second zone in one revolution of the
10 storage disk; and

storing information associated with the identified errors from a second segment of the first zone in a second segment of the second zone in a next revolution of the storage disk.

15 11. The method as claimed in any one of the preceding claims, wherein the data in the first zone is stored at a higher density than the information in the second zone.

12. A method of retrieving data stored on a track of a storage disk
20 comprising the steps of:

reading data from a first zone of the track;

reading information associated with errors in the read data from a second zone of the track; and

correcting errors in the read data based on the retrieved information.

25 13. The method as claimed in claim 10, wherein the information identifies locations of the respective errors in the first zone.

14. The method as claimed in claim 11, wherein the information identifies
30 a sector number and a bit number of each error.

15. The method as claimed in any one of claims 12 to 14, wherein the data in the first zone is stored at a higher density than the information in the second zone.

16. The method as claimed in any one of claims 12 to 15, wherein correcting the error comprises:

5 performing a XOR operation between the read data and the retrieved information.

17. The method as claimed in any one of claims 12 to 16, wherein the second zone is disposed at an end of the track following the first zone.

10 18. The method as claimed in any one of claims 12 to 17, wherein the first and second zones each comprises two or more segments alternatingly disposed along the track.

15 19. A controller system for a storage disk drive, the system comprising:
means for controlling the storage disk drive to store the data in a first zone of a track of a magnetic disk;
means for identifying errors in the stored data; and
means for controlling the disk drive to store information associated with the identified errors in a second zone of the track.

20

20. A controller system for a storage disk drive, the system comprising:
means for controlling the storage disk drive to read data from a first zone of a track of a magnetic disk;
means for controlling the disk drive to read information associated with errors
25 in the read data from a second zone of the track; and
means for correcting errors in the read data based on the retrieved information.