



- (51) **International Patent Classification:**
G11B 5/09 (2006.01) *G11B 5/596* (2006.01)
- (21) **International Application Number:**
PCT/US2016/014950
- (22) **International Filing Date:**
26 January 2016 (26.01.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
- | | | |
|------------|-------------------------------|----|
| 62/083,732 | 24 November 2014 (24.11.2014) | US |
| 62/083,696 | 24 November 2014 (24.11.2014) | US |
| 62/108,929 | 28 January 2015 (28.01.2015) | US |
| 14/688,515 | 16 April 2015 (16.04.2015) | US |
- (71) **Applicant:** SEAGATE TECHNOLOGY LLC [US/US];
10200 S De Anza Blvd, Cupertino, CA 95014 (US).
- (72) **Inventors:** GAO, Kaizhong; 2 Gadwall Lane, North Oaks,
MN 55127 (US). ZHU, Wenzhong; 13073 Elderberry Ct.,
Apple Valley, MN 55124 (US). RAUSCH, Tim; 19600
Cambridge Court, Farmington, MN 55024 (US). GAGE,
Edward; 5234 167th St West, Lakeville, MN 55044 (US).
- (74) **Agent:** SMITH, Teryl, L.; Holzeriplaw, PC, 216 16th
Street, Suite 1350, Denver, CO 80202 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

[Continued on next page]


(54) **Title:** NON-SEQUENTIAL WRITE FOR SEQUENTIAL READ BACK

FIG. 2

(57) **Abstract:** A storage device controller addresses consecutively-addressed portions of incoming data to consecutive data tracks on a storage medium and writes the consecutively-addressed portions to the consecutive data tracks in a non-consecutive track order. In one implementation, the storage device controller reads the data back from the consecutive data tracks in a consecutive address order in a single sequential read operation.



-
- *the filing date of the international application is within two months from the date of expiration of the priority period (Rule 26bis.3)*

NON-SEQUENTIAL WRITE FOR SEQUENTIAL READ BACK

Cross-Reference to Related Applications

The present application claims benefit of priority to U.S. Provisional Application no. 62/083,696, entitled “Interlaced Magnetic Recording in HAMR Devices” and filed on November 24, 2014, and also to U.S. Provisional Patent Application no. 62/083,732, entitled “Interlaced Magnetic Recording” and filed on November 24, 2014, and also to U.S. Provisional Patent Application no. 62/108,929, entitled “Non-Sequential Write for Sequential Read Back” and filed on January 28, 2015, and also to U.S. Patent Application no. 14/688,515, entitled “Non-Sequential Write for Sequential Read Back” and filed on April 16, 2015. All of these applications are specifically incorporated by reference for all that they disclose or teach.

Technical Field

This invention relates generally to a storage device controller and data management techniques, such as techniques for addressing data and for selecting data tracks to receive data.

Background

As requirements for data storage density increase for magnetic media, cell size decreases. A commensurate decrease in the size of a write element is difficult because in many systems, a strong write field gradient is needed to shift the polarity of cells on a magnetized medium. As a result, writing data to smaller cells on the magnetized medium using the relatively larger write pole may affect the polarization of adjacent cells (e.g., overwriting the adjacent cells). One technique for adapting the magnetic medium to utilize smaller cells while preventing adjacent data from being overwritten during a write operation is shingled magnetic recording (SMR).

SMR allows for increased areal density capability (ADC) as compared to conventional magnetic recording (CMR) but at the cost of some performance ability. As used

herein, CMR refers to a system that allows for random data writes to available cells anywhere on a magnetic media. In contrast to CMR systems, SMR systems are designed to utilize a write element with a write width that is larger than a defined track pitch. As a result, changing a single data cell within a data track entails re-writing a corresponding group of shingled (e.g., sequentially increasing or decreasing) data tracks.

Better designs are desired to increase ADC beyond that provided by CMR and SMR systems.

Brief Descriptions of the Drawings

FIG. 1 illustrates a data storage device including a transducer head assembly for writing data on a magnetic storage medium.

FIG. 2 illustrates data writes in an interlaced magnetic recording (IMR) system employing an example non-sequential write scheme for sequential data read back.

FIG. 3 illustrates a magnetic disc including a number of high-density data bands each separated from one another by one or more low-density data tracks.

FIG. 4 illustrates additional example writes in an IMR system employing an example non-sequential write scheme for sequential file read back.

FIG. 5 illustrates example operations for performing a non-sequential data write that provides for sequential data read back in an IMR system.

Summary

Implementations disclosed herein provide a storage device controller configured to receive data addressed to a consecutive sequence of logical block addresses (LBAs) corresponding to a consecutive sequence of data tracks on a storage medium and to write the data to the consecutive sequence of data tracks in a non-consecutive track order.

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter. These and various other

features and advantages will be apparent from a reading of the following Detailed Description.

Detailed Description

FIG. 1 illustrates a data storage device 100 including a transducer head assembly 120 for writing data on a magnetic storage medium 108. Although other implementations are contemplated, the magnetic storage medium 108 is, in FIG. 1, a magnetic storage disc on which data bits can be recorded using a magnetic write pole (e.g., a write pole 130) and from which data bits can be read using a magnetoresistive element (not shown). As illustrated in View A, the storage medium 108 rotates about a spindle center or a disc axis of rotation 112 during rotation, and includes an inner diameter 104 and an outer diameter 102 between which are a number of concentric data tracks 110. Information may be written to and read from data bit locations in the data tracks on the storage medium 108.

The transducer head assembly 120 is mounted on an actuator assembly 109 at an end distal to an actuator axis of rotation 115. The transducer head assembly 120 flies in close proximity above the surface of the storage medium 108 during disc rotation. The actuator assembly 109 rotates during a seek operation about the actuator axis of rotation 112. The seek operation positions the transducer head assembly 120 over a target data track for read and write operations.

The transducer head assembly 120 includes at least one write element (not shown) including a write pole for converting a series of electrical pulses sent from a controller 106 into a series of magnetic pulses of commensurate magnitude and length. The magnetic pulses of the write pole selectively magnetize magnetic grains of the rotating magnetic media 108 as they pass below the pulsating write element.

View C illustrates magnified views 114, 116, and 118 of a same surface portion of the storage media 108 according to different write methodologies and settings of the data storage device 100. Specifically, the magnified views 114, 116, and 118 include a number of magnetically polarized regions, also referred to herein as “data bits,” along the data tracks of the storage media 108. Each of the data bits (e.g., a data bit 128) represents one or more individual data bits of a same state (e.g., 1s or 0s). For example, the data bit 128 is a

magnetically polarized region representing multiple bits of a first state (e.g., “000”), while an adjacent data bit 126 is an oppositely polarized region representing one or more bits of a second state (e.g., a single “1”). The data bits in each of the magnified views 114, 116, and 118 are not necessarily illustrative of the actual shapes or separations of the bits within an individual system configuration.

The magnified view 114 illustrates magnetic transitions recorded according to a conventional magnetic recording (CMR) technique. Under CMR, all written data tracks are randomly writeable and have substantially equal written track width. “Written track width” refers to, for example, a width measured in a cross-track direction of a recorded data bit. Each of the individual data tracks is randomly writeable, but a maximum attainable ADC is reduced as compared to the shingled and interlaced magnetic recording technique described below.

As used herein, a data track is randomly writeable when the data track can be individually re-written multiple times without significantly degrading data on other adjacent data tracks. An adjacent data track is “significantly degraded” if reading the data track results in a number of read errors in excess of a maximum number of errors that can be corrected by an error correction code (ECC) of the data storage device.

The magnified view 116 illustrates magnetic transitions recorded according to a shingled magnetic recording (SMR) technique. Under SMR, consecutive data tracks are written to sequentially with a write element that affects two data tracks at once. Data is written to the magnetic medium 108 in units of “data bands” where each data band includes a series of consecutive data tracks and is separated from adjacent data band(s) by one or more “guard tracks” where no data is stored. Under SMR, no data tracks are randomly writeable, but a resulting ADC is increased as compared to a maximum ADC attainable when the above-described CMR technique is employed.

The magnified view 118 illustrates magnetic transitions recorded according to an interlaced magnetic recording (IMR) technique. IMR utilizes alternating data tracks of different written track widths arranged with slightly overlapping edges so that a center-to-center distance between directly adjacent tracks (e.g., the track pitch) is uniform across the surface of the magnetic medium 108.

For example, the IMR technique shown in the magnified view 118 illustrates alternating data tracks of two different written track widths. A first series of alternating tracks (e.g., tracks 158, 160, and 162) have a wider written track width than a second series of interlaced data tracks (e.g., 164 and 166). Each wide data track of the first series (e.g., the data track 160) is written before the narrow and directly adjacent data tracks of the second series (e.g., the data tracks 164 and 166). For example, the data track 160 is written before the data tracks 164 and 165. The write to data tracks 164 and 166 overwrites outer edge portions of the data track 160; however, the data track 160 is still readable.

Because each data track of wide written width is written prior to directly adjacent data tracks of narrower width, the data tracks of the wider written width (e.g., data tracks 158, 160, and 162) are also referred to herein as “bottom tracks,” while the alternating data tracks of narrower written width (e.g., the data tracks 164 and 166) are referred to herein as “top tracks.” According to one implementation, the top tracks are of a lower linear density (e.g., kilo bytes per inch, measured in the down-track direction) than the bottom tracks. Other IMR implementations utilize interlaced data tracks having more than two different linear densities and/or written track widths.

By manipulating the linear densities of the top and bottom data tracks and also manipulating an order in which the data tracks are written, IMR systems can be tuned to exhibit a number of characteristics superior to CMR and SMR systems.

In one implementation, the IMR technique illustrated in magnified view 118 provides for system performance that is better than the best performing SMR systems. If, for example, the linear density of the top data tracks is selected to be lower than an experimentally-defined threshold, then each of the top data tracks is randomly writeable. Since random writes cannot be performed in SMR systems, an IMR system that provides for random writes exhibits superior performance characteristics as compared to SMR systems.

In another implementation, the IMR technique illustrated in magnified view 118 provides for an ADC of the magnetic medium 108 that significantly exceeds the maximum attainable ADC in SMR systems. If, for example, the linear densities of the top and bottom data tracks are pushed to upper limits (e.g., beyond that of systems described above), the data tracks may cease to be randomly writeable but provide for a very high total ADC. For example, a linear density of the top data track 166 may be selected so that the top data track

166 can be written exactly once without significantly degrading data of the adjacent (e.g., previously-written) bottom data tracks 160 and 162. Many of the implementations described herein elaborate on this type of IMR system that provides a very high ADC by pushing linear densities of the top data tracks and bottom data tracks to upper limits.

5 In one implementation, portions of incoming data addressed to a sequence of consecutive logical block addresses (LBAs) are written to a series of consecutive data tracks (e.g., the data tracks 158, 164, 160, 166, and 162). Consequently, the data may be read back from the consecutive data tracks in order of the consecutive LBA sequence without lifting the actuator arm 109 to perform a seek operation. This type of read operation is referred to herein
10 as a “sequential read.” Notably, the nature of IMR does not allow for sequential writes (e.g., writing to consecutive data tracks in a consecutive order) without significantly degrading one or more of the data tracks. In one implementation, the technology disclosed herein provides for sequential reads of stored data by directing incoming data addressed to a consecutive sequence of LBAs to a series of consecutive data tracks in a non-consecutive track order.

15 FIG. 2 illustrates data writes to a magnetic disc in an IMR system employing an example non-sequential write scheme 200 for sequential data read back. A storage device controller receives data addressed to a consecutive sequence of LBAs corresponding to a consecutive sequence of data tracks (201-207) on the magnetic disc and writes the data to the consecutive sequence of data tracks in a non-consecutive track order.

20 In FIG. 2, the data addressed to the consecutive sequence of LBAs is annotated with the letters A-G, where alphabetical order indicates a sequential LBA order. In at least one implementation, the data addressed to the consecutive sequence of LBAs is a file with the data portions sequentially ordered A-G, from beginning to end.

In the following description, the terms “first portion”, “second portion”, “third
25 portion”, etc. refer to portions of the data addressed to various consecutive LBAs. If, for example, data is addressed to LBAs 0-299, a “first portion of the data” may refer to LBAs 0-99; a “second portion of the data” may refer to LBAs 100-199; and a “third portion of the data” may refer to LBAs 200-299. In FIG. 2, the letters A-G represent portions of the data addressed to consecutive LBA sequences. For example, the letter ‘A’ may refer to data
30 addressed to LBAs 0-99; the letter ‘B’ may refer to data addressed to LBAs 100-199; the letter ‘C’ may refer to data addressed to LBAs 200-299, etc.

In writing the data to the magnetic disc, the storage device controller first writes the first and third portions of the data (A and C) to alternating data tracks 201 and 203 separated by an interlaced data track 202 (e.g., a blank data track). After writing portions A and C, the storage device controller writes a second portion (B) to the interlaced data track 202. Then, in chronological order, the storage device controller writes a fifth portion of the data (E) to a data track 205; a fourth portion of the data (D) to a data track 204, a seventh portion of the data (G) to a data track 207, and a sixth portion of the data (F) to a data track 206.

To summarize the above-described write scheme: if a most-recently written data track is referred to as index N, the next data track written has an index of N+2 if the data track of index N-1 already stores data. However, if the data track of index N-1 is blank and does not already store data, the data track of index N-1 is written next immediately following the write to the data track of index N. This trend continues so long as some portion of the data remains to be written to the magnetic disc.

After all portions of the data are written to the magnetic disc, the series of consecutive data tracks 201-207 stores the consecutive portions of the data, A-G, in a consecutive order. To read the data back from the data tracks 201-207 in the consecutive LBA order, the storage device controller can perform a sequential read of the data tracks 201-207.

As shown in FIG. 2, the odd-numbered data tracks (e.g., 201, 203, 205, 207, etc.) are bottom data tracks having a larger written track width than the top, even-numbered data tracks (e.g., 202, 204, 206, etc.). Each of the top data tracks is written after a data write to the immediately adjacent bottom data tracks. For example, the data track 204 is not written to until after data is written to the immediately adjacent data tracks 203 and 205. Each of the top data tracks (e.g., 204) overlaps and overwrites edge data of the immediately adjacent bottom data tracks (e.g., 203, 205).

In one implementation, two different linear densities are employed to write the data to the magnetic disc. Data stored in the bottom data tracks is written at a first linear density that is higher than the linear density of data stored in the top data tracks. In one implementation, linear densities selected for use in the top and bottom data tracks are low enough that the top data tracks are randomly writeable.

In another implementation, the linear densities of the top and bottom data tracks are selected to be so high that individual tracks within the series (e.g., data tracks 201-207) are not randomly writeable. For example, a linear density of the data track 204 is so high that the data portion D cannot be re-written to the data track 204 without significantly degrading data on the immediately adjacent data tracks 203 and 205. In this case, individual tracks of the data tracks 201-207 cannot be updated without re-writing each of the other data tracks. For example, updating data of the portion D may entail writing the following portions in the following order: re-writing portion A to track 201; re-writing portion C to track 203; re-writing portion B to track 202, re-writing portion E to track 205; updating data of the portion D in track 204; re-writing portion G to track 207; and re-writing portion F to track 206.

Maximizing linear densities (as described above) can provide for a large increase in total ADC as compared to implementations where each individual data track has a low enough linear density and can be written to at random. However, the use of very high linear densities implicates a write scheme under which groupings of data tracks (e.g., two or more adjacent data tracks) are treated as individual write units.

For example, the data tracks 201-207 may represent a high-density data band 220 that is re-written in its entirety each time any data track within the high-density data band 220 is updated. The high-density data band 220 may be used to store a single file (as described above), or multiple files. Different high-density data bands on the magnetic disc can be of uniform or variable size.

In one implementation, each high-density data band on the magnetic disc is separated by an adjacent high-density data band by one or more low-density data tracks. For example, the data track 208 may be a low-density data track that stores data at a lower linear density than any of the data tracks within the high-density data band 220. Each of the low-density data tracks is randomly writeable and thus can be updated without updating or re-writing data of any other data track on the magnetic disc.

Designated low-density data tracks, such as a data track 208, can be used to separate adjacent high-density data bands. In one simplistic implementation, low-density data tracks separating high-density data bands are left blank until such a time that the magnetic disc becomes so full that user data cannot be directed elsewhere. In another implementation, one or more of the low-density data tracks store a portion of a file including other data stored

in one of the high-density data bands. In yet another implementation, the low-density data tracks are used to store infrequently accessed data, such as system data or infrequently accessed user data.

FIG. 3 illustrates a magnetic disc 300 including a number of high-density data bands (e.g., high density data bands 302, 304, 306) each separated from one another by one or more low-density data tracks (e.g., low-density data tracks 308, 310, and 312). Each of the high-density data bands 302, 304, and 306 includes multiple data tracks updated as a single write unit. For example, updating data in a single data track within the high density data band 304 may entail reading all other data tracks in the high density data band 304 to a temporary memory location and then re-writing the entire band to include the updated data in the affected data track.

In one implementation, each of the high density data bands 302, 304, or 306 includes data corresponding to (e.g., originally addressed to) a consecutive LBA sequence spanning the data tracks of the associated data band so as to allow for a sequential read of the data from the magnetic disc 300 according to the consecutive LBA sequence.

To record the data on the magnetic disc 300, a storage device controller writes to data tracks of a high-density data band in a non-consecutive track order. In one implementation, the high density data bands store data according to an IMR technique where alternating data tracks have varying written track width and/or linear densities.

Each of the low density data tracks (e.g., 308, 310, and 312) are randomly writeable and store data at a lower linear density than the data tracks in any of the high density data bands. A written track width of the low density data tracks may be the same or different from one another and one or more data tracks located in a high density data band.

FIG. 4 illustrates example writes in an IMR system employing an example non-sequential write scheme 400 for sequential file read back. In one implementation, the write scheme is used to write a file to a high-density data band of IMR data. The example writes 400 illustrate options for writing portions of a first file (e.g., A1-F1) and a second file (e.g., A2-I2) to data tracks located within a same high-density data band. In FIG. 4, consecutive file portions are annotated by alphabetically arranged letters indicating a sequential ordering of the file from beginning to end.

In writing a first file to the high-density data band, a storage device controller initially writes first and third portions of the first file (A1 and C1) to alternating data tracks 401 and 404 separated by a blank, interlaced data track 402, followed by a second portion (B1) of the first file to the interlaced data track 402. Then, in chronological order, the storage device controller writes a fifth portion of the first file (E1) to a data track 405; a fourth portion of the first file (D1) to a data track 404, and a sixth (e.g., a final) portion of the first file (F1) to a data track 407. Notably, the sixth and final portion (F1) of the first file is separated from the remaining portions of the file by a blank data track 406. Thus, so long as the data tracks 406 and 408 do not include any data, the sixth portion (F1) is randomly writeable to the data track 407.

Once the storage device controller receives and stores the data of the first file (e.g., A1-F1, as shown), the controller follows one of two possible addressing schemes to address and store portions of the second file to the same high-density data band. Under a first option, (option 1), the storage device controller moves the sixth portion (F1) of the first file from the data track 407 to the adjacent data track 406 so that all consecutive portions (A1-F1) of the first file are stored in consecutive data tracks 401-406. Thereafter, the storage device controller addresses consecutive portions (A2-I2) of the second file to consecutive data tracks 407 through 415. The consecutive portions A2-I2 are written to the data tracks 407 through 415 according to the following non-consecutive track order: portion A2 to 407, portion C2 to 409, portion B2 to 408, portion E2 to 411, portion D2 to 410, portion G2 to 413, portion F2 to 412, portion I2 to 415, and portion H2 414.

Under Option 2, the storage device controller allows the first portion (A2) of the second file to be interlaced between the two final portions (E1 and F1) of the first file. The last portion (F1) of the first file is not (e.g., as in option 1) moved from the data track 407 to 406 prior to writing the portions (A2-I2) of the second file. Rather, the storage device controller addresses the first portion (A2) of the second file to the data track 406 so that it is separated from other portions of the second file by the data track 407, which stores the sixth portion (F1) of the first file.

Remaining portions of the second file are written to the high density data band using a write methodology similar to that described above. For example, portions of the second file are written according to the following non-consecutive track order: portion A2 to

405, portion C2 to 409, portion B2 to 408, portion E2 to 411, portion D2 to 410, portion G2 to 413, portion F2 to 412, portion I2 to 415, and portion H2 to 414.

FIG. 5 illustrates example operations 500 for performing a non-sequential data write that provides for sequential LBA and/or file read back in an IMR system. A receiving operation 505 receives data addressed to a consecutive sequence of LBAs. In one implementation, the consecutive data tracks are included in a high-density data band generally written to as a single write unit and separated from an adjacent high-density data band on the disc by one or more randomly writeable data tracks storing data at a lower linear density.

A first writing operation 510 writes a first consecutively-addressed portion (e.g., a portion 1) of the data to a corresponding data track having an index N. A determination operation 515 then determines whether a data track of index N-1 already stored data of the file. If the data track of index N-1 does not already store data, a writing operation 520 writes a next consecutively-addressed portion (e.g., a portion 2) of the data to the data track of index N-1. If, however, the data track N-1 does already store data, another writing operation 525 writes the next consecutively-addressed portion (e.g., the portion 2) to the data track of index N+1. A re-defining operation 530 re-defines a most recently written data track (e.g., N-1 or N+1) as "N", and the operations 515 and 520 or 525 repeat until the entire data file is written to the magnetic disc.

The embodiments of the disclosed technology described herein are implemented as logical steps in one or more computer systems. The logical operations of the presently disclosed technology are implemented (1) as a sequence of processor-implemented steps executing in one or more computer systems and (2) as interconnected machine or circuit modules within one or more computer systems. The implementation is a matter of choice, dependent on the performance requirements of the computer system implementing the disclosed technology. Accordingly, the logical operations making up the embodiments of the disclosed technology described herein are referred to variously as operations, steps, objects, or modules. Furthermore, it should be understood that logical operations may be performed in any order, adding and omitting as desired, unless explicitly claimed otherwise or a specific order is inherently necessitated by the claim language.

The above specification, examples, and data provide a complete description of

the structure and use of exemplary embodiments of the disclosed technology. Since many embodiments of the disclosed technology can be made without departing from the spirit and scope of the disclosed technology, the disclosed technology resides in the claims hereinafter appended. Furthermore, structural features of the different embodiments may be combined in yet another embodiment without departing from the recited claims.

5

Claims

WHAT IS CLAIMED IS:

1. A method comprising:
5 receiving data addressed to a consecutive sequence of logical block addresses (LBAs) corresponding to a consecutive sequence of data tracks on a storage medium; and
writing the data to the consecutive sequence of data tracks in a non-consecutive track order.
- 10 2. The method of claim 1, further comprising:
writing the data addressed to a first consecutive portion of the consecutive sequence of LBAs to a data track of index N;
writing the data addressed to a second consecutive portion of the consecutive
sequence of LBAs to a data track of index N+2; and
15 writing the data addressed to a third consecutive portion of the consecutive sequence of LBAs to a data track of index N+1.
3. The method of claim 1, further comprising:
reading the data stored in the consecutive sequence of data tracks in a single
20 sequential read operation.
4. The method of claim 1, wherein the data is a file and the method further comprises:
writing data addressed to a last portion of the file to an interlaced data track
25 positioned between two blank data tracks.
5. The method of claim 4, wherein the data is a file and the method further comprises:
writing portions of another file to the two blank data tracks.
- 30 6. The method of claim 4, wherein the data is a file and the method further

comprises:

moving the data of the last portion of the file to one of the two blank data tracks prior to writing data in another one of the two blank data tracks.

5 7. The method of claim 1, wherein at least two adjacent tracks of the consecutive data tracks have different written track widths.

8. The method of claim 1, wherein the consecutive data tracks include alternating data tracks storing data at alternating linear densities.

10

9. A storage device controller configured to:
write data addressed to a consecutive sequence of LBAs to a series of consecutive data
tracks on a storage medium in a non-consecutive track order; and
5 read the data from the consecutive data tracks in a single sequential read operation.

10. The storage device controller of claim 9, wherein writing the data further
comprises:

writing data to a first data track and subsequently writing data to a second data track
10 separated from the first data track by one interlaced data track;
writing data to the interlaced data track after writing data to the first data track and the
second data track.

11. The storage device controller of claim 9, wherein the consecutive data tracks
15 include at least two adjacent tracks have different written track widths.

12. The storage device controller of claim 9, wherein the consecutive data tracks
include alternating data tracks storing data at alternating linear densities.

20

13. Apparatus comprising:

a storage device controller configured receive data addressed to a consecutive sequence of logical block addresses (LBAs) corresponding to a consecutive sequence of data tracks on a storage medium and write the data to the consecutive sequence of data tracks in a non-consecutive track order.

14. The apparatus claim 13, wherein the storage device is further configured to:

write the data addressed to a first consecutive portion of the consecutive sequence of LBAs to a data track of index N;

write the data addressed to a second consecutive portion of the consecutive sequence of LBAs to a data track of index N+2 if a data track of index N-1 already stores data; and

if the data track of index N-1 does not already store data, write the data addressed to the second consecutive portion of the consecutive sequence of LBAs to a data track of index N-1.

15. The apparatus of claim 13, wherein the storage device is further configured to:

read the data stored in the consecutive sequence of data tracks in a single sequential read operation.

16. The apparatus of claim 13, wherein the data is a file and the storage device is further configured to:

write data addressed to a last portion of the file to an interlaced data track directly adjacent to two blank data tracks.

17. The apparatus of claim 16, wherein the storage device is further configured to write

a first consecutive portion of another file to the two blank data tracks directly adjacent to the interlaced data track.


18. The apparatus of claim 16, wherein the storage device is further configured to

move the data of the last portion of the file to one of the two blank data tracks prior to writing data in another one of the two blank data tracks.

19. The apparatus of claim 13, wherein the consecutive data tracks include at least
5 two adjacent tracks having different written track widths.

20. The apparatus of claim 13, wherein the consecutive data tracks form a portion of a high density data band are separated from an adjacent high density data band by at least one low density data track.

10






FIG. 2




FIG. 3

+




FIG. 4

5 / 5




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/014950**A. CLASSIFICATION OF SUBJECT MATTER****G11B 5/09(2006.01)i, G11B 5/596(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G11B 5/09; G11B 5/02; G11B 7/24; G11B 5/55; G11B 5/584; G06F 12/00; G11B 5/596

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & keywords: logical block address (LBA), data track, consecutive sequence, non-consecutive track order, read, write, linear density, width, and similar terms.**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5949607 A (JOHANNES J. W. KALFS et al.) 07 September 1999 See column 3, line 46 - column 4, line 68; claims 1-3; and figures 1-3.	1-3, 9-10, 13, 15
Y		7-8, 11-12, 19-20
A		4-6, 14, 16-18
Y	US 2010-0027406 A1 (MICHAEL KRAUSE et al.) 04 February 2010 See paragraphs [0027], [0031]-[0035]; claim 1; and figures 2a, 5a-5b.	7, 11, 19
Y	US 4646168 A (YOSHIKI SONOBE et al.) 24 February 1987 See column 6, line 51 - column 7, line 2; and figures 7-9.	8, 12, 20
A	US 2007-0050593 A1 (HUAYUAN CHEN et al.) 01 March 2007 See paragraphs [0029]-[0048]; claims 1, 7; and figures 2-5.	1-20
A	US 2007-0174582 A1 (TIMOTHY R. FELDMAN) 26 July 2007 See paragraphs [0028]-[0039], [0065]-[0092]; and figures 2-3, 8-12.	1-20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

03 May 2016 (03.05.2016)

Date of mailing of the international search report

03 May 2016 (03.05.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

BYUN, Sung Cheal

Telephone No. +82-42-481-8262



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/014950

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 05949607 A	07/09/1999	CN 1168189 A EP 0803120 A1 EP 0803120 B1 JP 10-512705 A KR 10-1998-0701266 A WO 97-17700 A1	17/12/1997 30/01/2002 31/07/2002 02/12/1998 15/05/1998 15/05/1997
US 2010-0027406 A1	04/02/2010	AU 2007-331564 A1 CN 101553873 A EP 2092522 A1 JP 2010-514074 A KR 10-2009-0088408 A TW 200832393 A WO 2008-071653 A1 ZA 200903561A	19/06/2008 07/10/2009 26/08/2009 30/04/2010 19/08/2009 01/08/2008 19/06/2008 25/08/2010
US 04646168 A	24/02/1987	EP 0105502 A1 EP 0105502 B1 JP 1730626 C JP 4015521 B JP 59-063005 A	18/04/1984 08/04/1987 29/01/1993 18/03/1992 10/04/1984
US 2007-0050593 A1	01/03/2007	None	
US 2007-0174582 A1	26/07/2007	US 8756399 B2	17/06/2014