Shingled Magnetic Recording Technologies for Large-Capacity Hard Disk Drives

● SHIMOMURA Kazuhito

The information infrastructure society has been undergoing further major changes in recent years. Whereas most digital data were formerly created by individuals using PCs, the progress of social networks has led to the generation of huge volumes of data on a daily basis from smartphones and other portable terminals. Furthermore, the data produced by various sensors in the Internet of Things (IoT) has also been rapidly increasing. The total volume of data generated in the world in 2020 is projected to reach 44 zettabytes (ZB; 1 ZB = 10^{21} bytes), 10 times the volume in 2013. Larger capacity hard disk drives (HDDs) are therefore required in the market to store these huge volumes of data.

To address this demand, Toshiba has been developing shingled magnetic recording (SMR) technologies. SMR HDD record data in a narrower track pitch than conventional HDDs by overwriting a part of the adjacent track that has been recorded, like shingling a roof, using a read/write head of conventional size. We are working to realize new technologies for SMR that strengthen its advantages and compensate for its weak points.

1. Introduction

The volume of data generated globally is increasing at an explosive rate and is predicted to grow to 44 ZB by 2020, roughly 10 times the volume in 2013[(1)]. To keep all that data in storage, there is a strong demand for further increases in HDD capacity. Thus far, the storage industry has increased recording density of HDDs at an annual rate of over 15% through the improvement of key components such as read/write heads, magnetic disks, and system-on-chips (SoCs) and through the application of new technologies. Toshiba has continually developed innovative HDD components and utilized them to offer higher-capacity HDDs in a timely manner. However, component development takes a lot of time and cost. Since the volume of data produced by information infrastructure grows so rapidly, it is becoming difficult to meet the storage needs quickly.

In these circumstances, SMR is attracting much attention as a means of increasing the recording density of HDDs without relying on costly component development. SMR is suitable for use in low-cost, high-capacity HDDs because it allows the recording density to be increased using conventional components simply by modifying the track placement and the recording unit.

This report provides an overview of the SMR technology being developed by us and its future prospects, comparing it with the conventional magnetic recording (CMR).

2. SMR overview

An HDD moves a read/write head to an appropriate track of a proper disk to read and write data (**Figure 1**).

Because CMR HDDs can randomly read and write data from/to arbitrary sectors on a disk, they deliver excellent random-access performance. For this reason, CMR HDDs are widely used not only in PCs but also for online storage applications. Each track on a CMR HDD is separated from its adjacent tracks by gaps called "guard bands" so that a random-write access to a track will not affect data on the adjacent tracks. The width of

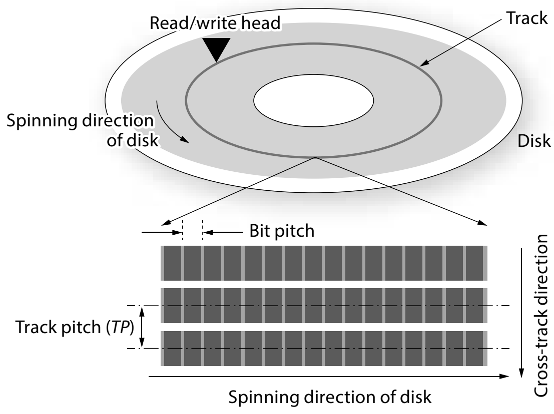

Figure 1. Outline of head and media of HDD.
HDDs position a read/write head over a spinning disk to store and retrieve data.

the write element of a head, *W*, is designed to be narrower than the track pitch, *TP* (**Figure 2**).

The recording capacity of a CMR HDD is a function of a bits-per-inch (BPI) density in the down-track (circumferential) direction and a tracks-per-inch (TPI) density in the cross-track direction. In order to increase recording density, the width of a write element and the *TP* of a disk must be reduced (**Figure 3**). For this purpose, the magnetic characteristics of both of them need further improvement.

However, at present, the width of a write element is as narrow as approximately 50 nm, and it is becoming difficult to further reduce its width without affecting recording capability. HDD manufacturers have been researching innovative technologies to break this limit, but it requires much time and cost to develop key components such as heads and disks.

Here comes SMR, which writes a new track, *N+1*, overlapping part of the previously written track, *N*. This means a narrower *TP* and thus allows for a higher TPI without reducing the width of the head. Because the tracks partially overlap similar to roof shingles, this approach is called shingled magnetic recording (**Figure 4**). Although both CMR and SMR have a restriction for *TP* due to the width of the read element, SMR provides far greater design freedom than CMR does.

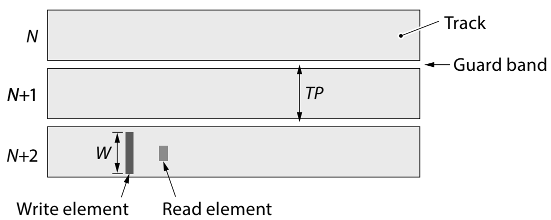

Figure 2. Schematic diagram of data tracks, write element, and read element of conventional magnetic recording (CMR) HDD.
CMR has guard bands between adjacent tracks. The width of the write element is narrower than the track width.

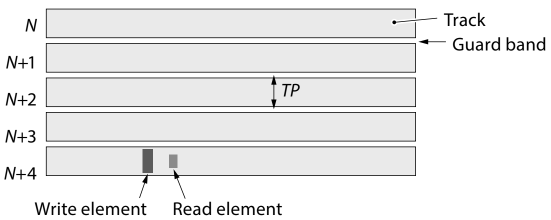

Figure 3. Schematic diagram of data tracks, write element, and read element of high tracks per inch (TPI) CMR HDD.
In order to increase TPI using CMR, the width of the write element must be reduced together with the track width.

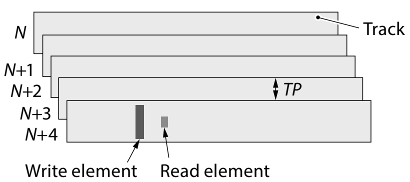

Figure 4. Schematic diagram of data tracks, write element, and read element of high-TPI SMR HDD.
SMR makes it possible to increase TPI to achieve higher recording density regardless of the width of write element.

3. Data recording

CMR temporarily buffers data from a host into memory. Regardless of whether writes are addressed to random or sequential locations on a disk, the controller firmware reorders the writes in such a manner as to minimize the overall access time and then performs writes to the disk (**Figure 5**)

In SMR, data are sequentially written to shingled tracks just as the roof of a house is shingled. Therefore, SMR does not have the random-access performance of CMR. To resolve this issue, SMR HDDs need a special writing procedure. Specifically, an SMR disk provides a media cache (MC) and bands of shingled tracks. Bands are defined as a unit consisting of shingled tracks that are sequentially written. There is a guard band between contiguous bands.

Upon a write request from a host, an SMR HDD temporarily writes the buffered data to a disk area called an MC. The addresses of data in the MC are recalculated so that the data can be written to sequential locations in shingled bands. To rewrite previously written data, an SMR HDD reads old data from a band and merges it with new data in the MC and then writes it back to a new band (**Figure 6**).

In order to improve write performance, we have also developed an algorithm that directly writes data to bands without involving the MC when a write request for sequential data from a host is larger than a certain size.

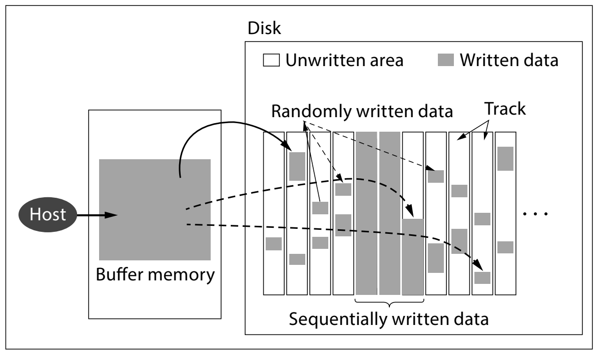

Figure 5. Data writing procedure in CMR HDD.
Data blocks held in a buffer memory are reordered to minimize the overall access time before they are randomly written to a disk.

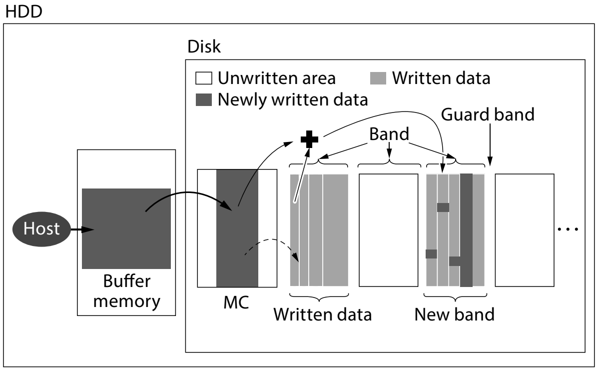

Figure 6. Data writing procedure in SMR HDD.
In response to write requests from a host, the SMR HDD receives the write data with a buffer memory and copies it to the MC temporarily. Then, the write data is written in a shingle pattern to new bands.

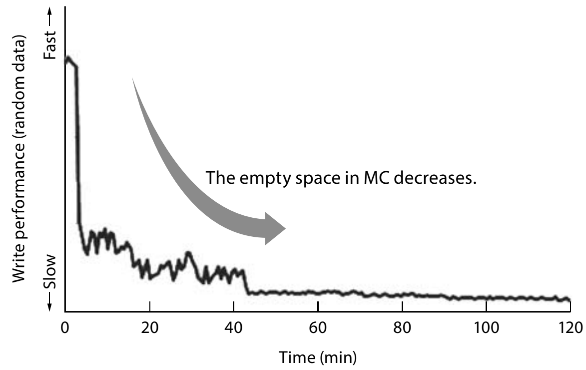

Figure 8. Random data writing performance of SMR HDD.
The random-write performance of the SMR HDD varies, depending on the amount of free space available in the MC.

4. Write performance evaluation

We have evaluated the data transfer rate of our SMR HDD prototype. Since SMR HDDs have a significant difference in performance between sequential and random writes, both of them have been measured. SMR HDDs have a read performance equivalent to that of CMR HDDs.

Figure 7 shows the performance of SMR and CMR HDDs when data were written sequentially from the outermost track of a disk to the innermost track. As an SMR HDD is basically designed for sequential writes, its sequential-write performance is on a par with that of a CMR HDD.

The SMR HDD, which is not well suited for random writes, improves its random-write performance by using an MC. However, if many random-write requests come in from a host, the MC is filled over time, degrading the random-write performance (**Figure 8**).

Because of its characteristics of recording method, an SMR HDD is suitable for sequential-read/write-intensive applications.

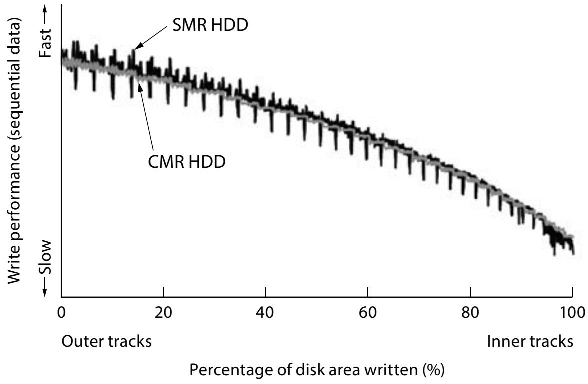

Figure 7. Comparison of sequential-write performance of CMR and SMR HDDs.
The SMR HDD has a sequential-write performance equivalent to that of the CMR HDD.

5. Combining high TPI and reliability

In the event of excessive vibration or impact during writing, the magnetic head could go off track, causing data on an adjacent track to be corrupted in the worst case. In order to reduce the risk of data corruption in a high-TPI SMR disk, the *TP* must be carefully determined.

An external disturbance could cause the write element to move sideways to an adjacent written track and erase the data in it. If this lateral displacement exceeds a threshold, D_0, the position control firmware stops writing to prevent the corruption of data on the adjacent track (**Figure 9**). An SMR HDD must be designed to satisfy the relationship betweenD_0 and TP in equation (1) below:

$$TP = 2 \times D_0 + O_{VR} + (R + E) \quad (1)$$

where, O_{VR} is a distance margin required so that the read quality of an adjacent track will not be adversely affected even if an external disturbance has caused a lateral displacement of the head exceeding D_0, erasing part of the data on the adjacent track. R is the width of the read element, and E is the width of magnetic data

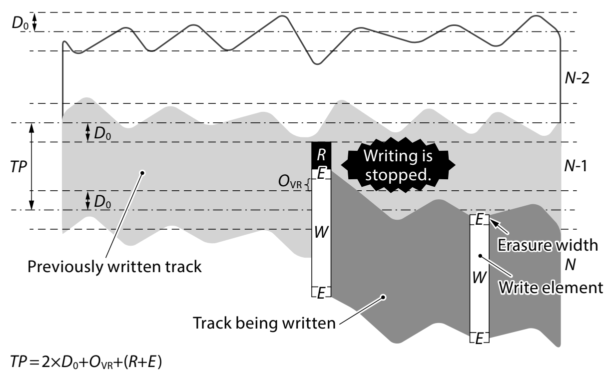

Figure 9. Track pitch design with emphasis on reliability.
The track pitch is carefully determined to prevent accidental data erasure in the event of an external disturbance.

erasure that occurs at the edge of the write element.

D_0 must be reduced to decrease *TP* and thereby achieve a high TPI. On the other hand, too small a D_0 value would cause an SMR HDD to stop writing data frequently, degrading write performance. Our SMR HDD incorporates innovative solutions for making full use of the advantage of the SMR HDD in order to achieve even higher TPI with a minimal effect on write performance.

6. Types of firmware and industry standardization

Although, depending on data characteristics, SMR HDDs do not always exhibit performance comparable to that of CMR HDDs, SMR will certainly provide the capacity growth spurt. Studies on the utilization of SMR HDDs have begun not only at individual HDD manufacturers but also across the entire industry sector involved in information infrastructure, including storage system manufacturers, large-scale data centers, and other users.

Furthermore, we have been engaged in the standardization of SMR together with other companies. The International Committee for Information Technology Standards (INCITS) has the T10 Technical Committee responsible for the Serial Attached SCSI(*1) (SAS) standards and the T13 Technical Committee responsible for the Serial Advanced Technology Attachment (SATA) standards, each of which has proposed SMR-specific command standards such as Zone Block Command (ZBC) and Zone-device ATA Command (ZAC) standards.

There are three types of firmware for SMR: drive-managed (DM), host-aware (HA), and host-managed (HM). Drive-managed SMR needs to run with an existing filesystem. While a host can handle a drive-managed SMR HDD in the same manner as a CMR HDD, its performance varies depending on the host's access conditions and is hard to predict. In contrast, host-managed SMR allows a host to manage the operation of HDDs in a manner suitable for the characteristics of SMR and thus delivers stable performance. For hosts supporting host-managed SMR, new technologies are being developed to make it possible for filesystems and kernels to run SMR HDDs in a desired manner. Host-aware SMR is not purely driver- or host-managed, but has characteristics of each (**Table 1**).

Table 1. SMR firmware categories

Category	Characteristics
Drive-managed (DM)	The host can handle a CMR HDD in the same manner as an SMR HDD. However, the HDD performance significantly varies according to the host's access conditions. SMR HDDs are downward compatible with CMR HDDs.
Host-aware (HA)	The host can optimize write operations using commands specifically designed for SMR HDDs. Even in the event of an SMR command is under violation, the SMR HDD accepts the command and performs a possible operation. The HDD performance sometimes varies significantly according to the host's access conditions. Host-aware SMR HDDs are downward compatible with CMR HDDs.
Host-managed (HM)	The host uses SMR-specific commands to optimize write operations. The SMR HDD rejects an SMR command that violates to the specification. The write performance can be estimated to some extent. SMR HDDs are not downward compatible with CMR HDDs.

(*1) SCSI: Small Computer System Interface

7. Conclusion

More and more SMR HDDs will be available as higher-capacity HDDs become necessary to meet customer and application needs.

CMR HDDs have been dominant for the past few decades. As SMR HDDs gain wider uptake, changes will take place in the design of such key components as magnetic heads and recording media. The host filesystem will also be optimized to obtain the best performance from SMR HDDs.

SMR HDDs will bring about a paradigm shift, involving system manufacturers and users, and find widespread use in information infrastructure across which a huge volume of data flows.

We will incorporate the SMR technology into next-generation high-capacity HDDs.

Reference

(1) Turner, V. et al. "The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things." IDC Website. Accessed May 25, 2015. http://idcdocserv.com/1678.

SHIMOMURA Kazuhito

Chief Specialist. HDD Products Engineering Department 1, Storage Products Design & Production Division, Semiconductor & Storage Products Company. He is engaged in the design of HDD products.

Note:
This publication may include Toshiba's original estimation, and/or data based on Toshiba's evaluation.