

# Hybrid Drive with Large Capacity and High-Speed Performance

●Mine BUDIMAN ●Eric DUNN ●Rick EHRlich

The hybrid drive, which combines bit cost advantage similar to rotating media with high performance close to that of NAND flash memory, has recently been attracting attention as a new storage device. The combination of a hard disk drive (HDD) and NAND flash memory makes it possible to deliver close to a solid-state drive (SSD) user experience with high-speed performance and ample storage capacity through improvement of the cache algorithm. Toshiba has developed a new cache algorithm that places frequently used data into the NAND flash memory, resulting in a hybrid drive providing close to SSD performance. As the newly developed cache algorithm is independent of the operating system, high-performance notebook PCs can be realized by replacing the HDD of existing notebook PCs with a hybrid drive as well as installing hybrid drives in new products.

## > 1. Introduction

As the Internet becomes more popular, large numbers of videos and large amounts of other types of data are being handled more frequently than before. With respect to the storage devices installed in PCs (e.g., HDDs), in addition to being able to store a massive amount of data, data access performance (e.g., OS data readout performance for quickly booting PCs as well as data read/write performance while PCs are in operation) has become very important. HDDs are still the major player in the storage device arena and are used in numerous fields, though there are limitations on performance improvement potentials. For this reason, SSDs, which feature high-speed data access, are attracting attention and the SSD market is expanding. However, HDDs still have the advantage in cost per gigabyte. Given these circumstances, Hybrid drives, which combine HDDs with NAND memory, have come into their own as storage devices capable of realizing both high capacities and high performance.

Toshiba is able to provide HDDs and SSDs, which can be used in various fields, to both enterprises and consumers entirely on its own. Now, by adopting a cache algorithm to add NAND memory to the cache memory in an HDD as a secondary cache, Toshiba has developed a Hybrid drive with high performance close to that of SSDs. In this article, we describe the Hybrid drive's architecture and give an overview of the technologies that we applied to develop the hardware and firmware.

## > 2. Toshiba's Approach to Hybrid Drives

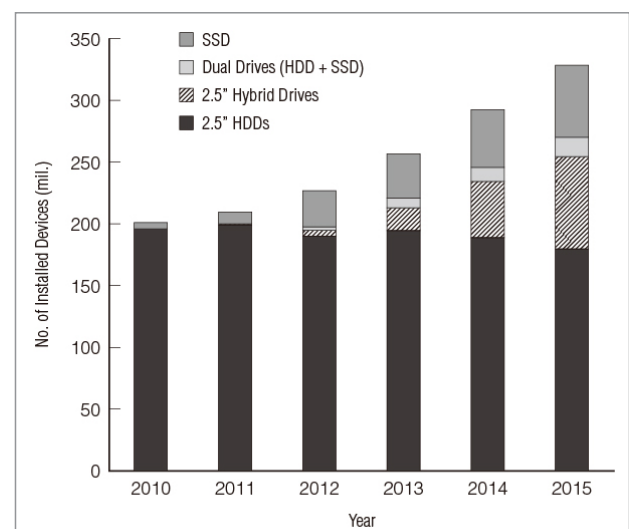
Toshiba has always led the development of state-of-the-art technologies for NAND memory. At the same time, we are promoting "Total Storage Innovation" for our Information and Communication Technology (ICT) society, which will grow in the future, by effectively combining all types of major storage devices (e.g., HDDs and SSDs). By promoting synergy derived from SSD technologies, NAND memory, and HDD technologies of all types, including 3.5-

inch HDDs, we are globally expanding the reach of storage devices, a cornerstone of our total storage innovation, into the consumer and enterprise markets.

As the amount of data increases in tandem with the popularity of the Internet and cloud computing, the capacity of storage devices such as the HDDs used in PCs is also increasing. Still, more and more SSDs are being employed in notebook PCs because of their advantages in terms of increased data access speeds, decreased PC start-up times, compact size, and lower power consumption.

**Figure 1** shows a forecast of what types of built-in storage devices will be used in notebook PCs. Full-fledged adoption of hybrid devices, which simultaneously realize both high capacities and high-speed data access performance while keeping bit costs low, is predicted to start in earnest in 2012 and account for about 25% of the market in 2015.

Given such market trends, the Hybrid drive that we



Source: IDC, "WW2012 HDD Market Update"<sup>(1)</sup>

**Figure 1 Forecast Share of Built-in Storage Products for Notebook PCs —**

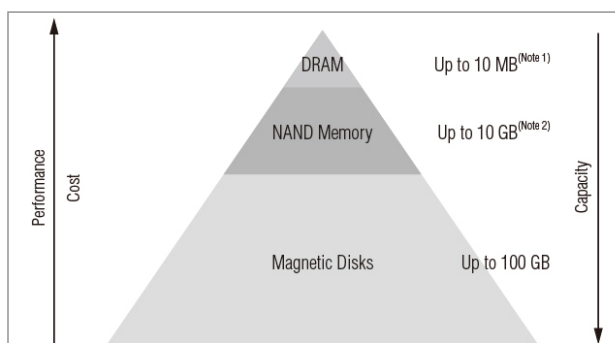
The number of Hybrid drives built-in to notebook PCs is predicted to increase greatly after 2012.

have developed realizes both a high capacity and high performance by combining our traditional NAND memory technologies, NAND memory handling technologies (instrumental in SSD technologies), and HDD product development technologies.

### > 3. Architecture

The basic structure of our Hybrid drive is shown in **Figure 2**. For the cache memory, the HDD uses low-capacity, high-speed DRAM (volatile memory) and a large-capacity magnetic disk (non-volatile memory). NAND memory (nonvolatile memory), which has medium capacity, read/write speeds, and bit costs compared to DRAM and magnetic disks, is added as a secondary cache in this HDD.

Our newly developed cache algorithm dynamically studies data access patterns and stores frequently accessed data in non-volatile NAND memory, allowing access performance to be improved even if no cache hits occur in DRAM during, for example, booting the PC. Thus, the algorithm can realize performance close to that of SSDs. PC systems recognize the Hybrid drive just as they do conventional HDDs and SSDs. Therefore, no additional software is needed and existing HDDs or SSDs can be replaced directly with this drive.



**Figure 2 Basic Configuration of Hybrid Drives** —


In Hybrid drives, read/write speeds can be improved by making use of 3-tiered storage system consisting of DRAM, NAND memory, and a magnetic disk.

### > 4. Hardware Development

In hybrid drive development, drives must be brought to market as quickly as possible while keeping hardware costs low. Thus, we decided to use components that were already in use in existing HDDs as much as possible.

To this end, we developed a NAND controller for controlling data read/write to and from NAND memory. This controller serves as a bridge chip positioned between the System on a Chip (SoC) and NAND chip. The NAND controller and Synchronous DRAM (SDRAM) are connected in parallel to the Double Data Rate 2 (DDR2) SoC interface; different address spaces are allocated to the NAND controller and SDRAM to allow the existing HDD and Hybrid drive to share the SoC.

In our Hybrid drive, however, the above-mentioned NAND controller and NAND memory additionally needed to be mounted on the existing HDD circuit board, yet we worried that there was an insufficient amount of space on the board's mounting area. Changing the area or shape of the board may lead to a redesign of the HDD. To avoid this, the necessary amount of mounting space was reduced by adopting small-sized DDR2 (rather than DDR1) for SDRAM. Furthermore, we secured the wiring area by changing the conventional 4-layer printed circuit board to a 6-layer board. As a result, components could be mounted on a board of the same size as that of a standard HDD as shown in **Figure 3**, and thus could be shared with the existing HDD case.



**Figure 3 EPrinted Circuit Board Assembly for Hybrid Drives** —

The NAND memory and NAND controller are mounted on a board the same size as that of HDDs by using small-sized SDRAM and increasing the number of layers in the multilayer board.

However, it was impossible to mount more than one NAND memory; hence, we had to select a single NAND memory with the required performance. The candidate NAND memories and a performance comparison of sequential write operations are shown in **Table 1**. With respect to NAND memory, single level cell (SLC) technology has a higher cost of capacity per unit than multi level cell (MLC) technology, but also much higher performance. If the number of dies (NAND chips) is increased, the drive's performance will be improved considerably because of parallel processing. Since this Hybrid drive could not include more than one NAND memory, we performed a cost/performance trade-off analysis and selected the 8-GB SLC (which uses four 2-GB dies). In typical HDDs, random access performance is restricted by seek-times and rotational delays; in most cases, command reordering techniques are used to minimize these speed decreases, but HDDs are still limited to 200-300 Input/Output per second (IOPS). NAND memory performance is restricted by the access speed to the NAND memory, but this can be improved by increasing the number of channels and the number of dies to make the drive operate in parallel. With such techniques, it is possible to enhance the random read/write performance of NAND memory up to several thousand IOPS, an order of magnitude better than that of a standard HDD.

**Table 1 Comparison of Sequential Write Performance of NAND Memories**

Memory Type	Sequential Write Performance* (MB/s)		
	2 dies	4 dies	8 dies
SLC 32 nm (2 GB/die)	–	137	237
MLC 24 nm (8 GB/die)	44	88	176
SLC 24 nm (4 GB/die)	160	248	248
eMLC 24 nm (8 GB/die)	44	88	176

eMLC: Enterprise MLC

\* Including ERASE before PROG, but not garbage collection (GC), wear leveling (WL), or Write Amplification (WA).

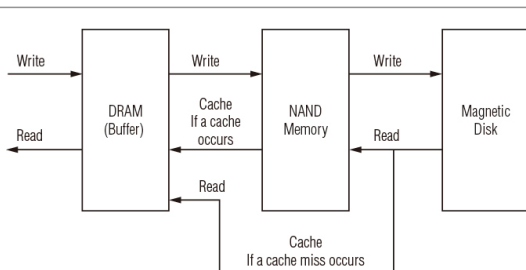
All configurations assume Multi Page Cache Read/Program.

## > 5. Software Development

The capacity of the NAND memory mounted on the Hybrid drive is 8 GB, approximately 1% of the capacity of a 1 TB<sup>(Note 3)</sup> HDD. Effective use of NAND memory as a cache requires careful consideration of what data to place in the NAND memory because of the limitations imposed by its capacity.

The Hybrid drive's design concept is shown in **Figure 4**. The Hybrid drive consists of the following three storage tiers: SDRAM, NAND memory, and magnetic disk, each of which has different read/write speeds. When the system requests that data be written, the Hybrid drive will write the data to the NAND memory before writing the data to the magnetic disk (data will be written to disk later). When the system requests that data be read from a specific address, the drive will check whether the data for the specified address is stored in NAND memory. If the data is present in NAND memory, the drive will quickly send the data to the system. If not, the drive will read the data from the magnetic disk and send the data to the system, then copy the data to NAND memory. Because of this algorithm, data that was recently written to the Hybrid drive or read out from the drive will be contained in NAND memory. If the system requests that data be read from the same address again, the drive can return such data to the system very quickly because the required data is stored in NAND memory.

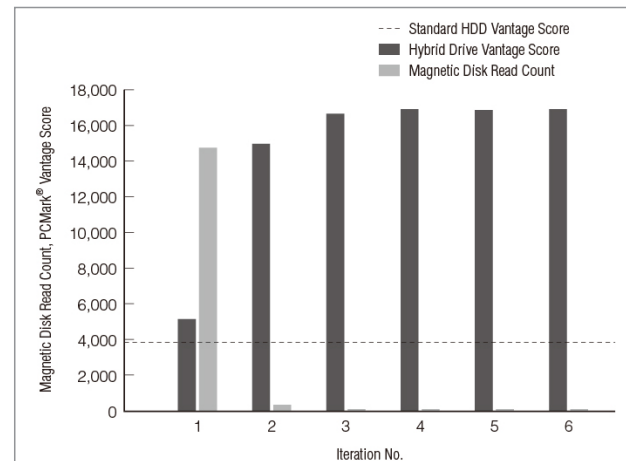
Using a Hybrid drive prototype with a 16-GB NAND memory, we conducted a series of six iterations of a performance benchmarking tool (PCMarkR<sup>(Note 4)</sup> Vantage HDD Suite). As for the benchmark results, the relationships between benchmark scores and magnetic disk read counts are shown in **Figure 5**. The



**Figure 4 Design Concept of Hybrid Drives —**

Data is controlled so that recently written and read out data is stored in NAND memory, allowing read/write speeds to be improved.

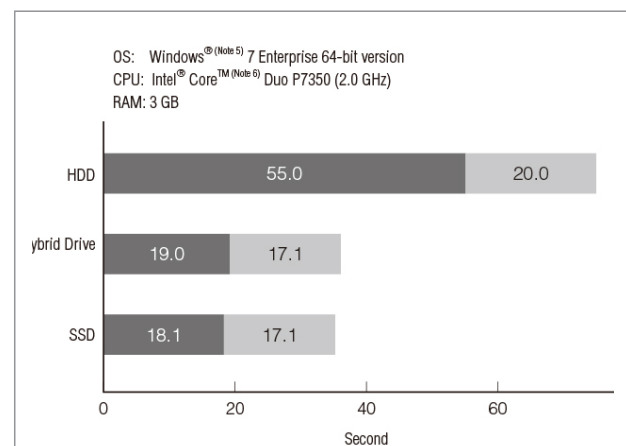
Hybrid drive must read data from the magnetic disk if the data is not stored in NAND memory. After the second measurement, the magnetic disk read count decreases dramatically, since most of the data is read out from NAND memory. In the first measurement, the benchmark score is only slightly better than that of a conventional HDD, but the scores for subsequent benchmarks exceed those for the first measurement by a factor of three. These scores are the result of studying the system's data access patterns and caching the necessary data into NAND memory.



**Figure 5 Benchmark Scores and Magnetic Disk Read Counts —**

Data storage locations are optimized during the first iteration, so repeating the benchmark reveals improved read/write speeds, resulting in an increase in score to about three times that of a conventional HDD.

**Figure 6** shows an example time comparison in which an HDD, Hybrid drive, and SSD were installed in three otherwise identical PCs. The time to load the same application was then measured. The load time required for the Hybrid device is much less than that of the conventional HDD and nearly as fast as that of the SSD. Based on these evaluation results, we can confirm the effectiveness of our Hybrid drive.



Application start-up times may differ depending on system configuration.

**Figure 6 Comparison of Application Load Times —**

On the Hybrid drive, applications started faster than they did on the HDD and nearly as fast as they did on the SSD.

## > 6. Conclusion

To further improve Hybrid drive performance, a specification for standardization of hints, which are data provided by PCs that give hints as to whether to keep data in NAND memory, is now being discussed at the Serial ATA International Organization (SATA-IO), the standard-setting organization for Serial Advanced Technology Attachment (SATA). In the NAND memory field, efforts to reduce costs by integrating components and miniaturizing processes are now underway.

Toshiba will continue to make intensive efforts to considerably improve the cache algorithm that intelligently uses the 3-tiered structure of SDRAM, NAND memory, and a magnetic disk, and also to develop second-generation Hybrid drives with even higher performance.

(Note 1), (Note 2), and (Note 3):

Definition of capacity: Toshiba defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1GB =  $2^{30}$  = 1,073,741,824 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, such as Microsoft Operating System and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

(Note 4) PCMark is a trademark of Futuremark Corporation.

(Note 5) Windows is a trademark of Microsoft Corporation in the United States and other countries.

(Note 6) Intel and Intel Core are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

---

### References

(1) IDC. WW2012 HDD Market Update. IDC. 2012



**Mine BUDIMAN**  
Toshiba America Electronic  
Components, Inc.



**Eric DUNN**  
Toshiba America Electronic  
Components, Inc.



**Rick EHRLOCK,  
D.SC.**  
Toshiba America Electronic  
Components, Inc.