Smooth Stream™ Technology in Hitachi Hard Disk Drives

HITACHI
Inspire the Next

Daniel Colegrove
Interface Standards
daniel.colegrove@hitachigst.com

Introduction

Hard disk drives developed for use in personal computers are now being used in audio/video devices such as MP3 players and personal video recorders. AV systems play and record streams of data. If the stream is interrupted, the interrupt may result in missing audio or video information.

Hitachi helped to develop standard AV stream support on hard disk drives and in August 2004, introduced the Deskstar® 7K400, the industry's first hard drive to incorporate this important capability. The Hitachi Travelstar 4K120 now follows the Deskstar 7K400 in the adoption of Hitachi Smooth Stream Technology, to become the first 2.5-inch drive optimized for audio/video applications.

In AV applications, there is no advantage to streaming data faster than the AV device requires. Data must be delivered consistently at the required rate. As a result, it may be better to have some small segment of incorrect data delivered in the stream than to have a long delay. Short delays may result in the loss of only a few pixels[1]. A long delay in the data stream would result in the loss of a larger block of data, which would be noticeable to a viewer. This new Streaming Command Set was developed for ATA drives, which allows AV products to change drive behavior to meet AV system requirements.

Disk drive behavior

Hard disk drives are rotating media devices with circular tracks which store data. They are equipped with a fast actuator that moves the read/write heads from track to track. As a result, drives operate with greater variable time delays between commands and data transfers, with pauses between data transfers. When a drive receives a command to read data, the actuator moves to the required track. The constantly rotating disk then moves the data area under the head and reading begins. If the data spans more than one track, a delay or skew is incurred, during which time the heads move from the end of one track to the start of the adjacent track.

Occasionally additional delays are required for error recovery. The delays can occur in multiples of a revolution time. The drive can detect an error in the data read using the error correcting code at the end of each sector. If the data is not correct, the drive will automatically wait for another disk rotation and re-read the data. The procedure for recovering from a data error can take several revolutions, and enact techniques such as micro-stepping off track. In the recovery process for computer systems where data quality is of the utmost importance, the time delay is acceptable, however in AV systems, a long delay to get perfect data is not the preferred behavior.

AV optimizations

The AV Streaming Command Set includes features that allow an AV system to derive maximum performance from the disk drive. The Configure Stream command enables the host system to tell the disk drive how many simultaneous streams will be used and whether they are read or write streams. This information triggers the disk drive to optimize its buffer management to accommodate the specific operation requirements.

AV operations usually contain several read commands that request data in sequential order between seeks. Drives are configured with a feature called Read Look Ahead that keeps the drive reading ahead, even after execution of the current command has completed. The data is retained in the drive buffer to provide a prompt response to a system request if the ensuing command is indeed sequential. The read stream in AV optimized drives contains a Non-Sequential, or NS bit. The NS bit informs the drive that the next command it sends will not be data that sequentially follows the current command. The next sequential data is not loaded by the drive into the buffer since the AV system has determined that the data will not be accessed.

1 Video information is stored as pixels. Pixel stands for picture element. It corresponds to the smallest thing that can be drawn on a computer screen. The amount of storage required for a pixel varies with the associated technology. Thus, simple black and white images require only one bit - 0 or 1. More complex technology such as color images require more storage. For example, 4 bits can store up to 16 colors; 16 bits can store 65,000 colors. HDTV requires even more storage for each pixel, since more details (colors and shades) are associated with the technology of HDTV.

Read continuous and write continuous

The Streaming Read/Write command support the Read Continuous (RC) bit or Write Continuous (WC) bit. When a stream command is issued with the RC or WC bit set, the drive will transfer the best data it can during error conditions, while completing the command within the designated time limit. If read data is transferred that is not completely accurate, an entry is placed in the Read Stream Error Log.

Logging the data performs two functions: 1) allows the AV stream to locate the defective data in a stream by accessing the log, and 2) enables the avoidance of lengthy error recovery procedures when an Error bit is flagged. The later allows the error bit to be suppressed during RC operations and yet provides a method by which the defective data may be retrieved.

WC operations are also supported. Although the probability of detecting an error is much lower than during write operations, WC is useful in that it allows a command to complete within the specified time limit by skipping inaccessible sectors and recording the write errors into the Write Stream Error Log.

AV normal behavior

The AV Streaming Command Set includes a set of tables that describes normal drive performance. Drives are formatted with multiple zones, each zone having more sectors per track progressing from the inner zone to the outer zone. These tables describe the zone map and other specific drive information including average seek times for read and write operations on a track to track basis and other critical parameters. With this information, an AV system can estimate the time it will take to reach a track and the time required to read or write the data.

Using the streaming performance data, an AV device can determine the average time it will take to read or write data—but data from disk drives arrives in bursts. There is no data transfer during seeking and skew, followed by a high data rate when the head is passing over the data. The pulsing nature of data transfers from the disk drive is smoothed out by system buffer memory. Fortunately, the data rate of disk drives exceeds the data rate required for AV data streams. This allows AV systems to support multiple AV streams simultaneously.

AV error recovery procedure controls

In computer applications, drives may enter into time consuming error recovery procedures to maximize the likelihood of returning correct data. In AV systems, data that is delayed is as useless as bad data. If the buffer runs empty, visible defects or a loss of signal will occur. The Streaming Command Set provides controls that set a time limit to interrupt error recovery. Using the performance data and the buffer capacity, AV systems determine how much time can be allotted to each command. By setting a time limit for each command, the AV system can manage the time allowed for drive error recovery, thereby minimizing disruptions in the delivery of the AV stream. The drive handles the errors based on the conditions prescribed by the host AV system.

It is important to allow drives some error recovery time. The recommended calculation is a minimum of the time to complete three full revolutions plus the data transfer time. This error control handling practice ensures that any error can be handled quickly without compromising the quality of the visual or audio experience.

Stream error handling

An additional optimization defined by the Streaming Command Set is the Handle Stream Error or HSE bit. AV applications may need data that must be correct, such as computer system data for software or file information The HSE bit allows an AV system to perform full error recovery procedures in incremental steps. Error recovery is enacted between the processing of streaming commands. If a command times out, the system can retry the command with the HSE bit set to 1. This causes the command to try again, starting at the level of error recovery reached during the previous try.

For example, a system may be running several streams of AV data and also need computer quality data. If an error occurs during the computer data read and the operation exceeds the allotted time limit, the data is not returned. The system then services the AV streams, ensures that the buffers are kept full, and then retries the computer data command with HSE set to 1. The drive starts where it left off processing the data command and continues until the data is accurately recovered or the time limit is reached. The system can continue this process until the computer data is fully recovered, without causing an interruption in the AV streams.

Summary

The new Streaming Command Set included in the ATA/ATAPI-7 standard[2] provides a comprehensive set of tools for disk drives used in audio/video applications. The ATA Streaming command Set includes:

- Streaming Performance Log—describes normal drive performance
- Error Recovery Time Limits—prevents long delays for error recovery handling
- Continuous Read/Write Controls—used to return/use partially correct data
- Streaming Error Logs—reduces time delays otherwise incurred in system error processing
- Configure Stream Command—allows optimum drive buffer management for AV streams
- Delayed Sector Log—documents relocated sectors that requires additional read/write processing time
- Handle Stream Error—enables full error recovery for computer data during AV streaming operations

These features and functions have been incorporated into Hitachi's Smooth Stream Technology. Together with their industry-leading performance as well as cool and quiet operation, Hitachi drives that support Smooth Stream are also optimized for consumer applications that require reliable storage for audio and video files.

2 Draft ATA standards are available at www.t13.org.
Completed standards are available at www.incits.org.

Hitachi Global Storage Technologies
3403 Yerba Buena Road
San Jose, CA 95135 USA

Produced in the United States 11/07.