DMA and Hard Disk Drives— Keys to Mastering the ATA/IDE Bus

HITACHI
Inspire the Next

Dr. William F. Heybruck
Senior Storage Engineer
bill.heybruck@hitachigst.com

Introduction

Today's IDE/ATA interface hard disk drives support both Direct Memory Access (or DMA) and PIO mode (or Programmed I/O) bus mastering protocols. DMA is a generic term referring to the direct transfer of data from a peripheral device to or from a memory location on the host system. In general, DMA is a high speed data transfer to or from the hard disk that allows the host to move the data directly to and from memory with very few state changes. By default, all IDE/ATA devices support PIO mode 0. This is the basic mode of registers and commands. PIO supports five different modes, however all PIO data transfers take place relative to the level of read and write strobe lines. Proper implementation of DMA in the host design will enable more efficient Central Processor Unit (CPU) utilization and faster overall system performance.

Identifying and setting the DMA mode

DMA enables faster data transfers than PIO mode as CPU intervention is not required for the transfer of data between the system and the device, which for the purpose of this paper is assumed to be a hard disk drive. Two versions of DMA are frequently used in today's IDE/ATA-based systems. The more traditional DMA version, or Multiword DMA, uses the read and write strobe lines to clock the transfer of data across the interface. The strobe line is level sensitive; data is usually transferred when the line is active low. The relatively new version called Ultra DMA was designed to improve the overall efficiency of the IDE/ATA interface. Ultra DMA enables data to be transferred on both the rising and falling edges of the interface clock, separates the strobes from the host and device, and embeds Cyclical Redundancy Checking (or CRC) in the data stream to improve integrity of data transferred at the high speeds possible in Ultra DMA mode.

The ATA specification provided fields in the Identify Device command to indicate what DMA version is supported by the hard disk drive and other DMA devices on the bus, as well as their DMA settings. Words 63 and 88 in the Identify Device response is used for this purpose, triggering the drive (and devices) to identify their capabilities and to send that information back to the host.

In addition to the two versions of DMA, various throughput modes are defined under both Multiword and Ultra DMA. Multiword DMA supports three modes, Mode 0, 1 and 2, and Ultra DMA six, or Mode 0 through 5. An IDE/ATA hard drive can be set to operate in any of these modes (if supported) through the Set Features command sequence.

To direct the hard drive (or "device") to use DMA, the Set Features command is issued with 03h in the Features Register; this will select the Transfer Mode Feature on the target device. The Sector Count register of the device will contain the appropriate mode information. Depending on which DMA version is designed on the target device, there are differences in how the device setting is actually handled. In Multiword DMA devices, the Sector Count register is set with 2Xh where X designates the Mode (0, 1 or 2). However in Ultra DMA devices, the Sector Count register is set with 4Xh, where X reflects the mode (0, 1, 2, 3, 4 or 5). The host can poll the drive to confirm its mode setting by reading the Identify Device response words 63 and 88, as detailed previously. Following this query, the host is expected to interact with the device in the same mode indicated during the Identify Device handshake.

Multiword DMA

Multimode DMA is the older and hence slower bus mastering version of DMA supported by IDE/ATA devices. Although still widely used, Multimode DMA does not take advantage of the maximum interface transfer speed capabilities on today's hard disk drives. Three modes of operation are available under Multiword DMA which differ by the timing of the data transfers, or cycle times (see Figure 1). The cycle times for Modes 0, 1, and 2, are strictly defined by the ATAPI Specification (ATA/ATAPI-7) to ensure compatibility across the many hard drives and devices using the IDE/ATA interface under the Multiword DMA bus mastering protocol.

Multiword DMA Mode	Cycle Time (nanoseconds)	Maximum Transfer Rate (MB/second)	Governing Standard
Mode 0	480	4.2	ATA
Mode 1	150	13.3	ATA-2
Mode 2	120	16.7	ATA-2

Figure 1: Data transfer modes defined by Multiword DMA

To avoid conflict on the interface bus and ensure the integrity of the data transfer, it is important that the Identify Device response word 63 be queried and the minimum t0 time (word 65) selected. The time value from the host must correspond exactly with the mode setting on the target device(s) in order to avoid compatibility issues on the bus.

In Multimode DMA operation, the host will issue a READ DMA or WRITE DMA command following the same DMA bus preparation sequence discussed above, and will send or receive the proper number of words of data to/from the device without intervention of the host system CPU. This usually entails prior configuration of the host DMA controller to ensure that the user data is placed in the proper memory location within the host.

The IDE/ATA device in Multimode DMA mode, will respond to a READ DMA command by reading the first sector of the count (if more than one), and placing the Read data in an onboard cache. The device will then assert DMARQ (or DMA REQUEST), indicating it is ready to initiate the Read data transfer. When the host responds with DMACK (or DMA ACKNOWLEDGE), the host will proceed to read the data starting from the falling edge of DIOR (DMA I/O READ) and read subsequent bytes or words from the device on falling edges of DIOR. The device may only remove the DMARQ signal when DIOR is inactive and all data has been transferred. Completion of the Read data transfer may be followed by the host removing DMACK, and a waiting period until the device is ready to transfer the next requested data segment. DMARQ is always removed upon completion of the data transfer and only when the device is ready to receive another command.

In the execution of a WRITE DMA command, the device will assert DMARQ when ready to receive data from the host. The host responds with DMACK, and begins the transfer of Write data by clocking the data into the device via DIOW (or DMA I/O WRITE). At this time, the device may remove DMARQ, signaling the host to hold the Write data transfer. The host cannot assert DIOW until the DMARQ signal is set to active by the device. Once the Write data transfer is complete, the device will remove DMARQ, which then signals the host to remove DMACK.

In the event of a data error occurrence during the transfer of Read or Write data, the Multiword DMA device has the option of either receiving all the data before posting the error STATUS response, or asserting INTRQ and removing DMARQ the moment the data error is detected. The host responds to the INTRQ signal by aborting the DMA transfer by removing DMACK, followed by a read STATUS from the device.

Ultra DMA

Much higher throughput rates are possible under Ultra DMA due to double transition clocking since the data is transferred on both edges of the interface clock transitions, called DSTROBE and HSTROBE. Like Multiword DMA, the transfer rates can be set to the processing capabilities of the host and the IDE/ATA device(s) attached to the bus. The host still uses the Set Features command, described above, to define the transfer rate mode to be used by the device(s). The host then prepares itself for the DMA transfer by configuring its DMA controller or subroutine parameters; this must be completed before the host is able to send any READ DMA or WRITE DMA command to the device.

Ultra DMA supports six data transfer modes, Modes 0 through 5, as shown in Figure 2, although only Modes 2, 4 and 5 are used with hard disk drives. Ultra DMA mode cycle times range from 240ns to 50ns, with Mode 0 operating at the slowest rate, and Mode 5 the fastest. The resulting transfer rates across the interface run from 20 to 100 megabytes per second. Hitachi's third generation Microdrive®, the model 3K6, supports Ultra DMA Mode 2 at the 33.3 MB/second data transfer rate, while the 1.8-inch Travelstar® C4K60 drive series for PC and consumer electronics applications supports Ultra DMA Mode 5 at 100 MB/second. The speed setting of the device is usually indicated by reference to the interface as "Ultra ATA/xx", where "xx" signifies the MB/second transfer rate.

Ultra DMA Mode	Cycle Time (nanoseconds)	Maximum Transfer Rate (MB/second)	Governing Standard
Mode 0	240	16.7	ATA/ATAPI-4
Mode 1	160	25.0	ATA/ATAPI-4
Mode 2	120	33.3	ATA/ATAPI-4
Mode 3	90	44.4	ATA/ATAPI-5
Mode 4	60	66.7	ATA/ATAPI-5
Mode 5	40	100.0	ATA/ATAPI-6
Mode 6	30	133	ATA/ATAPI-7

Figure 2. Data transfer modes defined by Ultra DMA

The processing of READ data commands is detailed in Section 9.13.1 of the ATA Specification, which may be better understood when referencing the following overview. This is intended as a theory of operation description and is not geared to be used in actual design work. The Ultra DMA device responds to a DMA READ command by asserting DMARQ, as in the sequence explanation given in the Multiword DMA section. However, the host response to the DMARQ command is considerably different in the Ultra DMA protocol. Here the host will respond by asserting STOP and removing HDMARDY (or HOST DMA READY), the signal which governs the two CS lines (CS0 and CS1) and three DA lines (DA0-DA2) in Ultra DMA mode. This can be done in any order before the host asserts DMACK.

DMACK then remains active throughout the corresponding DMA burst of Read or Write data. Data is clocked to the host (READ operation) on the edges of DSTROBE and from the host (WRITE operation) on the edges of HSTROBE. Unlike Multiword DMA, the data between the device and host under Ultra DMA operation is transferred on both edges of the strobe . However, the strobe edges are only valid if DMARQ, DMACK, STOP, and HDMARDY are in the proper state. Additionally, the strobes always begin in the high state; the first edge of DSTROBE is thus always High to Low.

As long as all the appropriate cycle times are followed, data may be continually transferred to and from the host according to the subsequent transitions in the data lines and STROBE edges. When all data is transferred, the device negates DMARQ and waits for the host negation of DMACK. The data bus is released at this time; the host will assert STOP, and STROBE will be placed in the negated (low) state. During the STOP cycle, this specific STROBE edge is ignored by the host, since the host is expected to negate HDMARDY soon after DMARQ is negated by the device.

CRC was introduced into the Ultra DMA protocol to improve the integrity of data being transferred at these faster rates. In order to indicate to the device that CRC is to be embedded into the data stream, the host must negate DMACK; the checking process begins only after DMARQ is negated. This triggers the sending device to enact its CRC algorithm, which then reads the CRC and compares its results against the information originally sent. In the event of a mismatch, the data is resent by the device. Following completion of the CRC compare, the device sets the status register with the appropriate signal indicating that either the transfer was ‘good’ or ‘no good’. The device responds by releasing STROBE to the inactive state (high), which then completes the data transfer from the device to the host system.

The processing of WRITE data commands is similar to the READ operation described above with the following exceptions:

- The DMA Write command is issued to the device from the host,
- DDMARDY is negated by the device after the host asserts DMACK,
- The host drives the data onto the interface bus, and
- The host changes the state of HSTROBE (first transition is always High to Low), no sooner than a delay from DDMARDY, and only when the data is active. (Delay timing details are provided in the ATAPI specification.)

The host continues to strobe data to the device on the HSTROBE edges until the Ultra DMA burst is either terminated or paused by the host. In the later case, the host may interrupt the data transfer simply by not generating HSTROBE edges; conversely, the host may recommence the data transfer at any time by generating a HSTROBE edge.

To terminate data transfers in Ultra DMA operation, the host ceases generation of the HSTROBE edges and asserts STOP. The device responds by negating DMARQ and waits should the host negate the STOP signal posting. If the STOP remains valid, the device negates DDMARDY (DEVICE DMA READY). If an odd number of words were transferred, the host will assert HSTROBE at this time. The host then places CRC on the data bus and negates DMACK. This triggers the device to compare the CRC information and update the status register accordingly. Upon completion of the CRC process, the device will post the DDMARDY signal. The host may then release other control lines. This termination process is used regardless of the number of words transferred, or whether the DMA command was complete. The status register is updated accordingly.

Summary

Many consumer electronics and handheld devices that incorporate memory storage devices such as IDE/ATA hard disk drives are configured with the PIO bus mastering protocol. Faster transfer rates can be derived from the IDE/ATA drive using either Multiword or Ultra DMA, with Ultra DMA offering the fastest transfer rates yet defined under the ATA Specification. Although the initial set-up process of either Multimode or Ultra DMA can be more complex, the long-term advantages of faster data transfers and more efficient usage of the disk drive can prove advantageous to the host system not only in higher data processing speeds, but also in enhanced battery utilization. Multiword and Ultra DMA were designed to be downward compatible with PIO devices.

Detailed ATA interface information, including PIO and DMA modes are available in the ATA Specification referenced below. Support and software utilities to ensure the seamless integration of these different devices are readily available through the Hitachi GST web site or any of the Hitachi Service Centers.

References and additional reading

1) ATA Specification, ANSI ATA/ATAPI-7, T13 document number ANSI NCITS 361-2002, (with 6 Packet Interface).

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

Produced in the United States 11/07.