

To: X3T10 Membership
CC: ATA Technology Users
From: Jim McGrath, Quantum
Subject: ATA Electrical Issues when using 2 ports

Enclosed is a DRAFT copy of a Quantum white paper on local bus electrical issues that covers the newly discovered issue involving the use of 2 IDE ports with the new generation of IDE bridge chips (or core logic). As has been documented previously, these bridge chips multiplex some of the IDE lines (Data and often various control lines) to minimize chip pin count. Unfortunately, this results in data buses with up to 4 devices and lengths up to 40 inches, making it impossible to use either ATA port due to unsatisfactory signal quality.

Fortunately, this problem can be resolved through some relatively simple external logic on the host motherboard. A discussion of this issue in more technical detail and proposed solution(s) is enclosed in this paper. Although it is ultimately the motherboard maker's responsibility to implement these changes, all ATA technology suppliers and consumers should be aware of this issue and its resolution. This will allow system companies to better specify motherboard requirements at the beginning of design, and allow easier identification of problems if faulty motherboard solutions are implemented.

Quantum would appreciate widespread distribution of this cover memo and the enclosed white paper. Quantum waives its copyright rights as long as all 12 pages are distributed and Quantum acknowledged as the sources. Specifically, people are requested to examine and comment on this document in writing. The preferred technique is to mark up the paper and fax it to:

Jim McGrath
c/o Mary Turco
Quantum
408-894-4990 (fax)

Email is also possible, in which cases mail should be sent to BOTH:

JMCGRATH@QNTM.COM
MTURCO@QNTM.COM

In addition to technical contributions, anyone interested in being included as endorsing the contents of the paper should also contact us. We will so note that in the paper. Our hope is to have the industry act with one voice on this topic to minimize confusion.

We would like to issue a final paper within 2 weeks. Anyone contacting us and leaving a US mail address will receive a master copy and the right to freely reproduce and distribute it.

ATA line termination issues on systems with high speed (Local Bus) interfaces

By the staff of
Systems Engineering
Quantum Corporation
3/9/95

Introduction

New PC motherboards and interface cards supporting local bus interfaces (both VLB and PCI) are running disk drives up to the maximum transfer rate specified by ATA mode 3 (11.1 MB/s max) and mode 4 (16.67 MB/s max). The ATA Bus is currently specified as an unterminated bus. This lack of terminations causes ringing on the signals traveling between the host and the disk drive (see figure 1). Previous bus speeds were sufficiently slow that this ringing was not a problem, but higher bus speeds - ATA modes 3 and 4, are now being exercised by local bus interfaces. This could lead to a system failure if the signals are not properly terminated.

The ATA interface specifies an 18-inch, 40-conductor ribbon cable as the means for connecting disk drives to the host interface. Up to two drives may be connected to the cable. The most common configuration is for the host to be at one end of the cable with the first drive located 12 inches away and the second drive located 6 inches further. For single drive configurations the disk drive is usually located at the end of the cable (18 inches from host). It is important to note that the ATA spec does not require this configuration. The host may be located at the center of the cable, or the drives may be spaced anywhere along the cable length (e.g. one drive 1 inch from host, the other 18 inches from host). These alternate connection schemes have been used by system vendors that are tight on space or have unusual physical configurations. ATA compatibility must be maintained for all valid configurations.

The ATA cable is not specified as having any terminations. With slow TTL edge rates this has not been an issue in the past. The advent of CMOS VLSI bridge chips has increased the edge speeds on the bus, causing ringing and crosstalk problems. A long-time rule-of-thumb for analog designers has been that if the propagation delay of the cable exceeds one-quarter of the rise time then cable termination must be used. In the case of the ATA bus, the worst case propagation delay of the cable is approximately 4 ns, so rise times of less than 16 ns could cause problems. Many local bus to ATA bridge chips available today have edge rates of 1 to 2 ns, despite the ATA-2 spec requirement of 5 ns.

RF circuits solve the issues of signal reflections and ringing by driving the cable with a source impedance equal to the characteristic impedance of the cable, and terminating the far end of the cable with this same impedance. ECL

Preliminary

circuit designers solve the problem by terminating only at the receiving end. Unfortunately, both of these cases assume a single source and a single termination. The SCSI interface designers understood this issue and specified that the physical devices at each end of the cable must be terminated. This has the disadvantage that the user must know the physical configuration and must install (or remove) the appropriate terminations.

The ATA interface designers did not want to burden the user with the issue of terminations, so the spec intentionally left out any requirements governing the location of the drives or the host adapter. Now that we acknowledge that some kind of termination is required, this places an additional constraint on system and disk drive designers.

One of the innovations in ATA technology that some suppliers are offering system customers is dual port - using the primary and secondary IDE disk controller addresses to allow the connection of up to 4 devices in a system, two per cable. This expands the number of devices (e.g. 2 IDE, 1 CD-ROM, 1 tape drive), and in practice makes ATAPI CD-ROM usage possible.

Meanwhile, the electrical specifications for ATA are all based on the assumption of 1 or 2 devices per cable, and cables limited to 18 inches in length. Violating these constraints, either by adding more devices to the cable or lengthening the cable, would create problems due to transmission line effects induced.

As conceived by the ATA and ATAPI standards, the new dual port feature work with these electrical limitations since two distinct cables are still used, each obeying the old electrical specifications.

ATA Bus Termination

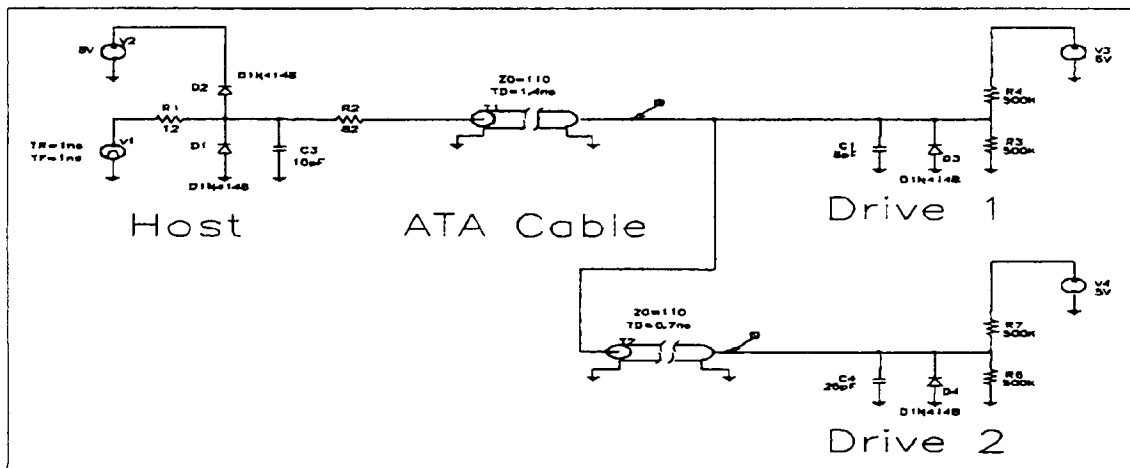


Figure 1 - Terminated host with two different disk drives

Preliminary

The ribbon cable commonly used for the ATA bus has a characteristic impedance of approximately 110 ohms. Theoretically, the optimum solution would be to drive the cable with a 110 ohm source and insure that the device at the end of the cable was terminated in 110 ohms. One problem with this solutions is that when a 110 ohm cable is driven with a 110 ohm source, the voltage is divided by two. Another problem is how to terminate the end of the cable and not have a termination in the middle when the user decides to install a second drive.

There is no technically "proper" solution to this problem. Redefining the ATA bus to use terminations would lose backward-compatibility which is one of the strong features of the ATA bus. The only recourse available to system and drive designers is partial termination at both the host and the drive. No one solution at a single point will work under all circumstances.




Figure 2 – Waveform at drive inputs

One of the common misconceptions is that limiting the slew rate of the source to 5 ns will fix the problem. The rule-of-thumb used by analog designers tells us that we need 16 ns edges to truly eliminate the ringing problem. The controlled slew rate helps, but does not solve the problem by itself.

Another misconception is that source termination is sufficient to eliminate the ringing problem. This is not the case. Simple simulations with a host and a single disk drive may give this false impression, but a more complete simulation with two drives reveals the problem. Figure 1 shows a SPICE model with a source, 12 inches of cable, a disk drive, 6 inches of cable, and another disk drive. The host is terminated with 82 ohms, but the drives are unterminated. Figure 2 shows the waveforms at the inputs to each of the two drives resulting

Preliminary

4

from the model of Figure 1. The signal at drive 2 is fine, but the signal at drive 1 shows a significant glitch at half of the source voltage.

This is easily explained by transmission line theory. The incident wave begins by traveling down the cable. The magnitude of the incident wave is the source voltage reduced by the voltage divider of the source impedance and the cable impedance ($110/(82 + 110)$). For now we will assume that this is about half of the applied voltage. The incident wave reaches the first drive, and then passes on down to the second drive. The second drive is at the end of the cable, so the incident wave reflects at this point. At the point of reflection the incident and reflected waves are added with no time delay, so the full signal amplitude appears. But the reflected wave still has to travel back to the source. As it does, it passes the first disk drive. Now the first disk drive sees the sum of the incident and reflected waves, and the voltage steps up to the applied voltage value.

The problem occurs because the incident and reflected waves occur at different times at the first disk drive. The greater the delay between the incident and reflected waves, the greater the problem. The reflected wave eventually arrives at the host where it is mostly dissipated in the source termination. Note that if the drives are more than 6 inches apart, the problem gets worse.

ATA Bus Timing

Terminating the ATA bus does not come for free. The impact of partial terminations at the host and the drive is to increase propagation delays throughout the system. The ATA document specifies that timing is referenced to the input pins of the disk drive. This means that most of the timing issues must be addressed by systems manufacturers and bridge chip designers.

The most significant timing issue is the propagation delay of the cable. This needs to be accounted for in the host-side timing. The SPICE model in Figure 6 shows a host with very fast rise times driving a cable with worst case loads. Two disk drives are assumed with the maximum allowed capacitive loading of 25 pF (ATA-2 working draft, Rev. 2f, section 4.5, table 3). The simulation results are shown in Figure 7. It can be proven that the period of the ringing is four times the propagation delay of the cable. This simulation shows that the worst-case cable propagation delay is 4.2 ns. Note that this is twice the value that would be calculated by assuming an 18-inch cable with a propagation velocity of 70%. The additional delay is due to the presence of the capacitive loads on the cable. This is important to system designers who must take into account worst-case cable delay when specifying the bridge chip timing.

From the host point of view all of the ATA timings must be corrected by the propagation delay of the cable to insure that the timing is correct at the input pins of the furthest drive. Figure 8 shows a typical result using the read cycle data setup time as an example. The ATA spec says that the setup time at the drive must be 20 ns (PIO Mode 3). This means that the host must allow for a setup time of 11.6 ns ($20 - 2 \times 4.2$). Note that this 11.6 ns time is referenced to the pins of the ATA interface at the host. If the host has a series partial

Preliminary

termination resistor then the bridge chip must include additional timing margin to account for the RC delay of that resistor. Quantum simulations show that the additional cable delay added by a series termination resistor at the host is approximately 5.6 ns for a 33 ohm resistor, and 11.0 ns for an 82 ohm resistor (two drive load, 25 pF at each drive, 18-inch cable). The extra delay of the 82 ohm resistor is one of the reasons that Quantum recommends 33 ohm series resistors at the host.

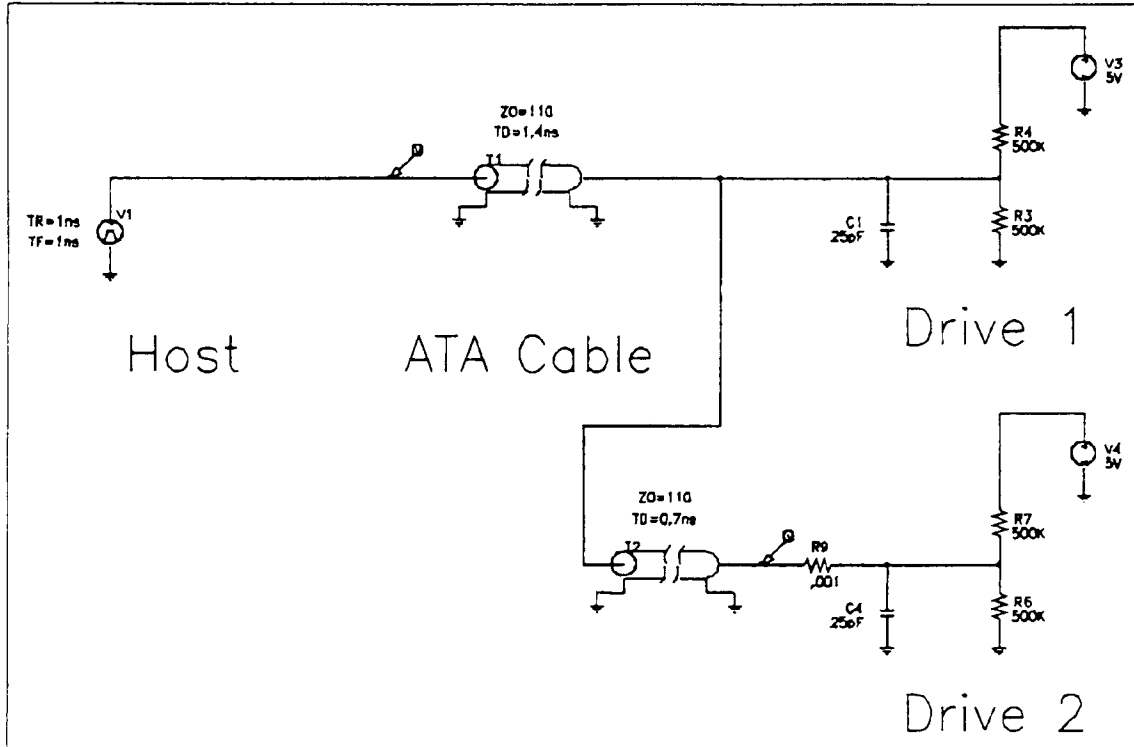


Figure 6 - SPICE model of ATA cable with worst case loads

Preliminary




Figure 7 – Simulation of unterminated ATA cable with worst case loads




Figure 8 – Host data setup time during a read cycle

One of the common mistakes made by chip designers is to use a lumped capacitance model for simulating the delay of the output cell. For ATA simulations this sometimes consists of adding the maximum capacitance allowed for the host and the drives (3 x 25 pF) to an estimated capacitance value for the cable (25 pF). The simulation is then performed with 100 pF capacitance on the output. This is not an accurate measurement of the timing. The proper simulation is to use an output capacitance for the motherboard, a host end termination resistor, and transmission lines to the drives (Figure 9). To illustrate how these are different, suppose that the output cell turns out to be 2 ns too slow. The designer (using a 100 pF model) might be inclined to increase the drive of the output devices. With enough drive into a purely capacitive load another 2 ns can be removed, bringing the output cell back into spec. But look at what happens

201

Preliminary

with the transmission line model. Increasing the drive of the output cell does not decrease the length of the cable nor increase the speed of signal propagation in the cable. The 2 ns required cannot be achieved by increasing the output drive. The time must come from somewhere else in the internal circuitry. Increasing the output drive only increases the edge speed, making the ringing worse at the drive end. (In reality some time is gained with the faster edge speed, but not nearly as much as would be predicted with a simple capacitive load model.)

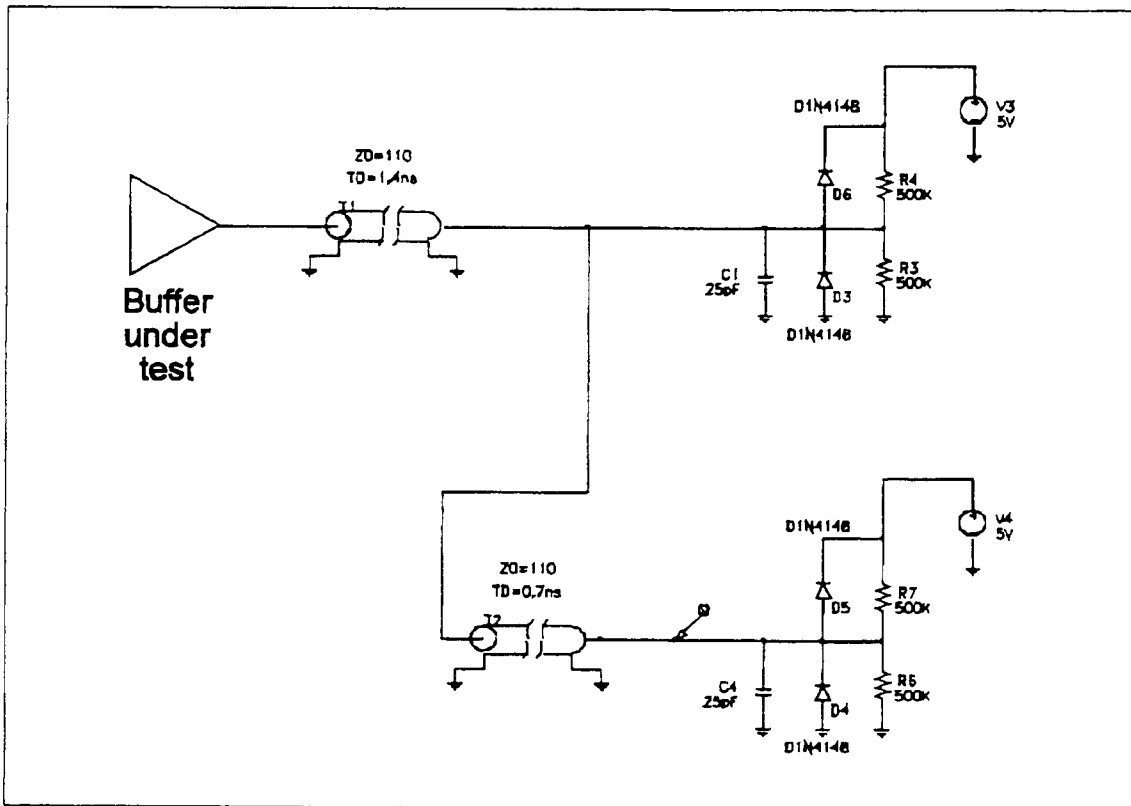


Figure 9 – Standard cable load for measuring ATA signal timing

Propagation delay times at the host (and at the drive) must be measured to the ATA standard of 0.8 V for high to low transitions and 2.0 V for low to high transitions. Many IC manufacturers measure to the "typical" switch point for TTL of 1.4 V. This is not appropriate for the ATA interface since virtually every chip manufacturer (both host and drive end) has included hysteresis for noise immunity. (To an extent, this hysteresis is the only thing that has kept ATA working as long as it has.) Since both the hysteresis window and hysteresis offset of a given receiver will move with process, voltage, and temperature, the only guaranteed switch points are the TTL high and low values (0.8 V and 2.0 V).

Disk drive manufacturers need to insure that any partial termination circuits they implement present an effective capacitance of 25 pF or less. What is an effective capacitance? From a practical point of view, any circuit is valid provided it does not increase the propagation delay of a worst-case cable. This

225

Preliminary

8

is because systems manufacturers are counting on a certain cable delay in their design. The easiest way to answer the question of acceptability is to do a simulation with the proposed termination and measure the delay. The simulation needs to include all possible drive combinations including generic drives with ordinary 25 pF loading. The simulations should also include read cycle timing with the worst-case load on the data lines. The difference here is that the drive manufacturer gets to make their measurements at the output of the drive, not at the far end of the cable.

Unfortunately the delay of a given termination circuit at the drive end also depends on the termination used at the host end. There is an interaction between these two terminations. Quantum simulations show that the incremental cable delay due to the 82 ohm termination at the drive circuit is 1.4 ns for a 33 ohm host termination and 5.0 ns for an 82 ohm host termination. This delay is in addition to the fundamental RC delay introduced by a series termination resistor at the host.

Disk drives internally need to account for any delay that their termination circuitry inserts. This delay must be added to the timing specifications of the interface controller chip such that the ATA timing specs are still met at the input pins to the drive.

Dual Port ATA Systems

New PCI-IDE bridge chips and core system logic chip sets have begun to support both Fast IDE and multiple ports. However, many implementations grossly violate ATA and will result in major system integration problems. Even though devices may not be at fault, these bad system implementations will cost the industry time, effort, and possibly sales if our customer's shipments are delayed.

In the interest of saving pins for these system chips, the chip makers have often multiplexed the pins between the ports. Some control lines obviously cannot be multiplexed (e.g. IRQ). But clever design can allow some control lines to be multiplexed (e.g. IOR and IOW), and the data lines can always be logically multiplexed. Unfortunately, they cannot be physically multiplexed without a proper care.

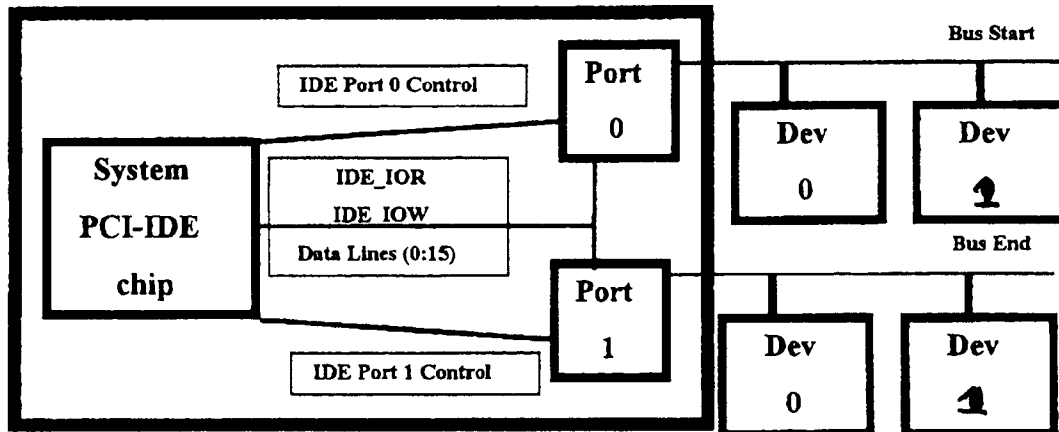
The restriction mentioned earlier on number of devices on a cable and length of a cable should really refer to "bus" not "cable." In the old, simple ATA world the bus was the cable (with some stubbing). But with 2 ports it is possible to create a system with 2 cables but one bus by multiplexing signals between the ports.

The figure below illustrates this case. With IOR, IOW, and data multiplexed, the physical bus for these lines begins at the point marked "bus starts" and ends at the point marked "bus ends." It contains the two ATA cables and the traces on the system board connecting the cables to one another. Stubs are present at each device (4 stubs) and one in the middle of the bus to connect to the system

220

Preliminary

chip. If the individual cables are 18 inches, then the total bus is 36-40 inches long (depending on the length of the traces on the system board) with 5 device loads. Under these circumstances the transmission line effects would generate errors for any device operating at Fast ATA speeds, and perhaps even at slower speeds.



This is one of the classic issues that is a system design problem impacting devices on the ATA cable. At a minimum technical white papers could be prepared to spread knowledge of this issue by assisting the system integrators in finding a clean solution to this potential problem such as buffering the signals on the motherboard. This will add a prop delay which should be considered in the overall system timing.

Another major problem is the implementation of DASP line on some motherboards. We have seen some cases that the primary and secondary port DASP- lines are wire-ORed together. This in itself is a problem. With these lines connected, if a master and a slave are connected to the primary port, while a master is connected to the secondary port, the secondary master will always incorrectly see a slave present on its port during a bootup. This could cause the master on the secondary port not to boot in the time required. This problem would be for all drive manufacturers who follow the ATA CAM specifications not only Quantum.

During bootup with a master and a slave drive, the slave drive asserts the DASP- line to inform the master that it is present. The slave drive needs to assert DASP- within 400 msecs after the RESET- line is released. The slave drive also has 30 seconds after the RESET- line is released to assert PDIAG- to indicate the slave has finished its diagnostics. A master drive will look for the DASP- line being asserted within the 400 msec window after the RESET- line is released. If the master drive sees the DASP- line asserted it assumes that a

Preliminary

10

slave drive is present and waits for up to 30 seconds for the slave to assert PDIAG- indicating the slave has completed its diagnostics.

Some CD ROM devices hold the DASP- line asserted for 86 msec after the RESET- line is released. Since the DASP- lines for the secondary port and the primary port are connected together, any master following the ATA specifications that is connected to the other port will incorrectly determine that a slave is present on its port.

Guidelines for system designers

There are not many guidelines in the ATA spec addressing the transmission line effects on the ATA cable. Here are few considerations

- * Make sure that control signals are properly driven.
Do not use any value less than 1K ohm for pull up resistors on ATA bus open collector signals such as IORDY and IOCS16. (ATA spec).
- * If series resistor termination is used, measure and analyze the other parameters on the same line in the system (load capacitance, distance from the source on the cable, propagation delay, rise and fall times of the buffer and the resistor value at the other end). Make sure that signal timing of the system is not compromised.
- * Series resistors of 22 to 47 ohm should be used on all of the control out lines closest to the ATA connector on the motherboard. For dual port implementation, one resistor for each port and each control out signal on the cable should be used.
- * Series resistors of 22 to 47 ohm should be used on all of the data lines closest to the ATA connector on the motherboard. For dual port implementation, one resistor for each port and each data line on the cable should be used.

Preliminary

11

- * Series resistors of 22 to 47 ohm should be used on all of the address lines closest to the ATA connector on the motherboard. For dual port implementation, one resistor for each port and each address line on the cable should be used. The best design practice for source termination is to have the chip drive two resistors, each of which is connected to the two cables (primary and secondary).

- * The chip driving the ATA signals should be located very close to the ATA connector. Technically if the chip is far from the connector this length should be added to the cable length, of which the sum should not exceed 18 inches. In theory, if the chipset has enough timing margin to account for the additional delay on the PCB (about 2 ns each way for 10 inches trace), and the capacitance of each trace does not exceed the 25 pF maximum of the ATA spec, and the crosstalk is not excessive, then this distance shouldn't matter. In practice, system designers should place the chip as close as possible to the ATA connector(s).

Directly connecting two cables to the chipset will not work. Separate termination resistors for each cable are required for any hope of system operation.

We recommend starting your termination circuit model with a value of 22 to 47 ohm series resistor on the control signals at the host system end. Adjust the value for the best suppression of ringing in the bus signal line you are working with.

The ATA specification requires that a driver be capable of sinking 12 mA of current and still maintain a valid TTL low (< 0.5 V). For signals that are driven by the drive toward the host, i.e. DD0 through DD15, DMARQ, and INTRQ; an 82 ohm resistor would provide an unacceptably large DC voltage drop. To remain within the guidelines of the ATA-2 specification, a resistor value of 22 to 47 ohms can be used.

Signals that are driven back to the host, such as DD0 through DD15 (during a read), DMARQ, and INTRQ will also be delayed due to the presence of a series resistor. Assuming that the host has 50 pF of stray capacitance, the incremental delay is less than 1.7 ns using a 33 ohm series resistor value. Delays of these magnitudes are not a problem in the context of typical ATA bus usage.

With the arrival of the new ATA-2 systems and adapters, it is imperative that termination circuits be re-evaluated, simulated, and tried out in the real systems before choosing the values of the components.

220