X3T10/95-168rs

ATA Signal Integrity Issues

Quantum Corporation
Systems Engineering
Steve Reames
Farbod Falakfarsa

10-Jul-95

Version 1.03

Contents

1.0 Introduction

The ATA Bus (a.k.a. IDE bus) is a disk drive interface originally designed for the ISA Bus of the IBM PC/AT. With the advent of faster disk drives the definition of the ATA Bus has been expanded to include new operating modes. Each of the PIO modes, numbered zero through four, is faster than the one before it (higher numbers translate to faster transfer rates). Modes 0, 1, and 2 correspond to the ATA interface as originally defined. PIO Mode 3 defines a maximum transfer rate of 11.1 MB/s and PIO Mode 4 defines a maximum rate of 16.7 MB/s. Additional DMA modes have also been defined, with Multiword DMA Mode 0 corresponding to the original interface, and DMA Modes 1 and 2 being faster transfer rates. Multiword DMA Mode 2 is the same speed as the new PIO Mode 4.

With this increased speed the weaknesses of the original ATA cabling and on-board interconnect scheme have become apparent. System manufacturers, chipset designers, and disk drive manufacturers must all take measures to insure that signal integrity is maintained on the bus. The areas of concern are:

1) ringing due to improper termination,
2) crosstalk between signals, and
3) bus timing.

The intended audience for this paper is digital and analog engineers who design circuits interfacing to the ATA bus. Familiarity with the ATA specification and a basic understanding of circuit theory is assumed.

1.1 The Problems

Early implementations of the ATA bus used LS-TTL parts to drive an 18-inch cable. The slow edges of LS-TTL and the short cable length worked well at the time. PIO Modes 3 and 4 demand higher performance. In an effort to cut cycle times, edge rates have inadvertently been increased, causing ringing on the cable and increased crosstalk between adjacent signals.

When an ATA host adapter was little more than a few buffers and some gates there was no issue with host adapter timing. With the advent of local bus architectures and faster transfer rates, timing issues have become more important. Propagation delay with worst-case loads must now be taken into account when designing host adapters.

1.2 The Goals

The recommendations in this document make the following assumptions. The word "drive" is used generically to describe disk drives and other peripheral devices on the ATA bus.

- Backward compatibility must be maintained. Old drives must work with new host adapters, and old host adapters must work with new drives.
- The ATA-2 standard must be followed as closely as possible. Without this, solutions implemented by different manufacturers will tend to diverge, creating incompatible systems.
- Solutions must be simple and inexpensive. The market for ATA products is very cost sensitive.

2.0 Termination

When analyzing the ATA bus, the standard 18-inch ribbon cable used to connect devices could be considered to be either a transmission line or a lumped LC circuit. Analog circuit designers generally use the rule of thumb that if the edge rate is less than four times the cable propagation delay, then it is a transmission line. Otherwise it can be considered to be a lumped LC.[1]

[1]Different ratios are used by different designers. A survey of textbooks shows that values of three times, four times, six times, and even sqrt(2*pi) have been suggested.

The cable used almost exclusively is a PVC-coated 40-conductor ribbon cable with 0.05 inch spacing. This cable can be modeled as a transmission line with a typical characteristic impedance of 110 ohms and propagation velocity of 60% c.[2] This gives a propagation delay of 2.5 ns. The edge rates from both hosts and drives are usually faster than 10 ns (4 x 2.5 ns), so a transmission line model applies.

2.1 The Problem

Many users have experienced problems with early implementations of PIO Mode 3 drives and hosts. Most failures in the systems observed can be attributed to signal integrity problems on the control lines that go from the host to the drive. The problem appears most frequently as ringing on the DIOR- (read command) and DIOW- (write command) lines.

During a read cycle when DIOR- is asserted, it is possible for the ringing to create a short duration deassertion pulse (figure 1). This pulse occurs early in the read cycle. Inside the ATA interface portion of the datapath controller chip is a FIFO buffer that contains the data to be read. The extra pulse on the DIOR- line advances the FIFO pointer by one. This results in losing one word of data. The host system read operation therfore receives one word too few, and the remaining bytes are shifted. A typical data sequence might look like . . .W7, W8, W9, W11, W12 . . . Notice that word 10 is missing from the returned data. This also means that the host will try to read one more word from the drive than the drive has remaining. Depending on the implementation of the BIOS, this may lock-up the system or simply return a byte of garbage at the end of the sector.

Pulse slivers due to ringing on the DIOW- line cause a similar problem during writes. The pulse sliver advances the FIFO pointer by one unexpectedly, writing an extra word of garbage into the FIFO. Subsequent data bytes are shifted by one word. A typical stored data sequence on the drive might look like . . . W7, W8, W9, XX, W10, W11 . . . In this example an extra word was inserted during the write cycle for word 10. From the drive's point of view, the host is trying to write 514 bytes rather than the expected 512 bytes. The drive will throw away the final word and probably flag an error. A properly written BIOS will detect this error and indicate a problem to the user.

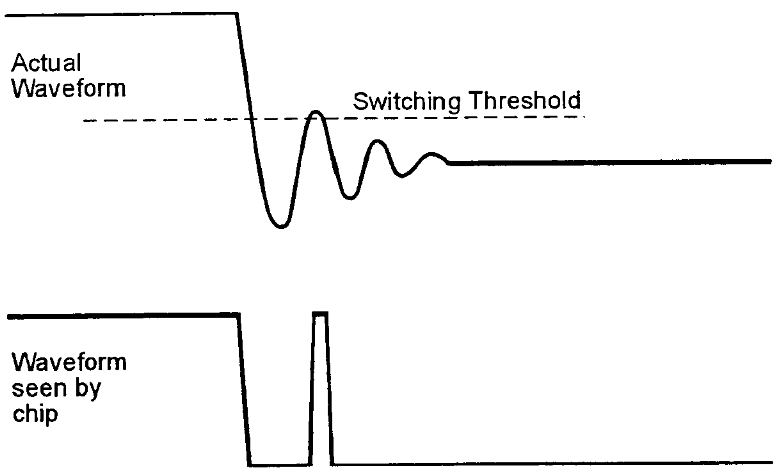

Figure 1 — Typical ringing on ATA bus and its effect

These are only two examples of a systemic problem. Ringing on any control signal, and possibly on data lines, can cause system failures or data loss. To address this problem it is necessary to examine the circuit structure of the ATA bus.

Figure 2 shows the seven basic driver/receiver structures that appear in ATA bus interfaces. The host circuitry appears on the left side of the diagram and the drive circuitry appears on the right. The first

[2]Measurements taken on a sample cable gave an impedance of 107 ohms and a delay of 2.6 ns (59% propagation velocity).

circuit in figure 2 shows the structure of the seven control lines that go from the host to the drive. A SPICE model of the circuit can be designed if we make some assumptions about the circuitry at the host and drive. Virtually all drives today use a CMOS VLSI chip as part of the bus interface. This high-impedance input can be modeled with clamp diodes to supply and ground and a typical input capacitance of 8 pF (Figure 3). Since the ringing problem is worse with CMOS VLSI bridge chips at the source, we will model the host as a voltage source with 1 ns edges, a 12 ohm output impedance, and clamp diodes to supply and ground. The ribbon cable is modeled as a 110 ohm transmission line. The resulting SPICE model appears in Figure 3.

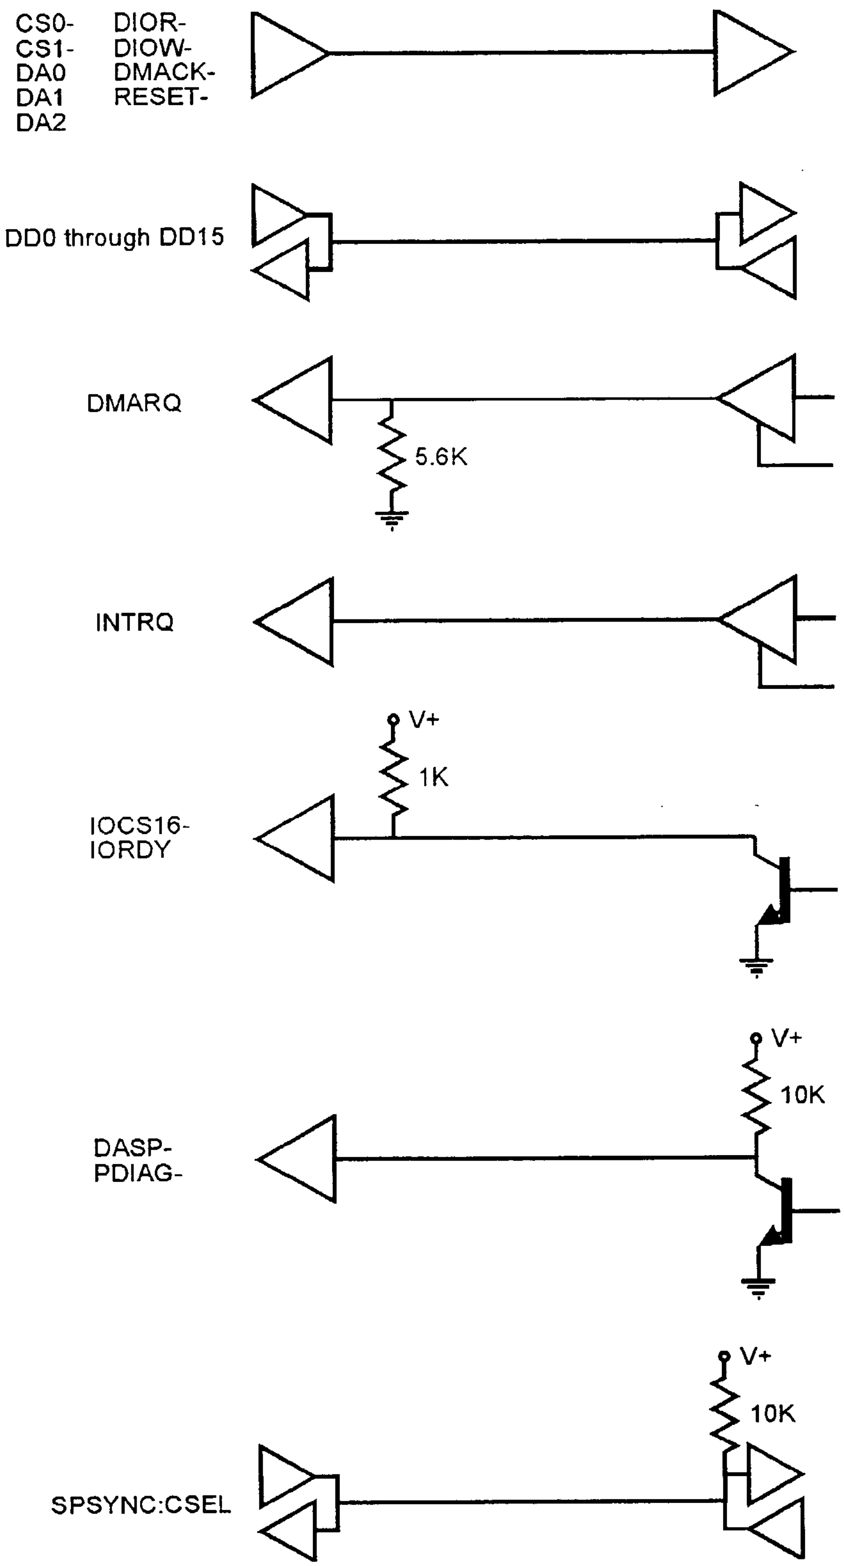

Figure 2 — The seven basic ATA bus structures
(host circuit on left, drive circuit on right)

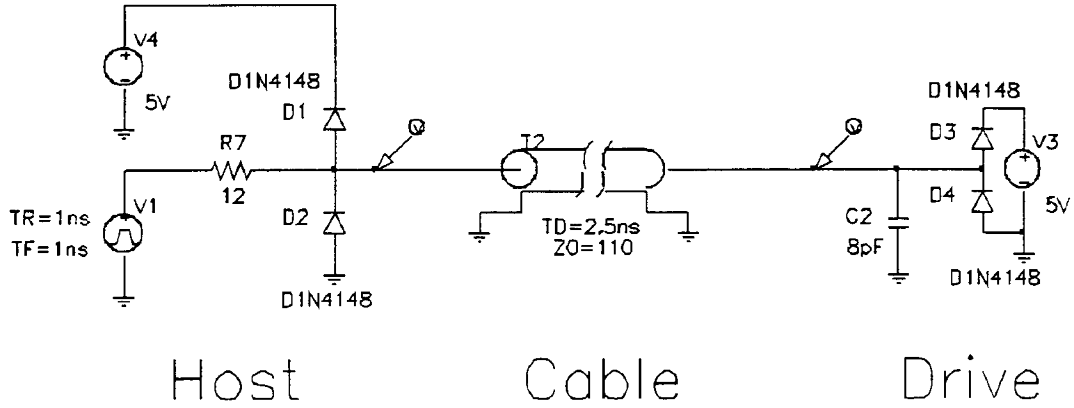

Figure 3 — Schematic of SPICE simulation model

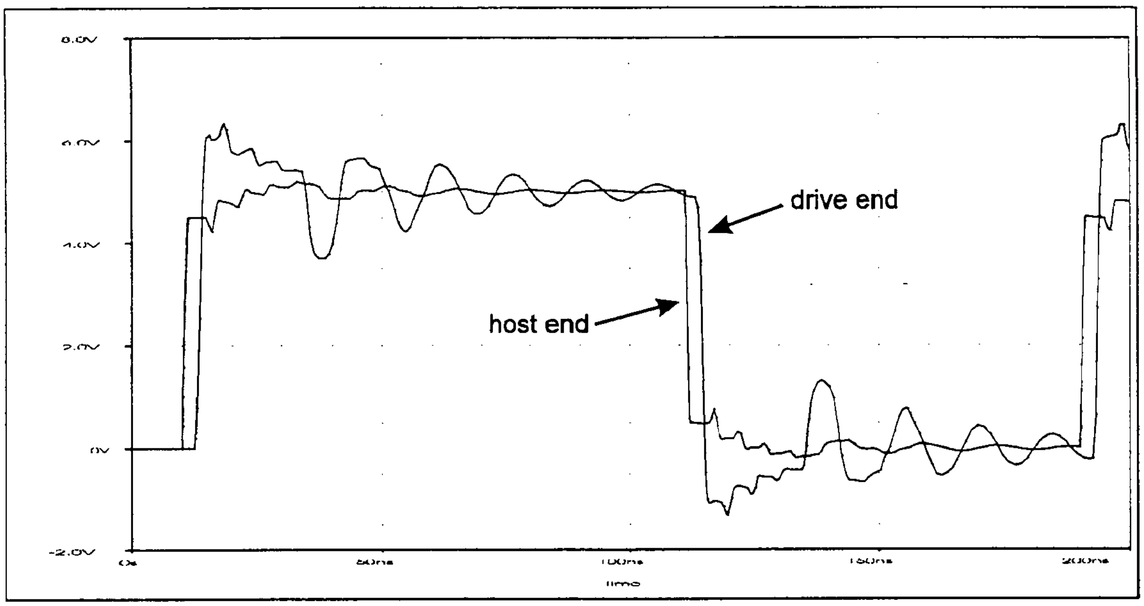

Figure 4 — Simulation waveforms at host and drive ends of cable

The simulation results in Figure 4 show the waveforms at both the host and drive ends of the cable. The signal at the drive end has ringing of sufficient amplitude to cause false triggering of the drive. This is confirmed by transmission line theory which indicates that ringing will occur whenever the source impedance is lower than the characteristic impedance of the cable, and the termination is of higher impedance than the cable. The greater the mismatch, the greater the amplitude of the ringing. The oscilloscope trace shown in Figure 5 confirms the results of our simulations.

The latest trend in ATA interface chips has aggravated the ringing problem. In an effort to decrease propagation delay, some bridge chip manufacturers have increased the output drive of the host in order to slew the output signal faster with the capacitive load of the cable. This has caused the edge rates and the output impedance to decrease, both of which increase the ringing at the drive end of the cable. The simulation in Figure 5 uses a generic driver and receiver — the problem of ringing is a fundamental characteristic of the ATA interface. This has not always been the case. To understand why ringing was not a problem in the past, read "Appendix A — How Did We Get Here?"

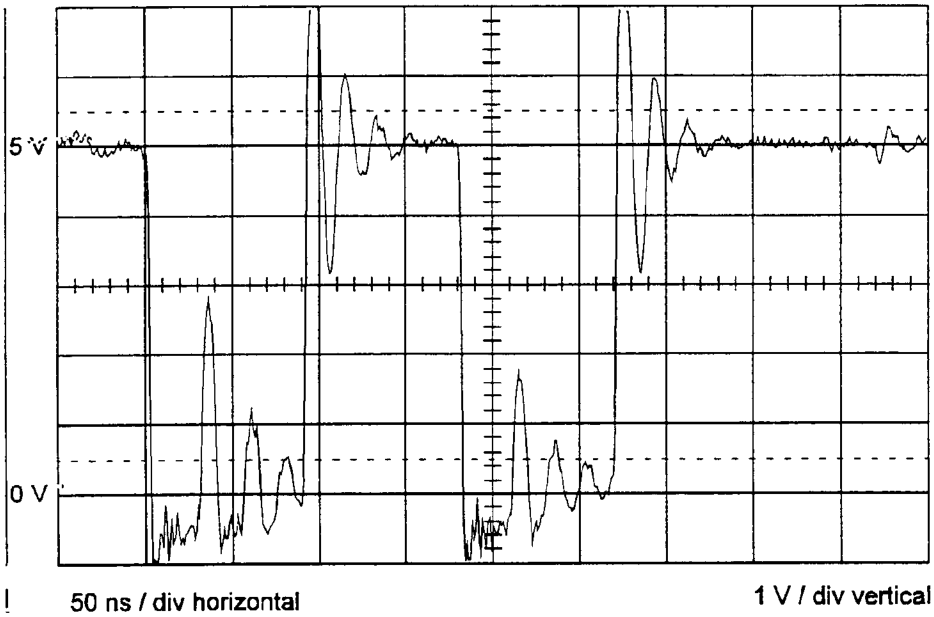

Figure 5 — Oscilloscope trace at drive end of DIOR- signal

2.2 What are the Options?

The proper solution is to terminate the transmission line. Either a series termination at the source or a parallel termination at the drive would be acceptable. Unfortunately, each of these solutions has problems of its own. A 110 ohm termination at the drive end would cause excessive DC loading. Having a termination on both drives in a master/slave configuration would result in too low of a load impedance, causing reflections on the cable again. And the drives are not necessarily at one end of the cable. It is perfectly legal in ATA to have the host at the middle of the cable and the drives at the two ends.[3]

Matching the source impedance of the host to the cable has similar problems. The impedance required is different when the host is in the middle of the cable as opposed to being at the end. Even with the host at one end of the cable, ringing can occur with a two drive configuration. This is because the input impedance of the drives is not infinite: they appear as a reactive load due to their input stray capacitance.

The SCSI interface standard avoids ringing by requiring terminations at each end of the physical cable and having each device drive the cable with a current sink. However, real SCSI configurations often have too many, too few, or improperly located terminations. Changing the ATA standard to a user-installed termination scheme would lose all backward compatibility and therefore is not considered to be a viable option.

One of the solutions used in the past to "fix" failing ATA configurations has been to place a capacitor at the input of the drive. Since the ringing is the result of a resonant system, we do not recommend adding purely reactive elements (capacitors and inductors) which simply change the frequency of oscillation. These elements may fix a given configuration of a drive and cable, but they

[3]Most disk drive manufacturers consider having the host in the middle of the cable to be an illegal configuration even though there is nothing in the ATA-2 specification to prohibit it. It does still occur in actual practice. It is debatable whether or not this configuration should be taken into account when designing cable termination. Some customers consider it to be bad form for a system to work when connected one way, but fail when the position of the host and a drive are swapped. Future versions of the ATA specification will most likely limit the host to residing at the end of the cable.

really just move the interfering resonance peaks to a different frequency, solving the problem only for that particular configuration. Proper solutions must include resistive elements to dissipate the energy stored in the transmission line.

No single solution meets the dual criteria of solving the ringing problem and being backward compatible with current systems. The suggested approach uses partial solutions in three different areas: partial termination at the host, partial termination at the drive, and edge rate control at both the host and the drive.

2.3 Design Goals

Before a solution can be designed the design goals must be explicitly stated. This leads us to the question of "How much ringing is acceptable?" To answer this question we must consider the design and specification of the ATA bus.

The ATA bus was originally designed to use standard TTL signals. TTL was designed with built-in noise margin. All drivers are required to have a "low" (zero) signal level of 0.5 V or less, and a "high" (one) signal level of 2.4 V or more. All receivers are specified to accept any signal below 0.8 V as a logical zero and any signal above 2.0 V as a logical one. This results in a low-side noise margin of 0.3 V (0.8 - 0.5) and a high-side margin of 0.4 V (2.4 - 2.0). Signals between 0.5 V and 2.0 V are in no man's land, and may be interpreted by the receiver as either a zero or a one. TTL compatible inputs typically use a switching threshold of 1.3 to 1.4 V.

Bus designers have long known that the noise margins of TTL are insufficient for signals passed on cables. To improve the noise margin inherent in TTL systems they have added hysteresis to the receiver input. Hysteresis changes the input switching threshold depending on the present state of the logic output of the receiver. For example, if the receiver is currently in a zero state, it might require an input voltage of 1.7 V before changing to a one. Once in a one state, the receiver might require the input voltage to drop below 0.9 V before changing back to a zero. Modern design practice dictates that all signals passing across a bus be received with hysteresis.

It is desireable that, even with ringing, the input signal remain less than 0.5 V after a falling edge and remain above 2.4 V after a rising edge. With CMOS drivers only the falling edge is of concern. This is due to the input switching threshold of TTL (typically 1.4 V) being closer to ground than to the supply. It turns out that designing to the 0.5 V requirement is too restrictive, so the looser requirement of 0.8 V has been used here. This relaxed requirement essentially removes the noise margin inherent in TTL and depends on receiver hysteresis for proper operation. As input hysteresis has been the norm in drive design for many years now, this limitation was not considered unreasonable.

Depending on system timing and other issues, a designer may elect to use a looser threshold of 0.9 V or a tighter one of 0.7 V. For these cases circuit simulation of the bus and receiver must be done to verify the design. The resulting termination circuits will have different values from those derived here.

2.4 Source Termination

A series resistor at the source (host) acts as a termination to the transmission line. When the value of the resistor matches the characteristic impedance of the cable (110 ohms) then the ringing is reduced to zero.[4] Resistor values less than 110 ohms will partially terminate the cable and reduce the ringing.[5]

The ATA specification requires that a source be able to sink 12 mA while maintaining a logical low output voltage of 0.5 V or less.[6] Adding a series resistor in the output of the driver causes the output logical zero voltage to increase with greater resistance. For example, if the unterminated logical zero output of the driver is 0.4 V, then a maximum series resistance of 8.3 ohms would be allowed ((0.5-

[4]This assumes that the output impedance of the driver is zero. In reality, optimum match occurs when the output impedance and the series resistor together equal the cable impedance.
[5]This also assumes that the host is at the end of the cable. If it is in the middle of the cable, then it sees two cables in parallel and the optimum impedance match would be 55 ohms.
[6]ATA-2 Working Draft, revision 3, section 4.5

0.4)/12mA). This DC voltage drop requirement acts in opposition to the higher resistance values required for cable termination.

If we assume that an unterminated CMOS output driver can drive a logical low to 0.2 V, then we have 0.3 V to drop across the series resistor. This gives us a maximum value of 25 ohms. A five-percent, 22 ohm resistor meets the requirement. Note that this places an addition requirement on the host interface chip: timing measurements must be made using a logical threshold of 0.2 V rather than 0.5 V as in the past.

Is a 22 ohm resistor adequate for reducing the ringing? The simulation was repeated using the same model as shown in Figure 3 with a 22 ohm series resistor added. The results of that simulation appear in Figure 6. The ringing has been significantly reduced from the previous simulation in Figure 4.

For maximum ringing we assumed that the drive(s) had CMOS input stages that did not provide significant DC loading. Yet for the series termination resistor calculation we assumed a logical low sink current of 12 mA. Both of these conditions cannot simultaneously occur in practice, but assuming the worst-case sink current gives us the best compatibility with older devices. The current proposal for the ATA-3 document is to decrease the DC current requirement from 12 mA to 4 mA. This will allow larger resistor values for the same DC voltage drop and will enable designers to more closely match the impedance of the cable.

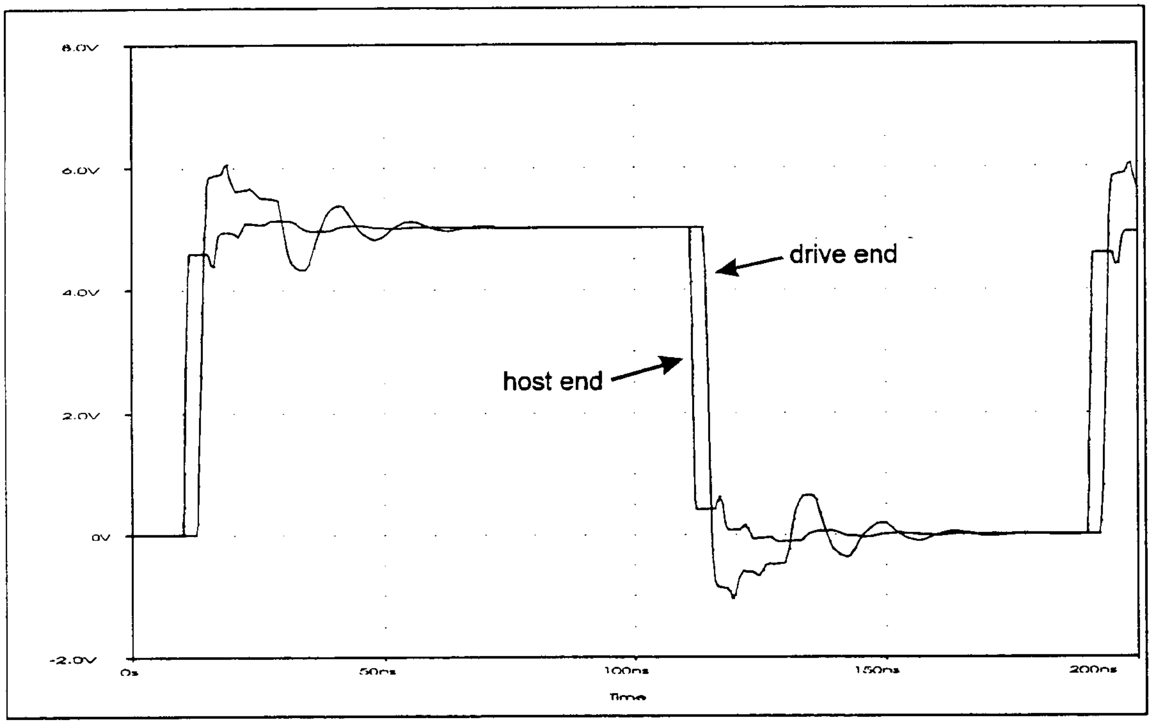

Figure 6 — Waveforms with 22 ohm series resistor at source

2.5 Drive Termination

Drive termination is more difficult than host termination. Viable solutions must work with one or two drives located anywhere along the cable. The host may be located at one end or in the middle of the cable. The host may or may not have termination. These and other considerations make drive termination a multifaceted problem.

The first constraint is the maximum DC loading allowed. The ATA specification requires that the host be capable of providing 400 µA of current while in a logical one state. Assuming each drive is allowed to take half of that amount, the minimum DC resistance allowed is 25K ohms.[7] For a 110 ohm transmission line, 25K ohms is as good as infinity. This means that any practical termination solution must not have significant DC loading.

One way of terminating the cable is with an "AC termination." This is a simple RC network that provides termination for high-frequency signals but does not load the line at DC (Figure 7). This circuit acts as both a cable termination and a filter for the ringing. The termination characteristics can be observed by looking at the ringing signal at the host when the circuit is connected or removed. When the circuit is in place, less energy is reflected back to the host, so the host waveform has less ringing. The lowpass filter characteristics of the circuit help decrease the amount of ringing presented to the interface circuitry of the drive.

Although this may appear to be an unusual method of terminating the cable, it is not without precedent. The IEEE P996 committee recognized the problems inherent in the design of the IBM PC/AT bus and recommended a series RC termination for increased "data integrity and system reliability." They suggested that the termination circuit be added to each end of the backplane or motherboard. The recommended values are 40 to 60 ohms for the resistor and 30 to 70 pF for the capacitor.

Deriving the optimum values for an ATA bus AC termination circuit is difficult. The easiest way of determining the values is to perform a number of trial-and-error SPICE simulations for different host and drive configurations. The values recommended by Quantum's System Engineering are 82 ohms and 10 pF. Simulations show that capacitance values between 8 pF and 20 pF work well. Since the input capacitance of many interface chips is between 8 and 10 pF, a discrete capacitor is often unnecessary. This reduces the cost of implementation on the disk drive. A conservative approach is to place pads so additional capacitance can be added if required.

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 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 run a SPICE simulation and measure the delay. The simulation should be run twice: once with a simple 25 pF load, and again with the proposed termination circuit. If the resulting delay of the proposed termination circuit is less than or equal to that obtained with a 25 pF load, then it meets the criterion for acceptance. The recommended termination of 82 ohms and 10 pF passes the test.

The major drawback of the RC termination circuit is that it adds delay to the signal. Since the ATA specification defines the timing at the input to the drive,[8] drive manufacturers must insure that their interface chip still works properly with the additional delay. The delay can be calculated for rising edges (2.0 V threshold) and falling edges (0.8 V threshold) with a fairly straightforward SPICE simulation. For the 82 ohm and 10 pF termination the delay is less than 1.5 ns.[9]

Will termination at both the host and the drive "over-terminate" the transmission line? Figure 8 shows the simulation results for drive termination with no host termination, and Figure 9 shows the same simulation with a 22 ohm host termination added. It is clear that termination at both the host and the drive results in the best signal integrity. To completely confirm the validity of the termination circuits more simulations must be performed with host termination and two drives with drive termination; two drives, one with and one without termination; host in the middle of the cable, etc. Quantum has performed hundreds of SPICE simulations and recommends these termination circuits with confidence.

Another option for controlling ringing at the drive is the use of a clamping circuit. Biased diodes have been shown to be excellent solutions, reducing the ringing to virtually zero. The advantage of clamp circuits is that they do not require any components in series with the signal, and therefore do not add any delay. This is particularly important for PIO Mode 4 operation. The disadvantage of clamping circuits is that they take considerably more space on the PCB and cost much more than passive elements. Some

[7]This assumes a CMOS output with a high output voltage of 5.0 V : 5.0V/200μA = 25KΩ

[8]ATA-2 Working Draft, revision 3, section 10.1

[9]0.7 ns for the rising edge, 1.2 ns for the falling edge, derived from simulations.

implementations have used clamping circuits on sensitive edge-triggered lines (such as DIOR- and DIOW-) and used passive terminations on less sensitive lines (such as data). Clamping circuits work well both with and without host-end termination and are worthy of further investigation.

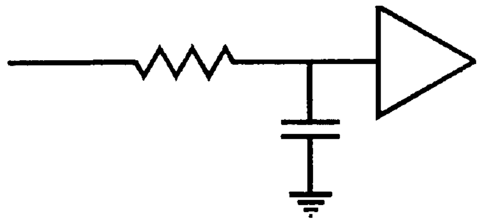

Figure 7 — AC termination circuit at drive end of cable

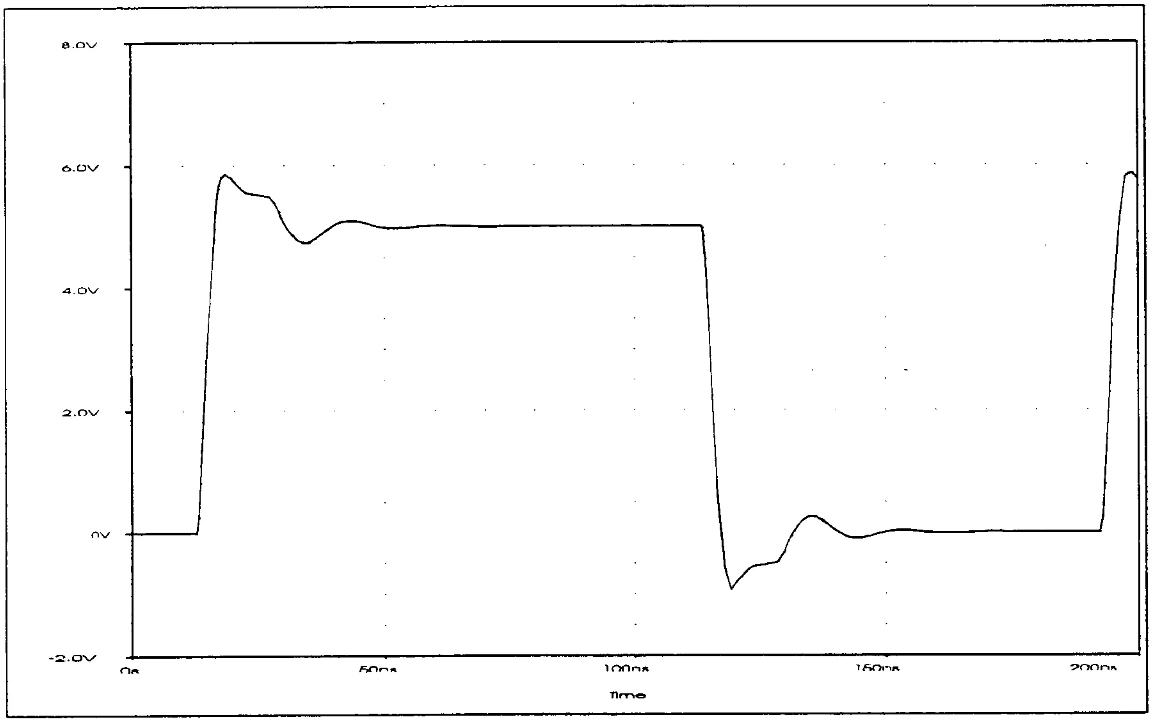

Figure 8 — Drive waveform with drive termination and no host termination

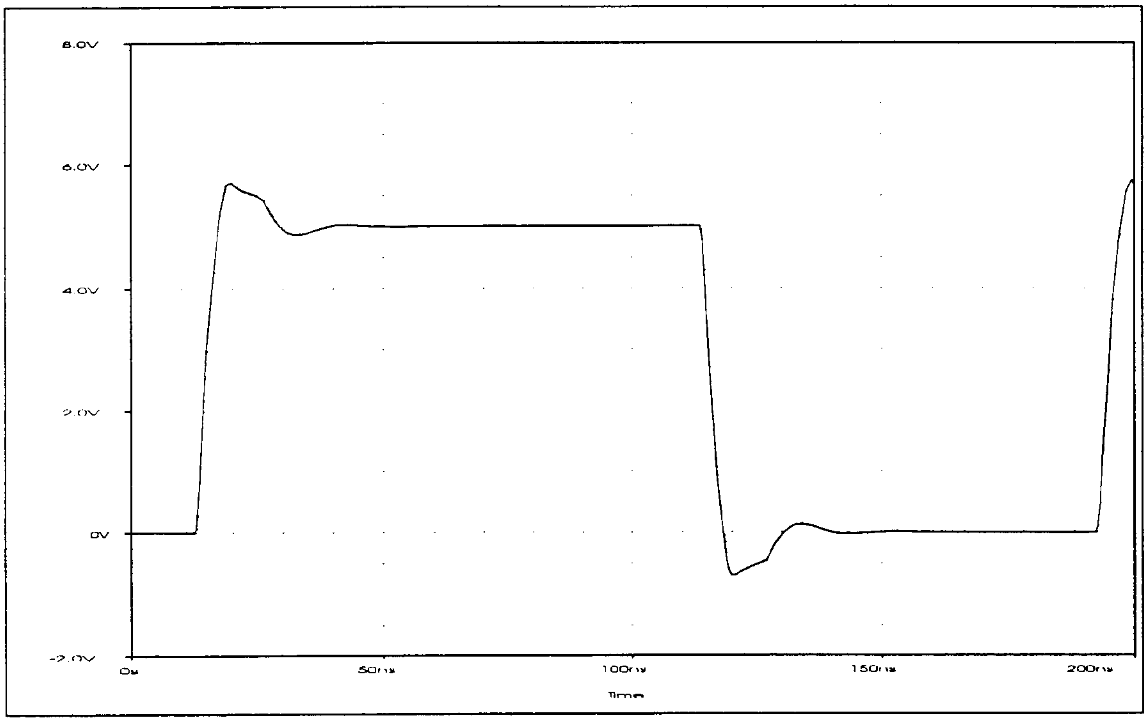

Figure 9 — Drive waveform with both drive and host terminations

2.6 Edge Rate Control

The ATA specification requires that all sources have a rise time of not less than 5 ns.[10] The original intent of this requirement was to avoid transmission line problems on the bus. One of the common misconceptions is that limiting the rise time of the source to 5 ns will fix the ringing problem.

A rule-of-thumb for analog designers has been that when the propagation delay of the cable exceeds one-quarter of the signal rise time, cable termination must be used.[11] In the case of the ATA bus, the worst case propagation delay of the cable is approximately 4 ns[12], so by this rule rise times of less than 16 ns require termination. Many local bus to ATA bridge chips available today have rise times of 1 to 2 ns, in violation of the ATA-2 requirement of 5 ns.

Risetime is poorly defined in the ATA specification. The ATA document says that the rise time must be a minimum of 5 ns into a 100 pF load. The easiest way to implement this from a chip designer's point of view is to decrease the drive of the I/O cell until the timing requirement is met. Unfortunately, very few systems in the real world ever approach 100 pF. Although the cable and the drives have maximum capacitance specifications, these capacitance values are never seen by the host. At DC and low frequencies the cable looks like a capacitor. But at high frequencies (or fast edge rates) the cable appears as a transmission line. A properly terminated transmission line appears to be a resistor with no capacitance or inductance. The capacitance of the drives is changed by the transmission line and appears as anything but pure capacitance at the output of the host. As a result, real-world systems rarely see more

[10] ATA-2 Working Draft, revision 3, section 4.5
[11] In reality this rule-of-thumb varies considerably. Various books use values of one-half, one-third, one-fifth, and even one over the square root of two times pi.
[12] 18-inch cable, 60% velocity factor; two drives, each drive having a maximum load of 25 pF.

than 25 pF of capacitive loading at the host. The reduced capacitance causes the I/O cell to slew faster, creating rise times less than 5 ns. In actual system testing, Quantum has found only a few systems that have rise times slower than 3 ns.

The best solution is to use special I/O cells that have slew rate feedback to keep the rise time at 5 ns regardless of load. These are more difficult to design than conventional I/O cells and consume more die area. This could be a problem for interface chip designs that are already pad ring limited. Another approach would be to use a conventional I/O cell that is designed to have 5 ns rise times into a 10 pF or 20 pF load. The total delay of the cell would be greater for heavier loads, but the maximum delay could be determined with SPICE modeling of a worst-case cable and load.

Rise time control is still an important tool for controlling ringing. Although it is not the total solution, simulations show marked improvement between sources with 1 ns rise times and sources with 5 ns rise times. Slower rise times give the added benefit of reduced crosstalk.

2.7 The Solution: A Combination

No one element — source termination, drive termination, nor rise time control — completely addresses the problem of ringing on the ATA bus. The solution is a combination of all three. Each item must be enough to exert some control over the ringing problem in order to maintain backward compatibility. With the faster transfer rates of PIO Mode 4 (and DMA Mode 2) it is unlikely that any systems will function reliably without improvements due to cooperation throughout the industry.

2.8 The Total Solution

The above discussion only addressed a particular group of signals driven by the host and received by the drive. There are other signals driven by the drive and received by the host that are equally susceptible to ringing. These signals must also be terminated, but in the opposite manner. The drive should insert 22 ohm resistors in series with signals it drives and the host should have an RC (or just R) receiving end termination.

The data lines are different in that they must move data bidirectionally. Strictly speaking, the data lines are not edge sensitive and are unaffected by ringing. This is true as long as the data signals have sufficient setup time to allow for bus settling. The settling time can be as long as 40 ns in severe cases. Excessive ringing on the data lines can induce spurious signals on adjacent control lines (crosstalk). Good design dictates that some type of ringing control be used on data lines, but perhaps not as much as on edge-sensitive control lines. A good compromise is to insert 22 ohm series resistors on data lines at both the host and the drive. The driving end sees the same source termination as before. The receiving end sees an RC network of 22 ohms combined with the input capacitance of the interface chip. This is enough to substantially reduce the ringing and minimize settling time.

The one remaining bus structure we have not discussed is the open collector outputs driven by the drive (IORDY, IOCS16-). This signal is being driven by a current source rather than a voltage source. Usually the transistors driving these signals are relatively slow and do not cause an excessive amount of ringing. The nature of these signals makes them relatively insensitive to ringing that might occur.

The following table summarizes the recommended changes to the signal lines on the ATA bus:

Signal Name	Host Termination	Drive Termination
DIOR-, DIOW-	22 ohm series	82 ohm & 10 pF
CS0-, CS1-	22 ohm series	82 ohm & 10 pF
DA0, DA1, DA2	22 ohm series	82 ohm & 10 pF
DMACK-	22 ohm series	82 ohm & 10 pF
RESET-	22 ohm series	82 ohm & 10 pF
DD0 through DD15	22 ohm series	22 ohm series
DMARQ	82 ohm & 10 pF	22 ohm series
INTRQ	82 ohm & 10 pF	22 ohm series
IOCS16-, IORDY	no change	no change
DASP-, PDIAG-	no change	no change
SPSYNC:CSEL	no change	no change

Note: For the 82 ohm & 10 pF termination, the capacitor may be deleted if the interface chip and PCB layout have at least 8 pF of capacitance.

2.9 Dual Port Cabling

One of the recent enhancements to the ATA bus has been the use of primary and secondary ports, allowing the user to attach up to four drives. The optimal way to implement dual ports is to have two completely separate interfaces that have no circuitry in common. This guarantees isolation between the ports and insures that no interference will occur.[13]

The advent of local bus bridge chips has introduced new driving forces to the dual port cabling issue. Implementing two independent ATA ports requires an inordinate number of pins on the bridge chip. Due to the cost of pins, some designs have combined the data lines of the two ports into one set of pins. Sharing the data lines (or any other lines) in this way without termination is asking for trouble. Simulations confirm that the ringing in such configurations is large and complex, particularly if the loads on the two cables are not balanced.

One alternative pin-saving solution would be to add a set of external buffers such as the 74ALS245. This would require three new control lines but would save sixteen data lines for a net improvement of thirteen pins. This also would require four packages on the PCB.

An economical solution is to add independent series resistors for each line (see Figure 10). Energy reflected back from the first cable passes through one termination resistor before getting to the host. The reflected signal is further attenuated as it passes through the second resistor and into the second cable. This signal is reflected from the end of the second cable (with loss), and must pass through the termination resistor again before arriving at the host. This provides sufficient attenuation of reflected signals to make ringing not a problem.

Proving that this scheme works is a significant task. Simulations must be performed for all combinations of one drive vs. two drives, drives with terminations and drives without terminations, drives with little load capacitance and maximum load capacitance, etc. With two ports the number of combinations is daunting. Extensive simulation work at Quantum has not yet revealed a combination of drives that fails to operate properly. Simulation results have been spot checked against actual systems.

[13] At least in the hardware. There is still the issue of software compatibility.

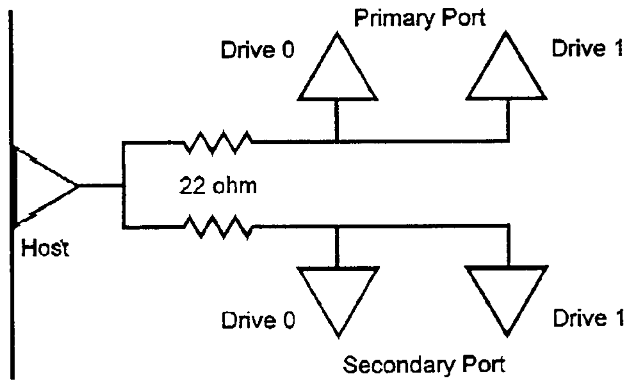

Figure 10 — Preferred connection for shared lines in dual port systems

Not all of the signal lines in a shared dual port interface can be shared. It is obvious that if the chip selects (CS0-, CS1-) and the data strobes (DIOR-, DIOW-) were shared, then it would be impossible to differentiate between the primary and secondary ports. A write to a drive on one port would cause the same action to occur on the other port, destroying the data on the other drive. The data strobe lines are sensitive edge-triggered signals while the chip selects act more like level-sensitive address lines. Quantum recommends sharing the less sensitive chip selects and not sharing the data strobes.

The purpose of the chip selects is to select the active drive. It makes sense to have independent chip selects for primary and secondary ports. This allows the DIOR- and DIOW- lines to be shared. The ATA bus signals that can and cannot be shared are shown in Table 1.

This table makes some assumptions about how the dual porting is being implemented. If the data lines are shared, there can not be simultaneous accesses to the primary and secondary ports. This allows the DMACK-, DMARQ, IOCS16-, and IORDY lines also to be shared. The INTRQ and DMARQ signals are driven by tristate buffers on the drives. Either the master or slave will have its tristate driver enabled depending on the state of the DEV bit in the Device/Head Register. Therefore INTRQ and DMARQ cannot be shared because either the master or slave drive will be driving these lines at all times. The primary port drives do not know about the secondary port drives, so sharing these lines would create a conflict. In theory the DMACK- line could be shared since it is driven by the host. In practice this is not recommended. It is likely that some drives respond unconditionally to the DMACK- signal, whether they ever requested a DMA cycle or not. This could lead to a conflict on a DMA cycle between a primary port drive and a secondary port drive during the data cycle. For these reasons the INTRQ, DMARQ, and DMACK- lines should not be shared.

Signal Name	Shareable	Not Shareable
DIOR-, DIOW-		X
CS0-, CS1-	X	
DA0, DA1, DA2	X	
DMACK-		X
RESET-	X	
DD0 — DD15	X	
DMARQ		X
INTRQ		X
IOCS16-, IORDY	X	
DASP-, PDIAG-		X
SPSYNC:CSEL		X

Table 1 — Possible sharing of ATA signals in dual port configurations

The DASP- lines cannot be shared. If there are two drives on the primary port, and one drive on the secondary port. With the DASP- lines connected, the single drive on the secondary port will incorrectly "see" the slave device on the primary port. This would be a problem for all manufacturers who follow the ATA CAM specifications. Similar problems can occur with the PDIAG- lines; they cannot be shared.

The SPSYNC:CSEL line is designed to be vendor specific. It may be worth considering the ability to have these lines separate or connected by a jumper since it is not possible to predict exactly how a drive manufacturer will use this signal.

Sharing lines between ports will lead to lower performance than fully independent ports. Newer operating systems are multithreaded, so it is possible for accesses to be occurring independently and concurrently on the primary and secondary ports. With shared signals lines concurrent access is not possible, hence the impact on performance. At the 1995 Windows Hardware Engineering Conference (WinHEC), Microsoft recommended the use of fully independent primary and secondary ports.

Note that the entries in Table 1 are not the last word on line sharing. For example, if separate DIOR- and DIOW- lines were provided for primary and secondary ports, then the CS0- and CS1- lines could be shared. Many combinations of shared and non-shared lines are possible provided that the functionality of the signals is carefully considered.

3.0 Crosstalk

Crosstalk is switching on one signal line causing induced signals in an adjacent or nearby line. Crosstalk has not been a significant issue in the past with slower edge rates; in newer systems the problem is often masked by ringing. Once the cable is terminated and the ringing is under control, then the presence of crosstalk becomes apparent.

3.1 Coupling Mechanisms

There are two mechanisms by which a signal couples into an adjacent line. The first is coupling capacitance, and the second is mutual inductance. As a switching signal wavefront propagates down the cable it couples energy into the adjacent line. Once this energy is in the second transmission line, it propagates in both directions: forward toward the receiver and back toward the source (Figure 11).

First we can examine the forward coupling components. The voltage induced in the second transmission line is proportional to the coupling coefficient, the inductance, and the rate of change of current in the primary side. This is a negative voltage: a positive current spike in the primary line results in a negative voltage spike in the secondary line.

The coupling capacitance between the two line causes a current pulse in the secondary line proportional to the capacitance and the rate of change of voltage on the primary side. A positive voltage step on the primary line causes a positive voltage spike on the secondary line.[14]

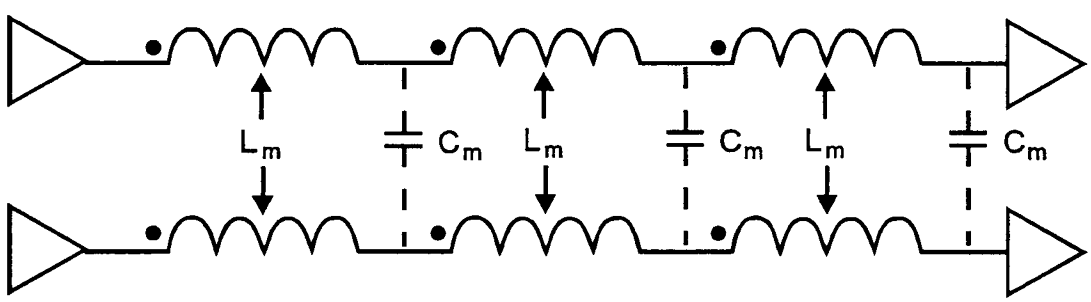

Figure 11 — Crosstalk coupling mechanisms

[14]This is a very cursory description of the crosstalk mechanism. For an excellent treatment of the subject, read "High Speed Digital Design: A Handbook of Black Magic" by Howard Johnson and Martin Graham, 1993, Prentice-Hall Inc., ISBN 0-13-395724-1

These two coupling mechanisms have some interesting characteristics. The polarity of the coupling is opposite for the mutual inductance and coupling capacitance. If the magnitudes of these effects are comparable, then they will cancel, resulting in no forward crosstalk. Unfortunately, accurately computing these values is difficult, and the easiest way to determine the actual amount of crosstalk is to measure it. The other noteworthy characteristic is that the magnitude of the coupled signal is proportional to the rate of change of the signal in the primary line. This is a major reason for controlling the slew rate on ATA bus drivers. Earlier we said that ringing on the data lines is not necessarily a problem. Here we see that fast edge rates and ringing on the data lines can couple by crosstalk into adjacent control lines, causing control sequence errors through mistriggering. It is unlikely that crosstalk from data lines would cause an observable failure in a laboratory environment. But the presence of crosstalk-induced voltage spikes on the control signals reduces the noise margin, and may increase the long-term error rate.

The amplitude of the coupled signal is proportional to the total amount of coupling capacitance and mutual inductance, and is therefore proportional to cable length. Once a line is terminated properly, ringing is no longer a function of length. This leaves crosstalk as the major factor limiting cable length.

Reducing crosstalk involves reducing the mutual inductance, reducing the coupling capacitance, or decreasing the source signal amplitude. Controlling the inductance and capacitance can be done by either keeping the length of the cable short or by increasing the distance between conductors. Placing a ground conductor between critical signals increases the separation of the signals and also adds a shielding effect from the intervening ground. In the ATA environment the only control we can exercise over the cable is to keep the length at 18 inches or less. The amplitude of the source signal cannot be reduced and still maintain ATA compatability, but we do have control over some elements of the source signal. Slew rate limitation reduces the high-frequency components of the source signal and therefore reduces the coupling of these components into adjacent lines. Terminating the lines reduces ringing which also decreases the amount of energy coupled at the ringing frequency.

4.0 Bus Timing

Terminating the ATA bus has its cost. Partial terminations at the host and the drive increases propagation delays throughout the system. The ATA document specifies that timing is referenced to the input pins of the disk drive.[15] This means that most of the timing issues must be addressed by systems manufacturers and bridge chip designers.

4.1 The Issues

The most significant timing issue is the propagation delay of the cable. This needs to be added to the host-side timing. The SPICE model in Figure 12 shows an unterminated host with very fast rise times driving a cable with worst case loads. Two unterminated disk drives are assumed with the maximum allowed capacitive loading of 25 pF.[16] The simulation results are shown in Figure 13. The period of the ringing is four times the propagation delay of the cable. This simulation shows a cable propagation delay of 5.6 ns. This is *twice* the value obtained by assuming an 18-inch cable with a propagation velocity of 60% c. The additional delay is due to the presence of the capacitive loads on the cable. This result is important to system designers who must take into account worst-case cable delay when specifying the bridge chip timing.

[15] ATA-2 Working Draft, revision 3, section 10.1
[16] ATA-2 Working Draft, revision 3, section 4.5, table 3

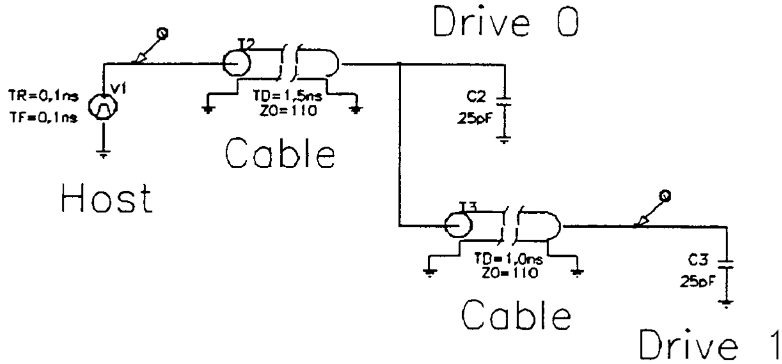

Figure 12 — SPICE model of ATA cable with worst case loads

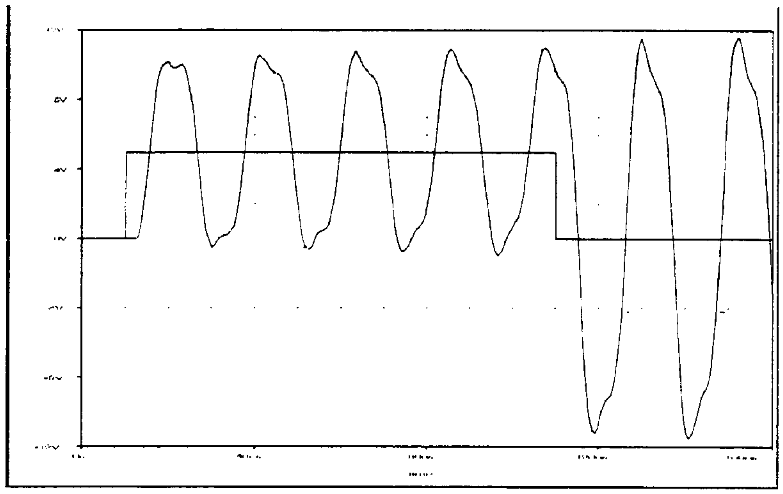

Figure 13 — Simulation of unterminated ATA cable with worst case loads

From the host point of view, all of the ATA timings must be corrected by adding the propagation delay of the cable to insure that the timing is correct at the input pins of the furthest drive. Figure 14 shows a typical corrected result using the read cycle data setup time as an example. The ATA document specifies a setup time *at the drive* of 20 ns (PIO Mode 3). The remaining setup time at the host is only 8.8 ns (20 - 2 x 5.6).

4.2 The Influence of Termination

If the host has a series partial 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 incremental delay added by a series termination resistor at the host is approximately 0.1 ns for a 22 ohm resistor and 1.7 ns for an 82 ohm resistor.[17] The extra delay of higher resistor values is one of the reasons that Quantum recommends 22 ohm series resistors at the host.

[17] Two drive load, 25 pF at each drive, 18-inch cable, 25 pF at host

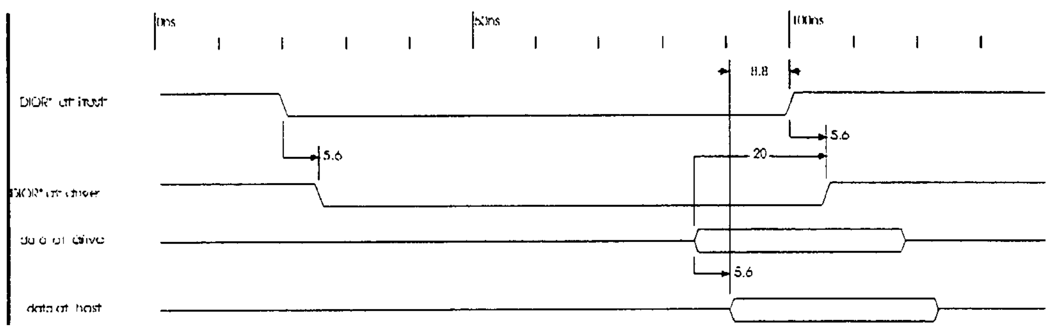

Figure 14 — Host data setup time during a read cycle corrected for a propagation delay of 5.6 ns.

The series resistor at the host must be located as close as possible to the ATA connector. To see the importance of this, SPICE simulations can be done with the stray capacitance on the driver side of the series resistor and again with the stray capacitance on the cable side. The ringing is reduced when the stray capacitance of the host is on the driver side of the series resistor. A related issue is the distance from the host adapter chip (or chipset) to the ATA connector. Some motherboards have the chip located up to 10 inches away from the connector. This effectively adds another 10 inches to the 18-inch ribbon cable, resulting in an equivalent cable length of 28 inches![18]

This additional length is not necessarily a problem. If the system manufacturer takes the extra trace delay into account in the application of the host adapter chip, and the total capacitance is kept below the ATA limit of 25 pF, then in theory no one should know the difference. Real-world experience indicates that this calculation is rarely done. The distance from the chip to the connector is not addressed in the ATA specification. Quantum recommends keeping the connector within 3 inches (by trace length) of the host adapter chip.

4.3 Calculating Rise Time

Chip designers often 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). Simulation is then performed with 100 pF capacitance on the output. This does not give an accurate measurement of the timing. A better approximation is to use an output capacitance for the motherboard, a host end termination resistor, and a transmission line to the drives (Figure 15).

To illustrate how these models are different, suppose that the output cell simulation is 2 ns too slow. The chip designer (using a 100 pF model) might be inclined to increase the drive current of the output devices. With enough drive current into a purely capacitive load, the 2 ns can be removed, bringing the output cell timing back into spec.

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 time reduction cannot be achieved by increasing the output drive current. The time must come from somewhere else in the internal circuitry. Increasing the output drive current only increases the edge speed, making the ringing worse at the drive end. (some time is gained with the faster edge speed, but not

[18] The traces on the PC board are anywhere from 50 to 200 ohm impedance, so the electrical length of the trace cannot simply be added to the 110 ohm ribbon cable. A SPICE simulation can be used to find the actual delay.

nearly as much as is predicted with a simple capacitive load model.) It is thus not possible to decrease the overhead by increasing the drive current beyond the 12 mA specified by the ATA document.

Most ASIC designers will find that simulations using the recommended model of Figure 15 will show their output cells to be faster than in models with a 100 pF load.

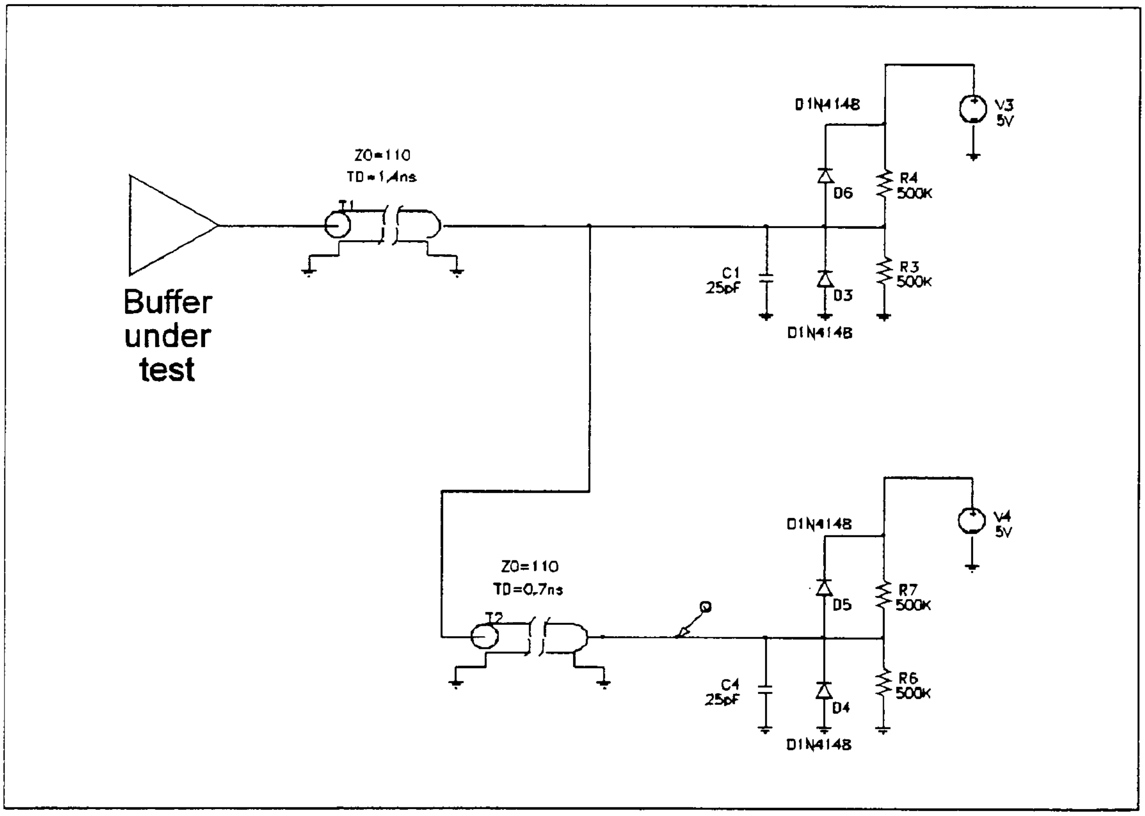

Figure 15 — Recommended load for modeling I/O cell propagation delay

4.4 Measuring Propagation Delay

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. 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).

5.0 Summary of Guidelines

This summary is a collection of reminders for drive, system, and chipset designers. They are separated into three groups by relevancy. It is pointless to remind a system designer to put termination on the drive, or to keep the rise time above 5 ns since the system designer does not have control over these parameters. But for this example, it might be wise for the system designer to show this list to their drive supplier and chipset supplier and ask whether these guidelines have been followed. The lists below are not intended to be a strict mandate, but a tool to help everyone build compatible, reliable, high-performance products.

5.1 Guidelines for Drive Designers

- Terminate signals received by the drive with a series resistor between 47 and 100 ohms (82 ohms recommended). Consider adding capacitors to ground on these lines if the input capacitance is less than 10 pF, or use active clamping circuits. Add series resistors between 22 and 47 ohms (22 ohms recommended) to bidirectional lines and driven lines. Place these resistors as close to the ATA connector as possible.
- Verify that the termination circuit used on received signals has less than 25 pF of equivalent capacitance.
- Perform a timing analysis to verify that ATA timings are met at the *input to the drive*. Include the time delay due to propagation and cable termination circuits.

5.2 Guidelines for System Designers

- Do not use any value less than 1K ohm for pull up resistors on ATA open-collector signals such as IORDY and IOCS16- (as per the standard).
- The ATA host adapter chip should be located as close as possible to the ATA connector. Keep the trace length between them less than 3 inches.
- After PCB fabrication, verify that the total input capacitance at the host is less than 25 pF.
- Use series resistors of 22 to 47 ohms on lines driven by the host (22 ohms recommended). Use series resistors of 47 to 100 ohms on lines received by the host (82 ohms recommended). Place these resistors as close to the ATA connector as possible.
- Perform a system timing analysis to verify that ATA timings are met at the input to the drive.
- For dual port implementations, one series resistor for each port should be used on every shared line (22 ohms recommended). These resistors should be placed as close to the ATA connector as possible.
- For dual port implementations, the signal lines CS0- and CS1- may be shared, or the signals DIOR- and DIOW- may be shared, but not both pairs.
- For dual port implementations, do not attempt to share DASP- or PDIAG- signal lines.
- For dual port implementations, perform a system timing analysis to verify that ATA timings are met at the input of the drive. In particular watch the assertion widths of DIOR- and DIOW- to insure that they meet the specification.

5.3 Guidelines for Chip Designers

- Design I/O cells to have rise and fall times of 5 ns or more under both minimum and maximum load conditions.
- Perform timing simulations using a transmission line load model, not a 100 pF capacitor model.
- Take worst-case cable delay into account when designing the ATA interface. Insure that ATA timing can be met at the drive-end of the cable. Provide typical application data with timing for system manufacturers.

Appendix A — How did we get here?

One of the frequently asked questions is why are we seeing problems with ringing on the ATA bus now when we didn't before. The answer is in the edge speed of the logic. How the edge speed changed can be found by looking at the history of the IBM PC.

When the IBM PC/AT was introduced in 1983, the 8 MHz 80286 processor quickly became the dominant platform. The AT bus (now called the ISA bus) became standardized around an 8 MHz processor speed. When the first ATA disk drive was introduced a few years later, it was designed as a simple extension of the bus — hence the name "AT Attachment" interface (Figure A.1). The idea was to remove the disk drive controller electronics from the PC and place them on the drive instead. What remained on the PC was a pair of bidirectional data buffers and an address decoder. This simple interface was most often implemented with a pair of 74LS245 buffers and a programmable logic device such as a PAL. These TTL devices had rise and fall times in the 5 to 6 ns range. Although this was fast enough to cause some ringing on the ATA bus, it was not so severe as to prevent the millions of successful system implementations.

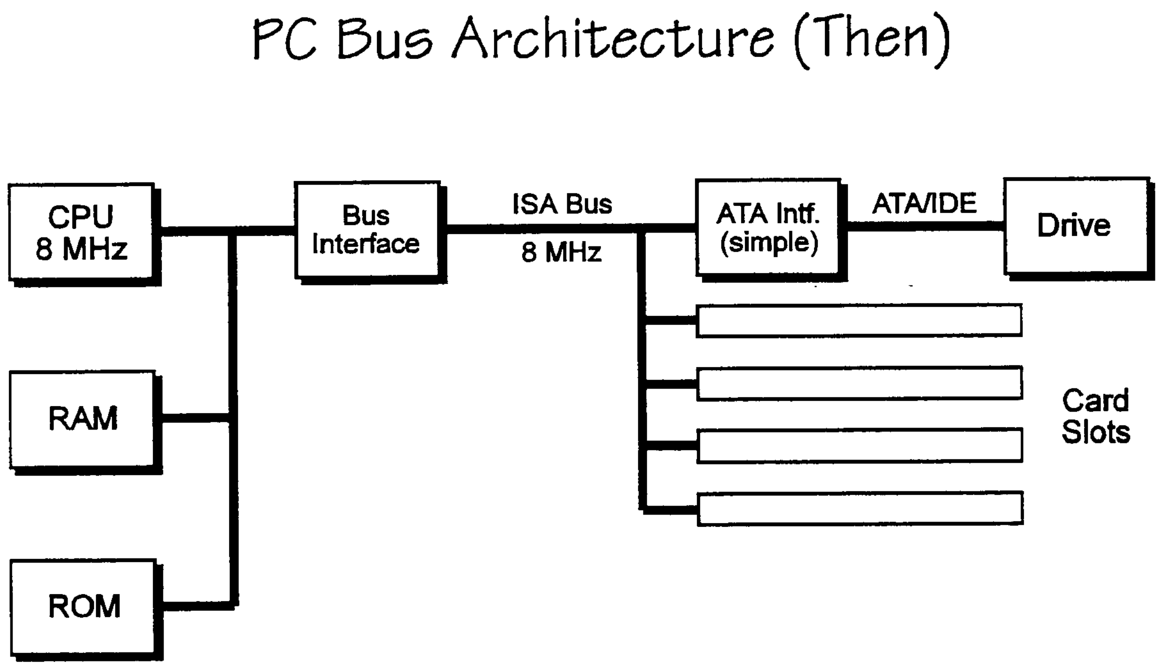

Figure A.1 — Original IBM PC/AT architecture

The architecture of modern PCs has changed somewhat from the original PC/AT. As the processor speed increased it became necessary to separate the processor bus from the ISA bus (Figure A.2). To maintain compatibility the ISA bus continues to run at an 8 MHz rate. Processors, and their associated busses, have increased from 8 MHz to 12, 16, 25, and now 33 MHz.[19] Disk drives have also increased in speed. The increase in rotational speed and linear bit density has increased the rate at which data comes off the heads, and the presence of cache on the drive makes data available at the access rates of RAM. This has created a data bottleneck at the ISA bus. The drive is faster, the processor is faster, but the data can't be moved from one to the other any faster.

[19]Processors that run faster than 33 MHz, such as the 486DX2-66, still tend to keep the external processor bus at 33 MHz.

PC Bus Architecture (Now)

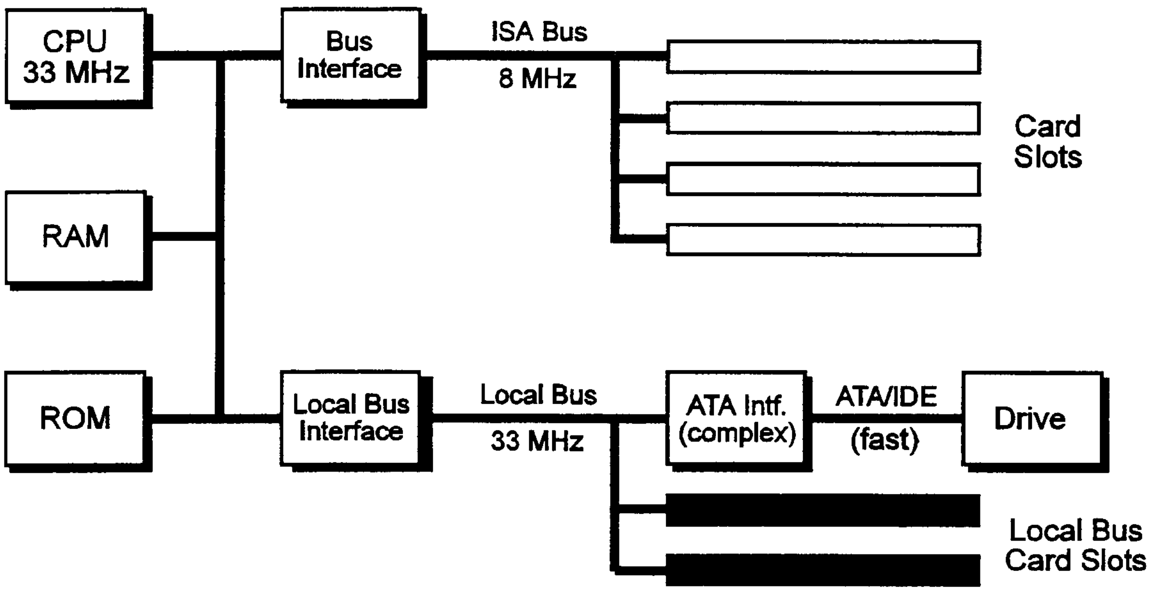

Figure A.2 — Modern PC architecture

This bottleneck inspired the invention of new, faster interfaces to the processor. Local busses are designed to run at the speed of the processor bus. Two local bus standards have emerged, the VESA Local Bus (VLB) and the Peripheral Component Interconnect bus (PCI). These local busses have the potential for faster data transfer from the disk drive to the processor.

To allow ATA disk drives to transfer data faster, the ATA standard had to be updated to allow faster transfer rates. These enhanced modes are still not as fast as a processor bus. To synchronize the data flow between 32-bit 33 MHz processor busses and slower 16-bit disk drives, a VLSI chip is required. Most of these bridge chips were implemented with fast CMOS processes to achieve the required bus speeds. As a result the edge rates on the ATA bus were often 1 to 2 ns, and sometimes less. These fast edges have aggravated the ringing on the bus to the point that system/drive combinations fail to work.

In summary, the reason for signal integrity problems appearing now, when they were absent before, is the advent of faster transfer rates on the ATA bus coupled with a change of IC process at the interface. The problems are not insurmountable, and in time it is likely that the ATA bus will be known as a robust, fast, and inexpensive interface.

Appendix B — Example of Drive-End Termination Timing Calculation

Let's assume that we have chosen to insert 82 ohm series resistors on all receive signals and 22 ohm series resistors on all transmit and bidirectional signals. We will also assume that the input capacitance of the interface chip is 10 pF. There are three different RC configurations that occur (Figure B.1). All receive signals will see an 82 ohm and 10 pF network. The data lines will see a 22 ohm and 10 pF network when the drive is receiving data. Signals driven back to the host (including data lines during a read) will see 22 ohms and 50 pF. The 50 pF assumption is the worst-case condition of both the host and another drive being located nearby (negligible cable length), and both of them having the maximum allowed input capacitance.

This careful thought process must be repeated for all fifteen of the ATA PIO timing parameters (and for DMA too). The easiest way to do this is to make a spreadsheet and enter the six values for RC delay and the interface chip timing parameters. Spreadsheet formulas can then compute the timing at the pins of the drive and highlight any that are not within spec. In this manner the difficult calculations need only be derived once and it becomes easier for others to check your work.

Appendix C — Example of Host-End Termination Timing Calculation

The host-end timing calculations are similar to the drive-end calculations described above with a few more complicating factors added in. The four different signal configurations are shown in Figure C.1. For this design we have decided to use 82 ohm series resistors on control lines received by the host and 22 ohm resistors on the data lines and control lines driven by the host. We will presume that the host adapter chip input capacitance plus stray capacitance is 15 pF.

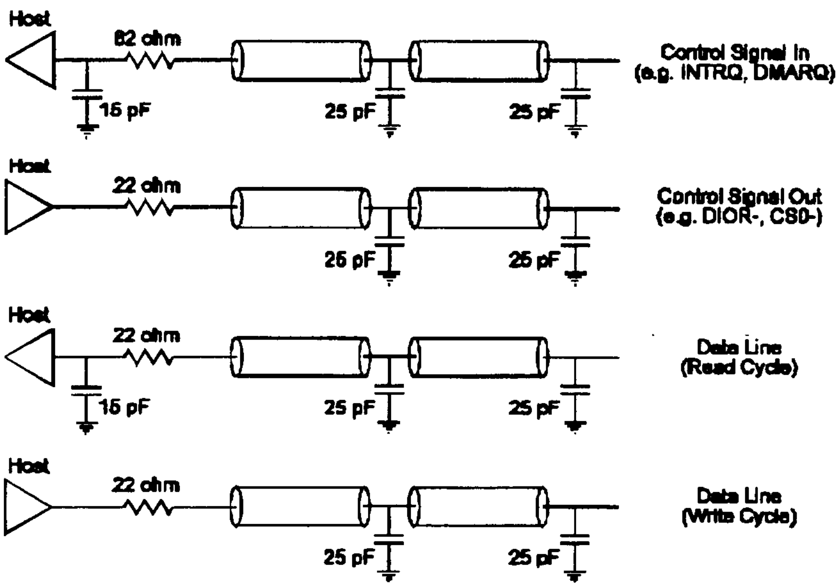

Figure C.1 — Host-end signal configurations with terminations

For these values the control signal out and the data out models look the same. We can use one simulation to determine both values. The greatest uncertainty is the delay through the cable for received signals. The total cable delay depends on the source impedance of the drive. This can be anything from zero to 82 ohms; the greater the impedance, the greater the delay. We will assume for this example that our drive vendor has read this document and has decided to use 22 ohms resistors. If we desire later that we want to make a worst-case assumption of 82 ohms, then we can add approximately 2 ns to our numbers.

Using SPICE models similar to the one shown in Figure C.2 we can derive the eight delay parameters required. A second drive appears in the model as a lumped capacitance of 25 pF which causes the maximum delay. The resulting values appear in Table C.3. By examining the values in the table it should be clear why the cable propagation delay is often referred to as being about 5 ns.

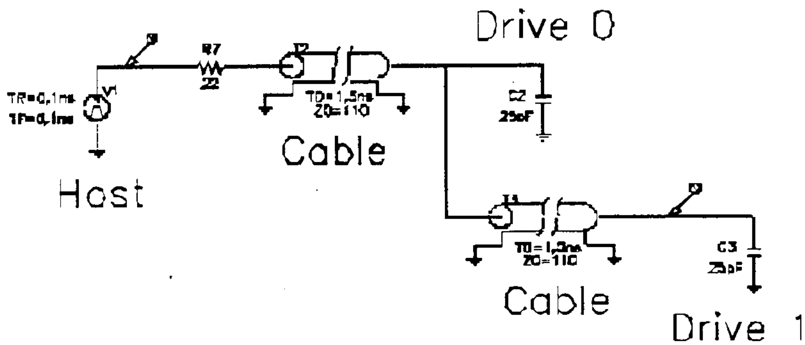

Figure C.2 — SPICE model for control signal out delay calculation

Symbol	Description	Value
Tphlci	Propagation delay, high to low, control in	6.4 ns
Tplhci	Propagation delay, low to high, control in	4.7 ns
Tphlco	Propagation delay, high to low, control out	5.9 ns
Tplhco	Propagation delay, low to high, control out	4.6 ns
Tphldi	Propagation delay, high to low, data in	5.7 ns
Tplhdi	Propagation delay, low to high, data in	4.3 ns
Tphldo	Propagation delay, high to low, data out	5.9 ns
Tplhdo	Propagation delay, low to high, data out	4.6 ns

Table C.3 — Typical host-end propagation delay times

The process of finding the ATA timing values is the same as for the drive-end example. The propagation delay times are added to and subtracted from the host adapter chip timings to obtain the timings at the input to the drive.[20] The resulting timing values can be compared against the ATA values to determine what mode the drive can operate at.

[20] In this case we are calculating the timing at the drive furthest from the host adapter.